

Publication date:

30 Jul 2026

Author(s):

Liam Deane, Principal Analyst

Omdia Market Radar: Cloud Platforms for Games – 2026

Table of Contents:

Summary.....	02
Omdia view	04
Recommendations	05
Defining the cloud platform market in the games industry.....	06
Key capabilities and vendor landscape	08
Vendor analysis	12
Appendix	30

Summary

Catalyst


The cloud platform critically underpins the games industry. Nearly every game released today is built, shipped, and kept running on cloud infrastructure. It hosts the servers behind online and multiplayer experiences, the build pipelines and version control systems through which games are assembled, the backend services that manage player identity and progression, the data and analytics systems on which commercial decisions rest, and the AI tools reshaping how all of this work gets done. Choosing among these platforms, and deciding how much of the stack to entrust to them, is now among the most consequential technology decisions a games company makes.

This Market Radar benchmarks nine cloud platforms against the requirements of the games industry. This is a pivotal moment for the sector, with competition growing every year. AWS has eliminated bandwidth charges on its managed game servers, and Microsoft has made PlayFab's core backend free for Xbox developers. More fundamentally, AI has become central to the value proposition. All major platforms now pitch their games offering through the lens of foundation models, generative pipelines, and agents.

This report navigates these trends, tracking the key developments in 2026 and the evolving capabilities of all of the major vendors in the market. These insights are critical both for games companies weighing where to place their infrastructure and platform bets, and for technology vendors whose products must find a place within an ever more cloud-centric games tech stack.

Market snapshot

Figure 1 summarizes the outcome of Omdia's assessment across the nine platforms and eight capability categories evaluated this year. The framework spans the full range of functions for which games companies engage cloud platforms—from hosting and multiplayer services through LiveOps, data, and AI—though each rating ultimately rests on the interplay between a platform's underlying infrastructure and the degree to which it has been productized for games.

Figure 1: Omdia heatmap for cloud platforms for games, 2026


© 2026 Omdia

Source: Omdia

Key messages

- Competition in the market in 2026 is increasing on two major axes. On economics, cloud platforms are making major changes, including free bandwidth, free backend tiers, and no-egress pricing aimed squarely at managing cost concerns, which developers consistently cite as a key challenge. On the other, they are competing on AI, rebuilding their games propositions around models, generative pipelines, and agents.
- AWS continues to lead the cloud services market for the games industry. Data from Omdia's Game Developer Collective surveys puts studio usage of AWS at 57%, comfortably the highest of any platform. AWS has also made one of the most significant commercial moves in the market: as of June 2026, network bandwidth is included at no additional charge for Amazon GameLift Servers instances, directly targeting the value proposition of bare metal specialists and lower-cost rivals.
- In terms of market share, Microsoft Azure and Google Cloud continue to contest second place, with some signs that the latter may be opening up a lead. Azure's games proposition has been enhanced by PlayFab Foundation Mode, which is set to make PlayFab's core services free for Xbox-committed developers, though the deep 2026 restructuring of Microsoft's Xbox business creates uncertainty around the company's broader games commitment.
- Google Cloud has assembled a particularly strong AI-led proposition for the games market. Its reorganization of the AI stack around agentic workflows, combined with custom silicon, frontier first-party models, IP indemnification, and named production deployments positions it strongly as AI moves from innovation and experimentation to the core operations of game studios.

- The Chinese platforms are using AI as a lever to gain greater global market traction. Tencent Cloud's HY 3D engine is being integrated into Western creative software, including Maxon's Cinema 4D, and Alibaba Cloud's Qwen family has become the most-downloaded open-source model lineup globally. AI is proving a more effective wedge into Western software ecosystems than these platforms' traditional infrastructure propositions ever were.
- Beyond the leading integrated platforms, the specialist and mid-tier field is being reshaped. OVHcloud joins this year's Market Radar as Europe's largest cloud provider with an explicitly gaming-branded product line; i3D.net has shifted its roadmap towards inference-optimized AI infrastructure; Oracle Cloud Infrastructure has transformed into an AI-infrastructure giant whose attention now sits further from the games market; and servers.com was rebranded as "Servers.com by Nexcess" as parent company CloudOne digital brings its brands closer together.

Omdia view

Games companies will spend \$16.3bn on operations technology in 2026, according to Omdia's *Games Tech Market Forecast*. Competition for this growing market has been heating up for several years and in 2026 has grown especially intense. This year, the numerous players are going on the offensive on cost. AWS's decision to include bandwidth at no charge with its managed game servers removed the least predictable line item in multiplayer hosting; Microsoft made PlayFab's core backend free for any developer shipping on Xbox; and OVHcloud and Oracle pressed their advantage on egress economics.

Notably, however, while buyer leverage has never been higher, cloud platforms remain remarkably sticky: in Omdia's most recent *Game Developer Collective* survey, 63% of developers said they were unlikely to change cloud platforms in the next six months, and cost remains the leading reason cited by the minority who would. The practical battleground is therefore for share of the multicloud strategies which more developers than ever are employing.

For the platforms themselves, revenue accrues to those who productize for the games industry. The games vertical is small relative to hyperscaler revenue, but it is disproportionately demanding—bursty, latency-critical, globally distributed, and operationally unforgiving—and those demands create defensible positions for vendors willing to invest in games-specific products. The platforms that shipped games-native services in 2026 consolidated or improved their standing, while those that drifted towards generic infrastructure saw their games relevance diminishing.

Finally, AI disruption has arrived. All major platforms have reorganized their games proposition around AI over the past year, from foundation-model access and generative content pipelines to agentic development tooling and in-game inference. The demand side tells a more complicated story, however. Omdia's *Game Developer Collective* data shows AI adoption among developers plateauing rather than accelerating, with doubts about the ultimate impact and efficacy of the technology remaining widespread in the industry even as AI capabilities continue to increase. This means that the platforms are, in effect, building ahead of demand. Eventually, this gap will close. The technical and economic logic of cloud-based AI in games is compelling, and production deployments in QA, moderation, and analytics are already substantial. But the pace at which it closes will be set by developers' trust and the economics of running AI at gameplay scale. For emerging players this

remains a moment of high disruption potential. For incumbents, meanwhile, the work of maintaining market leadership has never been harder.

Recommendations

Recommendations for game developers

- Broaden the criteria on which cloud platforms are judged. Raw infrastructure price-performance is now all but taken for granted. The most meaningful differences lie in managed games services, data and AI governance, egress and bandwidth economics, sovereignty options, and the quality of games-specific support. A selection process that scores only compute and storage pricing will miss key elements of where the platforms diverge.
- Re-run the numbers. The pricing moves of the past year—free bandwidth on Amazon GameLift Servers, PlayFab Foundation Mode, the spread of no-egress commercial models—have rendered cost comparisons made even twelve months ago out of date, particularly for multiplayer-heavy architectures where bandwidth and backend fees dominate. Studios carrying assumptions from their last procurement cycle are likely leaving money on the table.
- Focus AI investment on production needs. AI is delivering real value in development tooling, QA, moderation, and analytics, while player-facing AI remains largely pre-commercial. Platforms should accordingly be evaluated on the maturity, cost profile, and governance of those specific capabilities rather than on the breadth of their model catalogs.
- Revisit the backend build-versus-buy decision. The economics of backend services are changing. Cloud platforms are increasingly giving core services away to reinforce their ecosystems, while independent backend-as-a-service (BaaS) providers compete on neutrality and depth. The trade-off between integration convenience and platform lock-in deserves fresh scrutiny in light of this year's commercial moves.

Recommendations for cloud vendors

- The games opportunity justifies sustained, games-specific investment. Operations spending is growing at more than 10% annually, and this year's results demonstrate that games-native products—managed game servers, backend services, moderation, launch support—are what converts that spending into platform revenue. Vendors that treat games as an undifferentiated source of compute demand will find themselves competing solely on price, just as price competition has intensified.
- Compete on predictability. Cost is the leading reason developers cite for considering a platform change, but the underlying grievance is as often unpredictability (egress exposure, opaque tiering, surprise bills at launch) as absolute price. Flat-rate models, included bandwidth, and transparent launch-scaling economics have proven commercially potent this year, and they are available to vendors of all sizes.
- Close the gap between AI investment and AI adoption. Developer usage of generative AI has plateaued, and skepticism about its impact on quality is widespread. The vendors best

positioned to monetize AI in games will be those that productize specific, high-value workflows with clear economics and effective governance, which takes precedence over sheer breadth of model catalogs. For larger publishers, data-use guarantees and indemnification are becoming decisive selection factors.

Recommendations for other tech vendors


- Interoperability with the major cloud platforms is essential. Games companies assemble their tech stacks around one or more of the platforms, and tools that integrate cleanly enjoy a major advantage over those requiring custom integration work. As multicloud practice spreads, breadth of platform coverage matters as much as depth on any single platform.
- Backend and LiveOps specialists face a direct challenge and should choose their response carefully. With PlayFab's core services free for Xbox-committed developers and AWS strengthening the pull of its managed services offering, independent BaaS providers can no longer compete on price alone. Multi-platform neutrality, deeper specialization, superior tooling, and partnership with the infrastructure specialists are all viable strategies.
- The infrastructure-first platforms represent an underexploited partner opportunity. Several vendors in this assessment—Oracle, OVHcloud, i3D.net, servers.com—largely stop at infrastructure, leaving orchestration, backend, and LiveOps to be supplied by others. For middleware and services vendors, these platforms offer distribution with minimal competitive tension.

Defining the cloud platform market in the games industry

Definition and characteristics

Omdia's games technology taxonomy divides vendor revenue across five areas: development, operations, monetization, security, and non-core technology. Cloud platforms touch all five, but their center of gravity is operations—the technology through which live games are hosted, scaled, and managed. Within operations, Omdia tracks six sub-segments: game servers, backend services, multiplayer services, LiveOps, analytics, and identity management.

Operations is both the largest and one of the fastest-growing areas of games technology spending. Games companies will invest \$16.3bn in 2026, and \$24.8bn by 2030. Game servers are the primary driver of this spending, at \$7.9bn in 2026, or 49% of the total; backend services contribute a further \$2.5bn, or 15%. The remaining \$5.8bn was spread across analytics, LiveOps, multiplayer services, and identity.

Figure 2: Operations games tech revenue, 2021–30 (\$bn)



© 2026 Omdia

Source: Games Tech Market Forecast – 2026, Omdia

The sub-segments differ markedly in structure. Game servers and backend services are dominated by the platforms assessed in this report, whether as hyperscaler managed services or specialist bare metal. Multiplayer services will reach \$1.2bn in annual spending in 2026, but the cloud platforms share this segment with a crowded field that includes both game engine providers (Unity and Epic Games each field comprehensive multiplayer suites) and independent specialists. LiveOps, at \$1.3bn, is even more contested. The systems for managing content deployment, player engagement, virtual economies, and progression are supplied variously by cloud platforms, BaaS vendors, and point-solution specialists.

Analytics occupies a distinctive position in the market. At \$3bn in 2026, it is the largest operations segment outside servers. Even this figure understates the segment's significance, however. A games company's analytics infrastructure determines where its player data lives, and where the data lives increasingly determines where its AI workloads run. The platforms with the strongest data propositions, meanwhile, are converting analytics incumbency into AI adoption.

Looking at AI more generally, a striking feature of AI inference is the extent to which it shares a demand profile with game servers characterized by low latency, real-time state, bursty concurrency, and global distribution. Consequently, the platforms that master one are strongly positioned in the other. There are grounds for caution, however. Omdia's *AI in Games: State of the Market 2026* report found that while AI adoption surged rapidly in 2024 and early 2025, it has since plateaued below that peak rather than continuing to climb. Usage remains concentrated in coding assistance—cited by two-thirds of AI-using developers—with 2D asset creation and localization the next most common applications. Nearly all of this activity, whatever its pace, must run on cloud infrastructure given that neither model training nor production inference is practical on local hardware at any meaningful scale.

Figure 3: Cloud platforms currently in use by games studios, 2023–26


© 2026 Omdia

Source: Game Developer Collective Survey Database, Omdia

Measuring the platforms' relative positions in games is complicated by disclosure as none of the major vendors report games revenue separately. Omdia's *Game Developer Collective* surveys provide a good baseline, measuring the share of studios using each platform since 2023. On this measure AWS's leadership is long-standing and stable, with 57% of studios that use a cloud platform working with AWS as of 2026. The next tier has seen Google Cloud and Microsoft Azure trade places in second and third position over time, with some indications that Google has begun to pull in front in recent surveys. Platforms outside the big three held a combined share of around 23%, a figure that has trended gradually upward as specialist and regional providers gain visibility.

Two findings from the same surveys provide further insight into the competitive dynamics. First, multicloud is normal because studios, especially larger ones, spread workloads across providers for redundancy, regional reach, or workload fit. Second, switching is rare: 63% of developers say they are unlikely to change cloud platforms in the next six months, with cost the leading factor among the few inclined to move. Consequently, vendors are not, for the most part, trying to convert each other's customers wholesale. Rather, they are competing for marginal workloads within tech stacks that already span multiple providers.

Key capabilities and vendor landscape

The competitive field assessed in this report is characterized by a longstanding pattern with a single dominant leader, a trio of close rivals, and a challenger tier in flux. AWS's leadership compounds year over year as an ever-growing cohort of engineers familiar with the platform make it the default selection. Google Cloud, Microsoft Azure, and Tencent Cloud each counter with highly differentiated propositions—data and AI, backend services and ecosystem economics, and games-native tooling at scale, respectively. Beneath them, the challenger tier is continuing to evolve. Alibaba Cloud has

closed much of the capability gap to the leaders, Oracle has pivoted towards AI infrastructure, i3D.net has extended into AI inference, and OVHcloud has entered the assessment for the first time.

To compare platforms of such different scales and strategies on equal terms, Omdia assessed each against eight capability categories, summarized in **Figure 4**.

Figure 4: Cloud platforms for games Market Radar capabilities, 2026

Game development	How well does the platform support cloud-based game development?
Game servers	How capable are the platform's multiplayer and backend game servers?
AI and machine learning	To what extent can the platform support AI/ML-related applications?
Multiplayer services	To what extent can the platform offer tools to support multiplayer games?
LiveOps	How extensive are the platform's solutions for supporting live operations?
Database and analytics	How capable are the platform's data infrastructure and tools?
Security and compliance	How comprehensive is the platform's security and compliance offering?
Partners and integrations	How extensive and how well integrated is the platform's partner ecosystem?

© 2026 Omdia

Source: Omdia


Expanding on **Figure 4**, these key capabilities are defined for the purposes of this report as follows. The framework is organized around what games companies buy, and it has evolved with the market this year in two respects: security and compliance has been added as a standalone category, and AI and machine learning carries greater weight in the aggregate assessment than in prior editions.

- **Game development.** This category assesses how well a platform hosts the making of games. Considerations include cloud workstations and virtual machines suitable for content creation, build automation and CI/CD pipelines, version control hosting, asset storage, and the protection of source code and assets, together with the depth of integration with game engines and other core development tools and the degree to which these services have been adapted to games production workflows.
- **Game servers.** The platform's capacity to host multiplayer sessions and core game logic reliably, at low latency, and at unpredictable scale. Managed offerings that absorb orchestration and fleet management on the developer's behalf are weighted above raw infrastructure, and the assessment also considers DDoS protection, support for cross-platform play, geographic reach, and the ability to absorb the demand spikes that accompany launches and live events.
- **AI and machine learning.** The ability to support AI and ML applications has become a key criterion for cloud platforms in the games market. This category takes into account the platform's ability to support games customers at all levels of the AI tech stack, including custom model training and inference, provision of adaptable foundation models, generative content pipelines, AI governance features, and the ability to offer customized AI-powered solutions for game-related applications.
- **Multiplayer services.** The services and tooling supporting online gaming. This includes matchmaking, leaderboards, session management, and player communication through text and voice alongside anti-cheat capability. Platforms are assessed on the completeness of these services, and on whether they are delivered as games-native products or must be assembled from general-purpose components.
- **LiveOps.** LiveOps encompasses tools and services that support the ongoing operation of live games, beyond multiplayer functionality. This includes player authentication, content management and deployment and tools for managing in-game economies, player progression, and other systems. Platforms are rated on their ability to offer flexible, powerful tools that developers can customize, as well as user-friendly, ready-to-use solutions.
- **Database and analytics.** This category examines the infrastructure and tools available for data management and analysis. Platforms are assessed on their ability to support tasks such as A/B testing, player behavior analysis, and ML applications.
- **Security and compliance.** This category covers the security and regulatory foundations on which games companies depend, beyond the DDoS protection assessed under game servers. Key factors include identity and access management, encryption of data at rest and in transit, data residency and sovereignty controls, compliance with data protection regulations such as GDPR, and the breadth of security certifications and attestations held. Anti-cheat and content moderation continue to be assessed under multiplayer services; this category addresses platform- and data-level security.
- **Partners and integrations.** Every platform relies on a network of partners and complementary solutions, and this category assesses how well each works with the ecosystem around it. Particularly important are the breadth and quality of games-focused partners, the depth with which partner products integrate into the platform's own services,

and the coherence of each vendor’s strategy for balancing first-party product against third-party reliance in a market where interoperability failures carry high operational cost.

No single platform leads in all eight categories, and the ratings should be read as a map of fit rather than a single ranking. For instance, the right choice for a mobile studio consolidating its analytics infrastructure differs from the right choice for an AAA publisher provisioning launch capacity for a global multiplayer title. The heatmap in **Figure 5** presents the full assessment; the vendor analysis that follows sets out the reasoning behind each rating.

Figure 5: Omdia heatmap for cloud platforms for games, 2026



© 2026 Omdia

Source: Omdia

The heatmap is colored as follows:

- **Advanced capability:** The vendor demonstrates very strong capabilities and/or capabilities in alignment with what Omdia explored as part of this research.
- **Broad capability:** The vendor offers better-than-expected capabilities that are well-suited to the needs of most businesses.
- **Partial capability:** The vendor provides expected capability but lacks some of the advanced capabilities assessed as part of this research.
- **Functional capability:** The vendor provides limited or none of the expected capabilities explored as part of this category.

The categorization of each vendor is as follows:


- **Market leader:** This category represents the leading solutions that provide advanced capabilities across a wide range of areas explored and which we believe are worthy of a place on most technology selection shortlists.

- **Market challenger:** The solutions in this category offer some advanced or broad capabilities, have appropriate functionality across other areas, and should be considered as part of a technology selection process.

Vendor analysis

AWS

Figure 6: Omdia Market Radar recommendation – AWS



© 2026 Omdia

Source: Omdia

Why consider AWS?

- AWS's infrastructure spans 39 regions and 123 availability zones, extended by more than 30 Local Zone metropolitan areas and over 750 CloudFront points of presence, a footprint that keeps latency low and capacity deep for even the most demanding global gaming workloads. Recent infrastructure investments include the AWS European Sovereign Cloud, launched in January 2026, a physically and logically separate EU-only environment that gives European publishers an additional option for meeting data sovereignty requirements.
- Two decades in the games market have given AWS an incumbency competitors cannot easily replicate. The company claims more than 20 years supporting the industry's biggest titles—*Fortnite*, *Roblox*, and *League of Legends* among them—and states that over 750m gamers each month play games built on AWS. That history means that the industry has countless engineers who already know the platform, documentation and community knowledge accumulated over years, and there are proven reference architectures for nearly every games use case.
- Amazon is also a major player in the games market beyond just AWS. It owns Twitch, the leading game streaming platform; operates Amazon Luna, its cloud gaming service now bundled with Prime membership and included in Prime Video; and develops and publishes titles through Amazon Games. This first-party footprint keeps AWS close to the operational realities of the industry, and feeds directly into the design of its developer-facing services.

Focus and roadmap

- Game server capabilities remain the cornerstone of AWS's offering. This year has seen a major change to the pricing of this core capability, with network bandwidth now included at no additional charge for all Amazon GameLift Servers instances of generation 6 or newer, in

all regions except China. Bandwidth has historically been the least predictable line item in multiplayer hosting costs, making this a major change to one of the industry's biggest tech cost drivers. GameLift Servers continues to support scaling to 100m concurrent players per game, with Queues optimizing player routing, AWS Global Accelerator reducing latency and jitter, and GameLift Servers Anywhere enabling hybrid deployments spanning cloud, on-premises, and third-party infrastructure. New additions include Player Gateway, a fully managed relay-based DDoS protection layer for UDP game traffic, and full support for containerized game server fleets.

- Cloud-based game development and LiveOps are expanding areas of focus. Amazon GameLift Streams, launched in March 2025, streams game builds at up to 1080p/60fps to any browser and is being positioned for playable demos, instant-play advertising, and QA workflows, with Bandai Namco, Jackbox Games, Ludeo, and Xsolla among early adopters. The open-source Cloud Game Development Toolkit provides Terraform modules for build pipelines, Perforce deployment, and Unreal Engine Horde. On the LiveOps side, Amazon Cognito handles player identity, including federation with Steam, Xbox Live, and PlayStation Network; AWS End User Messaging, SNS, and AppSync power engagement channels; and Amazon Personalize enables targeted offers.
- AI infrastructure and model access are further fast-moving areas within AWS's portfolio. Amazon Bedrock now provides managed access to over 200 foundation models from more than 18 providers, including models from OpenAI, Anthropic, Meta, Mistral, and Amazon's own Nova family. This multi-model breadth is paired with an exceptionally broad accelerator fleet including Nvidia GPUs through to Blackwell-generation B200/B300 instances, plus third-generation custom Trainium silicon and Inferentia inference chips positioned on price-performance. Amazon SageMaker, with HyperPod for resilient distributed training, covers the full model lifecycle for studios building proprietary models.
- Games-specific AI applications increasingly feature productized use cases. Nova Sonic, a speech-to-speech foundation model on Bedrock, targets real-time voice-driven NPC conversation and dynamic narration, and AWS publishes deployable solution guidance for NPC dialog, AI content moderation, and AI-assisted QA, with SageMaker supporting distillation of small language models for on-device or on-server inference. Generative content pipelines span image, video, voice, and (via SageMaker-hosted models) 3D generation, with partners including Scenario, ElevenLabs, and Tripo. Cost and latency controls address the economics of running AI at gameplay scale, a major practical barrier for live-service adoption.
- AI developer tooling has been rebuilt around agents. Kiro, an agentic IDE launched to general availability in November 2025, supports spec-driven development, customization against proprietary C++/C# game codebases, and MCP-based integration with game engines, Perforce, and AWS infrastructure. Amazon Quick extends agentic assistance to non-engineering workflows including natural-language game analytics. Bedrock AgentCore provides the managed runtime for studios deploying their own AI agents, and AWS's roadmap points clearly towards agent-operated development and LiveOps workflows.

Market impact


- AWS enters this year's assessment, as it has every prior edition, as the reference point against which the rest of the market is measured. Omdia's *Game Developer Collective*

surveys found 57% of studios using a cloud platform were using AWS as of the first half of 2026, the highest of any provider, and its publicized customer roster—Riot Games, Epic Games, EA, Scopely, Behaviour Interactive, and Netflix Games among them—is generally understood to extend across most of the industry’s largest operators.

- The platform’s position rests on maturity and dependability as the industry’s default choice for over a decade, as well as scale and reach that provide confidence in reliability, together with a comprehensive suite of tools, partnerships, and integrations. Its partner network of over 300 games-focused vendors—including backend providers Pragma, Heroic Labs, AccelByte, and Beamable, and data partners Databricks and GameAnalytics—adds significant value to core offerings.
- The free bandwidth decision significantly strengthens AWS’s competitive position, helping to address its most persistent criticism in this market. It should be noted, however, that the benefit is limited to GameLift Servers instances rather than self-managed EC2 fleets, so studios running their own orchestration continue to pay standard egress rates.
- Weaknesses are comparatively few but real. The sheer sprawl of the AWS portfolio imposes a navigation cost, particularly on smaller studios without dedicated cloud teams. Mainland China, meanwhile, continues to be served only through partner-operated regions and sits outside the bandwidth offer.

Google Cloud

Figure 7: Omdia Market Radar recommendation – Google Cloud



© 2026 Omdia

Source: Omdia

Why consider Google Cloud?

- Google Cloud ranks third among cloud platforms globally, behind AWS and Microsoft Azure, but its infrastructure bears comparison with either, stretching across 43 regions and 130 zones linked by Google’s private fiber backbone and more than 100 network points of presence worldwide, providing the scalable, low-latency foundation games workloads demand in nearly every major market. The significant exception remains mainland China, where Google Cloud has no presence.
- The defining development of the last year is the reorganization of Google Cloud’s AI offering around agentic workflows. At Google Cloud Next 2026, the company launched the Gemini Enterprise Agent Platform, the evolution of Vertex AI, combining its model selection, model building, and agent building capabilities with new features for agent integration, DevOps, orchestration, and security. For games companies, this consolidates the AI stack into a

single platform from foundation model access through agent deployment, governance, and observability, with direct SDK integrations for Unity and Unreal to bring AI agents into the developer workflow.

- The platform's commitment to open source remains a key differentiator. Agones (game server orchestration) and Open Match (matchmaking) integrate with the major engines and run on Google Kubernetes Engine (GKE), minimizing vendor lock-in, while a partner ecosystem of BaaS providers (Heroic Labs, AccelByte, Diarkis) offers prebuilt backend services for studios that prefer managed solutions over custom builds.
- Google Cloud's data and analytics capabilities remain among the most advanced in the industry. BigQuery is widely regarded as the de facto standard for games analytics among the largest publishers, complemented by Cloud Spanner for globally consistent player data, AlloyDB, Memorystore, and Bigtable for specialized workloads, and Firebase for mobile-first analytics, authentication, and A/B testing. Increasingly, these data assets feed directly into AI workloads with use cases including churn prediction, personalization, and AI-driven anti-cheat models built on BigQuery telemetry.
- Enterprise-grade AI governance is emerging as a differentiator for risk-conscious publishers. Google Cloud offers dual copyright indemnification covering both training data and generated output; SynthID watermarking for content provenance; and Model Armor guardrails against harmful content and prompt injection, addressing IP and safety concerns that remain barriers to production AI adoption in games.

Focus and roadmap

- Google Cloud's core games infrastructure remains built on GKE and Agones for managed game server hosting and orchestration, with Bare Metal Solution for specialized workloads and Google Distributed Cloud for hybrid, multicloud, and edge deployment. Four-way autoscaling on GKE handles the bursty load patterns characteristic of live service games, while Cloud Spanner scales the database tier automatically with a 99.999% availability SLA.
- AI is now unambiguously the center of the roadmap. The Gemini Enterprise Agent Platform provides tools organized around four pillars (build, scale, govern, and optimize) including the Agent Development Kit, Agent Studio, Agent Registry, Agent Identity, and Agent Gateway, with access to more than 200 models through Model Garden, including Gemini 3.1 Pro alongside third-party models such as Anthropic's Claude family and open-source models via Hugging Face integration. Real-time, in-game inference is served from GKE co-located with game servers, using efficiency-optimized models for latency-sensitive use cases such as AI NPCs and in-game assistants.
- Generative content pipelines have matured into a full first-party suite including Nano Banana 2 (Gemini 3.1 Flash Image) for 2D asset and texture generation, Veo for cinematic video, and Lyria for music and audio, alongside partner integrations (Atlas for agentic assistance in Unreal/Unity, Dreamlands for text-to-3D environments, ComfyUI and Fal for custom workflows). On hardware, Google's custom TPU silicon alongside Nvidia Blackwell-generation GPU instances gives it unusual flexibility to optimize price-performance for AI workloads—a key advantage over pure GPU resellers.


- Security has been significantly reinforced. Google completed its \$32bn acquisition of cloud and AI security platform Wiz in March 2026, combining it with Mandiant threat intelligence and the BeyondCorp zero-trust architecture. For games companies, the pitch is protection of the AI development lifecycle itself—player databases, models, and IP—at a time when both attack surfaces and AI-related IP exposure are growing.
- Google Cloud’s strategic vision remains “the ecosystem for living games” but the emphasis has shifted from AI as a feature to agents as the central focus. The company identifies AI agents as the most significant innovation area in games, spanning production efficiencies (automated QA, asset generation) and player-facing experiences (dynamic, unscripted NPCs).

Market impact

- Google Cloud states that its platform is used by all ten of the top ten games companies globally by revenue, and Omdia’s *Game Developer Collective* continues to place its games industry share among the leaders. Google Cloud now consistently places as the second-most used cloud in the games industry, giving it momentum in a market where AWS’s incumbency has historically been the default.
- The strongest evidence of AI traction comes from named customer deployments. Capcom has built AI agents on the Gemini Enterprise Agent Platform (in combination with proprietary technology) that run more than 30,000 hours of autonomous playtesting per month, using visual inspection, predictive, and institutional-knowledge agents. Notably, Capcom maintains that AI-generated assets will not appear in its game content, illustrating that even flagship adopters currently confine AI to production tooling rather than shipped creative content—an important qualifier on the pace of AI adoption.
- Player-facing AI remains largely experimental. Ubisoft’s “Teammates” prototype demonstrates voice-driven AI companions and an agentic in-game assistant, but remains an R&D experiment rather than a shipped title. Other notable deployments include Nexon’s AI-powered harmful-image detection for trust and safety, and Klang Games’ agent-driven MMO simulation.
- Google Cloud’s data and analytics tools are widely regarded as industry-leading, and Firebase continues to anchor its strength in mobile, where authentication, messaging, remote config, and A/B testing are deeply embedded in Android-centric development workflows. These established strengths now double as the on-ramp for AI adoption, since studios’ player telemetry already resides in BigQuery.
- Despite its absence from mainland China and a smaller overall share of the cloud computing market than AWS and Azure, Google Cloud has a highly coherent AI-led proposition. Its combination of custom silicon, frontier first-party models, an integrated agent platform, IP indemnification, and named games-industry production deployments positions it strongly as AI moves from the innovation agenda to the operational core of game development and LiveOps.

Tencent Cloud

Figure 8: Omdia Market Radar recommendation – Tencent Cloud



© 2026 Omdia

Source: Omdia

Why consider Tencent Cloud?

- No cloud platform is more deeply embedded in the Chinese and broader Asian games market than Tencent Cloud, and that position is increasingly a springboard to global expansion. Its footprint has grown to 23 geographic regions and 66 availability zones (up from 21 and 56 in 2025), with more than 1m servers deployed and over 3,000 acceleration nodes globally. A third availability zone in Frankfurt, operational from mid-2026, further expands its European infrastructure.
- Tencent is by some measures the world's largest games company, and its cloud platform is built on infrastructure originally developed for its own titles, a differentiation that general-purpose clouds cannot replicate. The games industry accounts for an unusually high share of Tencent Cloud's revenue, ensuring the vertical's claim on sustained corporate attention and investment. Games-industry tooling validated on titles such as *Honor of Kings* and *PUBG Mobile* is available to external customers.
- AI has become the organizing principle of Tencent Cloud's games proposition. In early 2026, the company introduced a wave of AI-powered solutions spanning how games are developed, played, and protected, and repositioned its flagship Game Multimedia Engine as the Game Multimedia Engagement solution (GME), an expansion into an integrated, AI-fueled suite spanning the game lifecycle.
- The platform pairs game-native AI applications with a native foundation model stack. The Hunyuan/HY model family covers text, image, video, 3D, and role-play modalities, while TokenHub provides unified, OpenAI- and Anthropic-compatible API access to both Hunyuan and major third-party Chinese models (including DeepSeek, Zhipu GLM, Kimi, and MiniMax). This is underpinned by TI-ONE for training and managed inference, Hyper Computing Cluster with RDMA networking for large-scale training, and the TACO-LLM inference acceleration engine.
- Tencent's expertise in China's live service game market, which is in many respects more advanced than in the West, continues to translate into a deep LiveOps offering. However, partner and software-ecosystem integrations in many cases remain stronger in the Chinese market than for Western equivalents, and advanced AI accelerator availability outside mainland China is comparatively limited, which may constrain international customers with heavy training workloads.

Focus and roadmap

- Generative AI content production is the clearest area of acceleration. Hunyuan 3D launched globally in November 2025, generating commercial-grade 3D assets from text, images, or sketches, with over 3m community downloads on Hugging Face and more than 150 enterprises in mainland China integrating the model via Tencent Cloud, including Unity China. It integrates with Unity, Unreal Engine, and Blender and sits within a broader AI production toolchain including the likes of GiiNEX (combined generative and decision AI), Superman, MagicDawn, and Miora. This connects generated assets into professional game pipelines covering rigging, skinning, animation, lighting, and asset management.
- Runtime and in-game AI is the second pillar. GME now offers AI-enhanced voice with built-in ASR, interruption handling, and directional pickup, acting as a high-performance voice gateway connecting developers to LLMs and TTS for gameplay such as “AI teammates”, alongside Magic Voice AI voice-changing and real-time translation. TRTC Conversational AI targets sub-300ms audio transport and sub-1,000ms dialog latency, the dedicated Hunyuan-role model supports character dialog and NPC interaction, and the Agent Development Platform (ADP) provides RAG, workflow, and multi-agent orchestration for virtual NPCs and studio knowledge assistants.
- AI also permeates operations, QA, and security. The Acorn AI Game Testing Agent (WeTest) supports natural-language-driven automated testing and defect analysis; ACE builds a full-scenario real-time anti-cheat system via AI-powered behavioral detection; and AI-assisted analysis has been embedded in observability (TCOP), log analytics (CLS), and FinOps tooling (AI cost analysis and budget optimization).
- Core game server infrastructure remains highly capable and continues to develop. Orchestration spans TKE and the PGOS dedicated-server lifecycle manager; EdgeOne combines over 3,000 edge nodes with integrated DDoS/CC/bot protection and claimed three-second attack detection-to-scrubbing; and the distributed cloud portfolio (CDC, TDCC, Tencent Edge Zone) covers hybrid, multicloud, and edge deployment.
- Data and analytics tooling remains a strength, anchored by TcaplusDB for game state data, the TCHouse warehouse family, WeData governance, and the GEM game experience management dashboards. The ABetterChoice A/B experimentation platform, used by *PUBG Mobile* and Ubisoft among others, supports parameter experiments and staged rollouts without repackaging.


Market impact

- Tencent Cloud’s games customer base has an increasingly international profile. Southeast Asia in particular accounts for a growing share of games customers, with North America, East Asia, and Europe each also generating substantial business. Southeast Asia has emerged as a stronghold of Tencent’s non-Chinese business.
- AI is proving to be Tencent Cloud’s most effective path of entry into Western software ecosystems. Maxon, for instance, announced a strategic partnership to integrate the Tencent HY 3D Global AI engine into Cinema 4D, and Germany’s 3D AI Studio and Lithuania’s CGTrader are building on the same engine.
- The infrastructure growth story continues with two additional regions and ten new availability zones added in the past year, European capacity expansion in Frankfurt, and

near-global coverage across North America, Latin America, Europe, and Asia & Oceania. Combined with competitive pricing and games-specific commercial programs (startup credits, WeChat Mini Game consumption rebates, an overseas expansion accelerator), this solidifies its position among the leaders in this year’s Market Radar.

Microsoft Azure

Figure 9: Omdia Market Radar recommendation – Microsoft Azure



© 2026 Omdia

Source: Omdia

Why consider Microsoft Azure?

- Microsoft Azure data center infrastructure is the match of any other provider, a footprint that Microsoft’s AI-driven buildout continues to extend. For gaming workloads, that reach translates directly into latency and resilience, with capacity close to players across every major games market.
- PlayFab, Microsoft’s managed backend platform, remains a key differentiator and has seen sustained investment over the past year. The PlayFab Unified SDK, included in the October 2025 GDK, brings together all PlayFab client libraries including Multiplayer, Party, Game Saves, and Services into one cohesive product with standardized authentication, async operations, error handling, and memory management, addressing a longstanding integration pain point. PlayFab continues to power major titles including *Minecraft*, *Sea of Thieves*, and *Helldivers 2*.
- Microsoft’s presence across the games industry remains unmatched in breadth. Xbox, Windows, and the Activision Blizzard King portfolio give it direct operational insight into game development and live operations at the largest scale. However, this presence is now a more complicated asset than in prior years, given the deep restructuring of the Xbox business in 2026.
- Hybrid and multicloud capabilities continue to be an Azure strength. Azure Arc enables unified management across cloud, on-premises, and edge environments, and Azure’s interoperability makes it a common component of multicloud strategies.
- Azure’s general-purpose AI capabilities, anchored by the OpenAI partnership and the Azure AI Foundry platform, are potent and increasingly relevant to games in development workflows. However, games-specific AI productization has narrowed rather than deepened over the past year, and Azure’s AI strength is more as a world-class general-purpose infrastructure rather than a games-native AI portfolio.

Focus and roadmap

- PlayFab's 2025–26 roadmap has centered on cross-platform progression. PlayFab Game Saves reached general availability in late 2025, building on Xbox Game Saves to deliver cross-platform progression with online/offline handling and conflict management, with costs covered by Xbox for titles in the Xbox ecosystem. This sits alongside a rebuilt Statistics and Leaderboards service and expanded identity support, including native Battle.net player logins.
- Another significant move is PlayFab Foundation Mode, which entered public preview in early 2026. Foundation Mode gives Xbox creators access to PlayFab's core game services across all platforms at no additional cost, with no Azure subscription required. This continues the transition of PlayFab from a standalone commercial product into a core element for the Xbox ecosystem, and undercuts competing BaaS providers on price for any developer shipping on Xbox.
- Azure offers modern data pipelines via PlayFab Data Connections, Azure Data Explorer, and Microsoft Fabric. The standalone Game Development Virtual Machine marketplace offering has also been retired, with Microsoft Dev Box now the primary cloud workstation route.
- AI investment continues but with a more disciplined scope. Muse and Azure AI Foundry are notable developments, while consumer-facing experiments have been pruned, as with the Gaming Copilot feature for consoles which was cut in 2026. The clearer through-line is Azure OpenAI and Foundry as general-purpose AI infrastructure that games companies adopt for content pipelines, player support, and analytics, rather than bespoke gaming AI products.

Market impact


- *Game Developer Collective* surveys have historically placed Azure as vying with Google Cloud for second place in terms of market share, though recent surveys have shown a slight slippage in Azure's market share. Nonetheless, its interoperability and hybrid capabilities keep it well-represented in multicloud arrangements and among third-party server orchestration providers.
- The turmoil in Microsoft's games business has deepened this year. The division has seen a leadership overhaul together with the elimination of more than 3,000 jobs and a plan to divest at least four development studios. This follows an earlier major round of cuts in July 2025.
- Importantly, however, this retrenchment is in Microsoft's games content business rather than its cloud business. Azure itself is healthy, buoyed by AI demand, and the Xbox reset arguably strengthens the case for Azure's games platform services. Foundation Mode and the stated pivot towards open development tools indicate that Microsoft's games ambitions are consolidating around services more than content. The risk for prospective customers is less about Azure's infrastructure commitment and more about the durability and prioritization of games-specific product teams amid repeated reorganizations.
- Azure's credibility in the games market has, however, been hit by the retirement of games-specific products (PlayFab Insights, the Game Development Virtual Machine, Gaming Copilot for consoles), the shift of more data-engineering responsibility onto customers, and the difficulty of assessing how games-specific product teams are being prioritized amid

repeated reorganizations. While Azure’s infrastructure and PlayFab’s recent delivery cadence remain strong, it is not clear that Microsoft is keeping pace with its rivals’ rate of games-specific investment across the wider capability set.

- PlayFab remains one of the most comprehensive backend platforms in the market, and Foundation Mode helps to shore up its competitive position against BaaS challengers, at least for developers shipping on Xbox. The negative press from the troubled *Helldivers 2* of 2024 has faded, and the platform’s recent delivery cadence has been improved.

Alibaba Cloud

Figure 10: Omdia Market Radar recommendation – Alibaba Cloud



© 2026 Omdia

Source: Omdia

Why consider Alibaba Cloud?

- Alibaba Cloud remains one of the leading cloud platforms for games companies operating in mainland China and its relevance increasingly extends well beyond China as well. The platform operates 105 availability zones across 32 regions, having opened new data centers in Japan, Malaysia, France, and Mexico as part of a \$53bn commitment to AI infrastructure, making it a viable option for developers across Asia and, increasingly, for global customers targeting the Chinese and broader Asian markets.
- The platform’s games-specific stack has matured considerably. GameServer Engine (GSE) provides fully managed fleet hosting and FlexMatch-compatible matchmaking, while OpenKruiseGame (OKG), Alibaba’s open-source Kubernetes framework for game servers, is in production at miHoYo, Lilith, FunPlus, and other major publishers, offering in-place server updates without dropping sessions. GameShield, its games-dedicated DDoS product, is the de facto standard among China’s top publishers, and its Edge Node Service spans more than 3,000 edge points of presence for latency-sensitive workloads.
- AI is increasingly central to Alibaba Cloud’s value proposition. The Qwen model family surpassed 2bn cumulative downloads on Hugging Face by mid-2026, overtaking Meta’s Llama as the most-downloaded open-source model family, and the company’s Bailian/Model Studio platform gives games companies access to the full Qwen lineup via a single API. Alibaba Cloud reports that 1,500 of its 4,000 games industry customers now consume its Model-as-a-Service (MaaS) offerings, an unusually high AI attach rate for the sector.
- For studios serving both mainland China and global audiences, Alibaba Cloud offers cross-border network capabilities that Western hyperscalers cannot match, including optimized

CN2/MPLS routing into China combined with its Global Accelerator backbone. This comes alongside compliance support spanning China's ISBN licensing regime, MLPS, and minor-protection rules as well as GDPR and regional certifications.

Focus and roadmap

- Alibaba Cloud has repositioned itself as an “AI-native” cloud, and its roadmap for games reflects this. Its Lingjun intelligent computing clusters support large-scale training on Nvidia H800/H20 and A100-class hardware, while proprietary inference engines and optimization techniques such as KV-cache reuse and quantization target the cost of running AI in production games.
- Generative AI content pipelines are a major focus. Qwen-Image and Qwen-Image-Edit support concept art and marketing creative; the Wan video model family generates cinematics and short-form promotional clips; CosyVoice and Qwen-TTS provide multilingual voice synthesis and cloning for NPC dialog; and 3D generation tools produce character and environment assets from text or image input.
- Alibaba is also building towards meaningful deployment of real-time, in-game AI inference. Model serving solution PAI-EAS provides managed elastic inference with sub-150ms latency for mid-sized models, edge inference runs across the ENS footprint for AI NPC and translation use cases, and the Bailian agent platform supplies pre-built templates for AI game masters, companion NPCs, and customer support. Though in-game use cases remain limited, Alibaba Cloud cites a publisher processing more than 5bn chat moderation calls per quarter using Qwen models, and customers with daily API volumes in the billions.
- Core infrastructure for games continues to advance in parallel. The company reports repeatedly handling launches above 10m concurrent players, with game-aware autoscaling (including player-count and room-load scalers in OKG), free million-CCU load testing under its Game Launch Acceleration Program, and war-room support for major launches. Its data stack provides excellent support for LiveOps and analytics.
- Notably, Alibaba Cloud argues that its LiveOps breadth—which now spans identity, payments integration, push, RTC, chat, moderation, matchmaking, and analytics as first-party services—reduces the third-party stitching required on rival platforms, and its open-source support (OKG, Qwen under Apache 2.0, Nacos, Dragonfly) mitigates lock-in concerns for studios wary of proprietary game services.

Market impact


- Alibaba Cloud is notably strong in China's MaaS and agents market, and projects per-client MaaS spending to grow between 100% and 1,000% in 2026. Its case studies suggest games companies are moving beyond experimentation into production AI workloads on the platform, with Qwen's open-source availability lowering the barrier to adoption.
- The customer base remains heavily concentrated in Asia & Oceania. Alibaba has seen growing multicloud wins against AWS, Azure, and Google Cloud, particularly on price-performance, and its June 2026 launch of a France region (its third European hub alongside Germany and the UK) signals intent to compete more seriously in Western markets.
- Nevertheless, challenges to broader adoption persist. Western publishers may weigh geopolitical and data governance considerations in platform selection, while the partner

and ISV ecosystem remains oriented towards China and the broader region. Meanwhile, the most advanced AI hardware (Lingjun clusters) is concentrated in Chinese and Singaporean regions, with European expansion still in progress. For studios targeting China and Asia & Oceania, Alibaba Cloud has become one of the strongest options in the market. Establishing the same position in the West remains a longer-term task.

- Alibaba Cloud’s capability assessment now places it closer to the leader group than to any other challenger. While Alibaba Cloud’s games business remains regionally concentrated with the large majority of its current games customers in Asia & Oceania, if its Western traction develops as its capabilities suggest it could, Alibaba Cloud could soon be a global contender.

i3D.net

Figure 11: Omdia Market Radar recommendation – i3D.net



© 2026 Omdia

Source: Omdia

Why consider i3D.net?

- i3D.net is the leading specialized managed private cloud focused primarily on the games industry. Owned by Ubisoft but operationally independent, the company serves a customer base that includes Funcom, Saber, Bungie, Mainframe, and Grinding Gear Games alongside its parent. Its core proposition is bare metal game servers, now operated across 60 physical points of presence spanning six continents, complemented by a games-specific software stack including its GLAD anti-DDoS system, FairFight anti-cheat, game server orchestration, and Kubernetes and Terraform support.
- The platform’s network is a major differentiator. i3D.net operates one of the most peered networks in the world, reaching 65% of global IP prefixes without transiting tier-1 carriers, which translates into predictable performance during launches and traffic spikes. This capacity was demonstrated during Funcom’s launch of *Dune: Awakening*, during which 96% of traffic was routed directly from i3D.net to end-user ISPs. Dedicated network acceleration also enables sub-50ms service to mainland China from Hong Kong.
- Flexibility has improved substantially. Flex Metal provides on-demand bare metal servers deployable within 10 minutes on an hourly, pay-per-use basis with no long-term commitment, narrowing the gap with public cloud elasticity. Hybrid and multicloud deployments link i3D.net bare metal to AWS, Google Cloud, Azure, and Tencent Cloud via WireGuard tunnels or direct connects, and a per-game-instance pricing model offers studios a hardware-agnostic, highly predictable cost structure.

Focus and roadmap


- AI infrastructure is a growing part of i3D.net’s portfolio. The company offers GPU-as-a-Service on Nvidia hardware (A10, B200, and B300 across its private cloud) alongside token-as-a-service inference running on custom silicon from partners including Positron and Axelera.ai. Hosted open-weight models span Llama 3.1, Gemma 2, Qwen 2.5, and GPT-OSS variants, exposed through an OpenAI-compatible API, with fine-tuned and proprietary model hosting evaluated case by case. Partnerships with inference silicon specialists Positron and Axelera.ai, plus “NeoCloud” on-ramps connecting large GPU clusters to public or private networks, position i3D.net as primarily inference support.
- This inference focus aligns directly with where in-game AI is heading. Real-time NPC and agent inference is latency-sensitive and cost-sensitive at scale, a workload profile perfectly suited to i3D.net’s edge-proximate, ASIC-accelerated infrastructure targets. The company’s EU-only platform data residency (Netherlands and France) adds a sovereignty angle relevant to European studios deploying AI workloads.
- The core game server platform continues to advance incrementally, featuring automated demand-based scaling with self-healing resource substitution, multicloud burst scaling, engine-agnostic SDK integration with Unity and Unreal, and full-stack observability from per-server metrics to live DDoS scrubbing status. Remote GPU workstations and managed build machine farms extend the platform into cloud-based game development, though as managed infrastructure rather than managed software services.
- Trust and safety is a quieter growth area. FairFight now provides anti-toxicity moderation for in-game text chat, and a partnership with Modulate extends moderation coverage to voice.

Market impact

- i3D.net’s position as a games-first specialist among generalist clouds remains intact, and it continues to develop its software and services capabilities. LiveOps services such as identity, matchmaking, engagement, and content delivery are addressed through third-party integration rather than first-party product, and its analytics offering is limited to operational and security dashboards. This clarity is a strength for studios that want performance and cost control without platform lock-in, but it cedes the integrated-backend buyer to hyperscalers and BaaS providers.
- The AI expansion meaningfully changes i3D.net’s competitive relevance. By betting on inference-optimized ASICs rather than competing for scarce training GPUs, the company offers a differentiated cost profile for the production inference workloads that games will increasingly run. Its AI stack remains narrow relative to hyperscalers, but for studios deploying open-weight or proprietary models into live games, it is now a real option.
- Commercially, the company continues to punch above its size. High-profile launches such as *Dune: Awakening* demonstrate its ability to support major releases, though multicloud remains the norm among its customers, with studios commonly pairing i3D.net bare metal for steady-state game server workloads with hyperscaler services for elastic backend needs, a pattern i3D.net actively accommodates.

OVHcloud

Figure 12: Omdia Market Radar recommendation – OVHcloud



© 2026 Omdia

Source: Omdia

Why consider OVHcloud?

- OVHcloud is Europe's largest native cloud provider, crossing €1bn in annual revenue in 2025. Its defining characteristic is a vertically integrated model. The company designs its own servers, builds and operates its own data centers (more than 46 across four continents) and runs its own fiber backbone. This end-to-end control underpins two of its core selling points for games companies: predictable pricing and a hardened, self-operated network.
- OVHcloud maintains an explicitly gaming-branded product line. Its dedicated game bare metal servers, offered from eight global locations, are built on high-frequency AMD Ryzen 7 9800X3D and Ryzen 9 9950X3D processors with water-cooling, and include protocol-aware Game DDoS protection with more than 30 predefined game profiles. Servers can be delivered in as little as 120 seconds, with unmetered bandwidth outside Asia-Pacific and no setup fees on longer terms.
- Commercial terms are a key differentiator. OVHcloud charges no egress fees anywhere in its portfolio, including on S3-compatible Object Storage and on data moved to other clouds. For games workloads this significantly improves cost predictability relative to consumption-based egress models, and directly lowers the cost of multicloud operation.
- As the primary sovereign alternative to US hyperscalers in Europe, OVHcloud offers strict region-based data residency, a broad certification portfolio, and SecNumCloud-qualified sovereign zones in France operated exclusively by EU staff. For European studios and publishers with data-protection or regulatory requirements, this is an increasingly important capability.

Focus and roadmap

- OVHcloud's games strategy centers on a tiered compute portfolio rather than managed game services. Dedicated game bare metal is the core line, complemented by a VPS range widely used for entry-level and indie hosting of titles such as *Minecraft*, *Counter-Strike*, *Rust*, and *Valheim*; Public Cloud instances (including GPU) for elastic burst capacity; and High Grade/Scale bare metal for HPC-class backend workloads.

- The game hosting reseller market is a significant and distinctive part of OVHcloud's games business. Reseller-friendly terms have made its game servers a foundation for hosting providers. For example, *Minecraft* hosting specialist Shockbyte runs on OVHcloud's automated server infrastructure, alongside providers such as PebbleHost.
- Scaling follows a hybrid pattern. The recommended architecture pairs a bare-metal baseline with Public Cloud burst capacity, orchestrated through the free-control-plane Managed Kubernetes Service, on which customers deploy open-source tooling such as Agones and Open Match. OVHcloud offers no proprietary game server orchestrator, matchmaking, or fleet management. Rather, orchestration is Kubernetes-native and customer-configured.
- AI, naturally, is a growing focus, framed around openness and European data residency. The stack spans Nvidia GPU instances (L4 through H100/H200) and GPU bare metal, managed AI Notebooks, Training, and Deploy services, and AI Endpoints, a serverless API providing access to pre-trained open-weight models including Llama, Mistral, and Qwen families, with image generation, speech-to-text, and moderation models available for games use cases.
- Infrastructure investment continues at the regional level. Having built out 40 Public Cloud products, OVHcloud's focus is now rolling them out in a three-availability-zone configuration, now live in Paris and Milan, with further regions planned in North America and Asia & Oceania, plus more than 30 edge Local Zones for latency-sensitive deployments.


Market impact

- OVHcloud occupies a position between the hyperscalers and gaming infrastructure specialists. Its general cloud portfolio is far broader than that of specialists like i3D.net or servers.com, but its games-specific software offering is thinner. It lacks integrations the specialists offer, while also eschewing the managed games services (matchmaking, LiveOps, identity) of the hyperscalers. Its natural role is as cost-efficient infrastructure within hybrid or multicloud games stacks, and as the backbone of the hosting reseller ecosystem.
- The sovereignty trend is a meaningful tailwind. European sovereign cloud IaaS spending is growing, and OVHcloud has won positions in the European Commission's Cloud III framework and the European Central Bank's digital euro project. Sovereignty is a weaker purchasing driver in games than in the public sector, but for European publishers handling player data under GDPR, and for companies wary of extraterritorial data access, OVHcloud's default-qualified position is a potentially key differentiator.
- The no-egress-fee model has considerable significance beyond cost. Egress charges are a principal mechanism of cloud lock-in, and their absence makes OVHcloud unusually easy to adopt as a secondary provider. This aligns well with the games industry's near-ubiquitous multicloud practice among larger companies. The absence of a first-party CDN means customers typically also rely on third-party edge delivery, tempering the benefit for global content distribution.
- OVHcloud's games market share remains modest. Nevertheless, its combination of gaming-branded infrastructure, distinctive commercial terms, and sovereign positioning

makes it a viable component of European and cost-conscious cloud strategies, and its addition to this year’s Market Radar reflects its growing visibility in the games infrastructure market.

Oracle Cloud Infrastructure

Figure 13: Omdia Market Radar recommendation – Oracle Cloud Infrastructure



© 2026 Omdia

Source: Omdia

Why consider Oracle Cloud Infrastructure?

- Oracle Cloud Infrastructure (OCI) has undergone one of the most dramatic transformations of any cloud provider over the past year. Once a mid-tier platform competing on bare metal performance and price, OCI is now a key part of the AI infrastructure buildout with performance obligations reaching a record \$638bn in mid-2026 on the back of AI mega-deals with OpenAI, Meta, and Nvidia. For games companies, this repositioning changes the calculus of what OCI represents as a vendor: vastly greater scale and capital intensity, but with attention concentrated on a small number of frontier AI customers.
- OCI’s core appeal to games companies remains its cost-to-performance ratio. The platform’s premium bare metal servers, RDMA cluster networking, and generous networking terms make it well suited to bandwidth-heavy games workloads such as game server hosting and content delivery, where egress charges on rival platforms can dominate the bill. This positions OCI as a compelling cost-optimization play within a multicloud strategy, particularly for studios running their own server orchestration.
- The platform’s GPU access is now among the strongest in the market. Games companies experimenting with generative AI in production pipelines, LiveOps personalization, or in-game AI features can potentially gain access to frontier-class hardware without hyperscaler-level pricing.
- Oracle’s multicloud posture has matured from marketing into a fully differentiated product. As of mid-2026, Oracle Database@AWS spans 20 locations, the Google Cloud equivalent covers 15 with two more confirmed, and Oracle Database@Azure spans 33 regions, meaning games companies can run Oracle’s database services inside the cloud they already use.

Focus and roadmap

- OCI’s roadmap is dominated by AI infrastructure at unprecedented scale. Oracle is one of three founding partners in the Stargate project, the \$500bn AI infrastructure initiative with

OpenAI and SoftBank, and is actively building AI data centers across the US. At its AI World conference, Oracle announced the OCI Zettascale10 supercluster, delivering 16 zettaFLOPS across as many as 800,000 Nvidia GPUs, alongside its Acceleron RoCE network fabric. While these capabilities target frontier AI labs rather than games companies, the resulting capacity, networking innovation, and hardware access flow through to the broader platform.


- The AMD partnership provides a distinctive second GPU track. OCI will be a launch partner for the first publicly available AI supercluster powered by AMD Instinct MI450 Series GPUs, with an initial deployment of 50,000 GPUs starting in calendar Q3 2026. For games companies, AMD-based shapes offer a price-performance alternative and open-source ROCm compatibility, consistent with OCI's positioning against premium-priced hyperscaler AI services.
- OCI's games-industry proposition remains infrastructure- rather than solution-centric. The platform still does not offer managed game server orchestration or matchmaking, while its gaming pitch centers on compute instances spanning VMs, RDMA bare metal, and GPU VMs, low-cost storage tiers, inexpensive data egress, and security-first design principles such as isolated network virtualization. Studios adopting OCI should expect to bring their own orchestration and multiplayer stack, or to rely on third parties.

Market impact

- Oracle's financial trajectory has fundamentally changed the company's risk profile. Oracle lifted its capex plan to \$56bn for the fiscal year, up from \$35bn, as it races to build out Stargate and other AI facilities. The reported cancellation of a planned expansion at its flagship Texas facility with OpenAI is a reminder that large contracts are not immune to renegotiation. Games companies evaluating OCI for long-term commitments should weigh both the platform's expanded capacity and the concentration risk embedded in its AI-heavy backlog.
- The pivot towards AI infrastructure cuts both ways for games customers. In accepting that Oracle will not displace AWS or Azure as a general-purpose cloud, the company has positioned itself as a specialized provider of GPU-dense, HPC-optimized infrastructure competing primarily with CoreWeave, Lambda, Crusoe, and the hyperscalers' own AI-dedicated regions. The games industry, identified as a strategic growth area in previous years, now sits further from the focus of Oracle's attention.
- Within multicloud strategies, OCI's role is expanding from cost-optimization niche to providing interoperability. Oracle plans to connect Oracle Interconnect with AWS Interconnect, giving customers managed connectivity to run applications and move data seamlessly between OCI and AWS, and its Multicloud Universal Credits program allows database services to be purchased across OCI, AWS, Azure, and Google Cloud under a single currency. For games companies this lowers the barrier to using OCI. The absence of mainland China regions remains a limitation, meanwhile, though not one unique to OCI.

servers.com

Figure 14: Omdia Market Radar recommendation – servers.com



© 2026 Omdia

Source: Omdia

Why consider servers.com?

- servers.com sells dedicated bare metal infrastructure as a service, and games represent its single most important market. Its proposition remains unmanaged, high-performance infrastructure for companies with the technical expertise to run their own environments, or those working with partners for orchestration and backend. The company positions itself as one of the largest bare metal hosting providers in gaming, supporting a large number of game launches each year and underpinning many of the industry’s managed hosting and private game server hosting companies.
- Parent company CloudOne Digital appears to be moving towards consolidating its properties, with the business recently rebranded as Servers.com by Nexcess, forming the bare metal foundation of the new Nexcess “Specialty Cloud” ecosystem, with the company stating that its infrastructure, platform, team, and support experience remain unchanged.
- servers.com’s product lineup now comprises Scalable Bare Metal (SBM) for hourly, on-demand server rental; Enterprise Bare Metal (EBM) for customizable longer-term deployments; AI Compute (AIC) for GPU-accelerated workloads; and managed Kubernetes, alongside networking services such as load balancing and firewalls. For games companies, this maps neatly onto the baseline-plus-burst infrastructure pattern: EBM covers cost-optimized baseline concurrent users while SBM provides on-demand capacity for player peaks.
- servers.com does not offer managed game server hosting or orchestration itself, instead operating a partner model. It refers customers seeking managed services to partners including Gameye, Nitrado, AccelByte, and 4Players. This makes it a strong candidate for the bare metal component of hybrid and multicloud setups, but companies wanting a managed platform experience will need to assemble it from partners.

Focus and roadmap

- The core offering remains bare metal servers under an IaaS model, with full customer control over hardware, operating systems, and software. Provisioning speed has become a headline capability, with servers able to be spun up or down in minutes via API, Terraform provider, or the customer portal, closing much of the historical agility gap between bare metal and public cloud for handling launch spikes and seasonal demand.

- The company operates ISO 27001 and ISO 9001 certified data centers across the US, EU, UK, Singapore, and Hong Kong. This footprint covers the major games markets, though it lacks the density of hyperscalers and of games-specialist rival i3D.net.
- AI Compute (AIC) is a notable recent product-line addition. AIC provides dedicated, single-tenant GPU hosting built on the same hardware base as EBM, tailored to specific machine learning use cases, training pipelines, and model sizes. It supports major ML frameworks including PyTorch, TensorFlow, JAX, and ONNX, as well as vector databases, with provisioning via API, CLI, or Terraform for CI/CD and MLOps integration. For games companies, this potentially positions servers.com to capture inference and AI-adjacent workloads alongside game servers.
- Following the Nexcess integration, the roadmap emphasis is on scale and adjacency. The combined Nexcess entity claims more than 100,000 servers serving over 185,000 customers, and positions the merger as giving servers.com customers access to a wider infrastructure footprint and managed cloud expertise.

Market impact

- servers.com continues to occupy a distinct position as unmanaged bare metal beneath much of the games hosting ecosystem. Its role as infrastructure supplier to managed hosting providers and orchestration platforms (Gameye and Nodecraft are both publicized customers) means its games industry footprint is larger than its direct customer roster suggests.
- In principle, the tighter integration with Nexcess gives servers.com access to managed-services capability (via Liquid Web) and greater scale. In practice, however, there is real integration risk. Games customers who chose servers.com for its focused, specialist character will be watching whether that survives absorption into a larger, mid-market-oriented cloud group.
- The economic argument for bare metal has only strengthened over time, with cloud cost discipline now a persistent theme across games infrastructure buying. servers.com's combination of predictable pricing, short contracts, and rapid provisioning positions it well for the growing cohort of studios repatriating stable, predictable workloads from public cloud while retaining hyperscalers for burst capacity.

Appendix

Methodology

This report uses the results of standardized questionnaires sent to each of the featured vendors, focusing on their capacity to support key games industry use cases. The data received from these questionnaires was supported by information gathered from publicly available sources (both published by the vendors themselves, and from third parties) to build a view of their capabilities in each of the areas assessed. Omdia's analysis of the market is underpinned by data from the *Games Tech Market Forecast – 2026*, which contains detailed data on spending by games companies on cloud services, the results of surveys from Omdia's *Game Developer Collective* panel, and on tracking of the broader ecosystem of tech vendors and services in the market drawn from Omdia datasets

such as the *Games Tech Market Landscape* and *Games Tech Mergers, Acquisitions, and Investments Database*.

Further reading

[*Game Developer Collective Survey Database*](#) (July 2026)

[*Games Tech Market Report 2026*](#) (June 2026)

[*Games Tech Market Landscape Database – 1H26*](#) (April 2026)

[*Games Tech Market Forecast – 2026*](#) (March 2026)

[*AI in Games: State of the Market 2026*](#) (December 2025)

Author

Liam Deane, Principal Analyst, Games Tech

askananalyst@omdia.com

Citation policy

Request external citation and usage of Omdia research and data via citations@omdia.com.

Omdia consulting

We hope that this analysis will help you make informed and imaginative business decisions. If you have further requirements, Omdia's consulting team may be able to help you. For more information about Omdia's consulting capabilities, please contact us directly at consulting@omdia.com.

Copyright notice and disclaimer

The Omdia research, data and information referenced herein (the "Omdia Materials") are the copyrighted property of TechTarget, Inc. and its subsidiaries or affiliates (together "Informa TechTarget") or its third party data providers and represent data, research, opinions, or viewpoints published by Informa TechTarget, and are not representations of fact.

The Omdia Materials reflect information and opinions from the original publication date and not from the date of this document. The information and opinions expressed in the Omdia Materials are subject to change without notice and Informa TechTarget does not have any duty or responsibility to update the Omdia Materials or this publication as a result.

Omdia Materials are delivered on an "as-is" and "as-available" basis. No representation or warranty, express or implied, is made as to the fairness, accuracy, completeness, or correctness of the information, opinions, and conclusions contained in Omdia Materials.

To the maximum extent permitted by law, Informa TechTarget and its affiliates, officers, directors, employees, agents, and third party data providers disclaim any liability (including, without limitation, any liability arising from fault or negligence) as to the accuracy or completeness or use of the Omdia Materials. Informa TechTarget will not, under any circumstance whatsoever, be liable for any trading, investment, commercial, or other decisions based on or made in reliance of the Omdia Materials.

CONTACT US

[omdia.com](https://www.omdia.com)

askananalyst@omdia.com