

THG / INGENUITY  
Cloud Services

BUILD FOR WHAT'S NEXT

The cloud fabric for  
scalable, multiplayer  
experiences





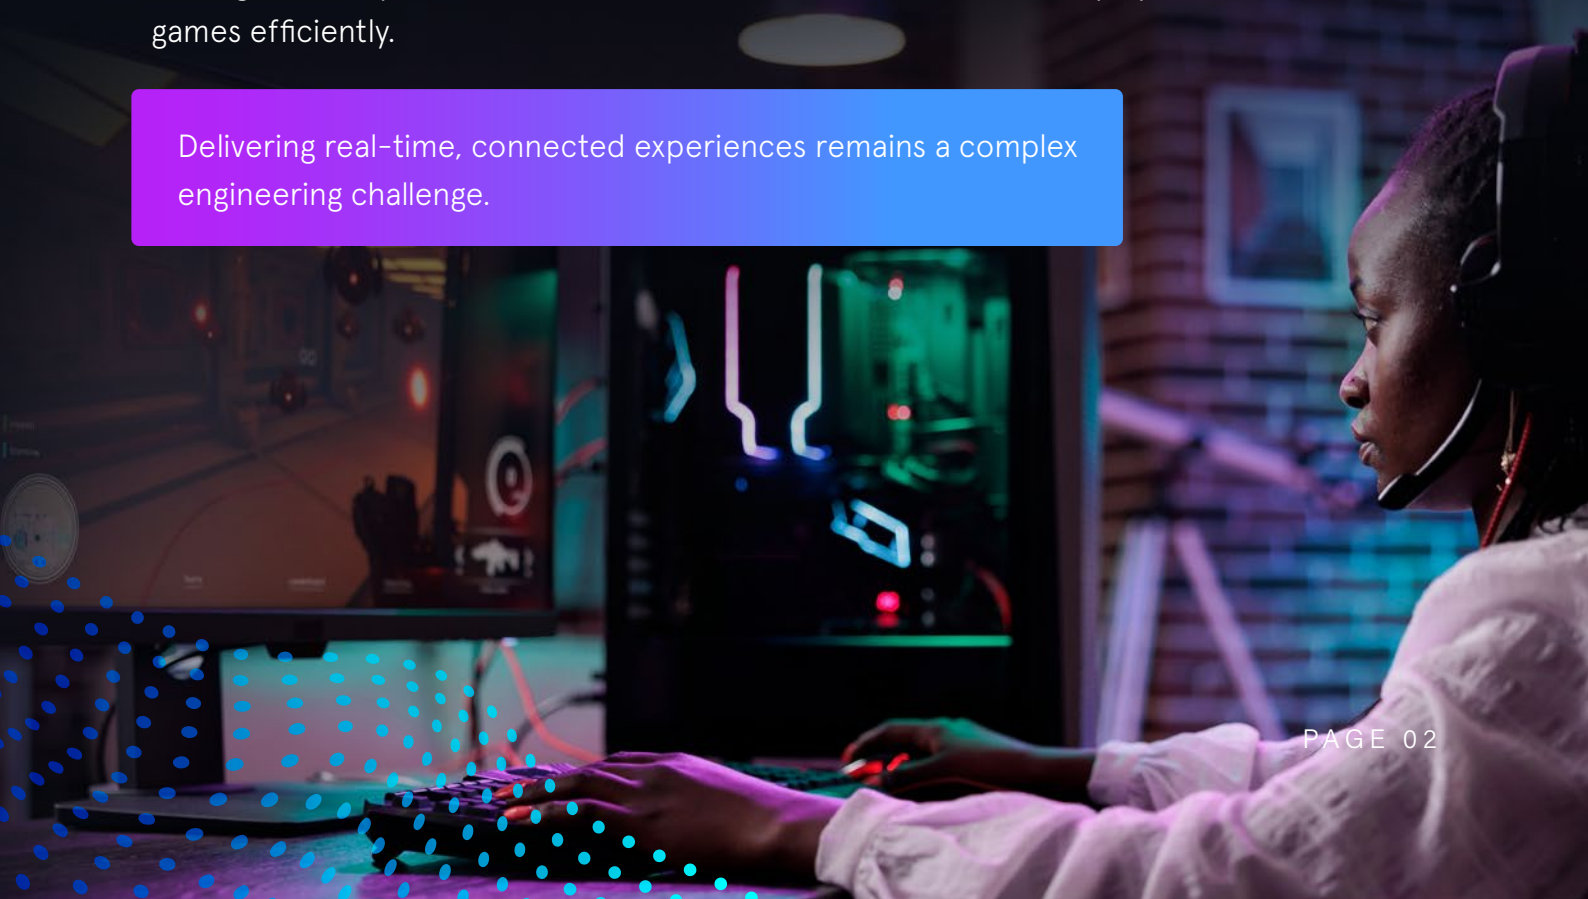
## Mastering speed to market vs. quality gaming experiences

Launching new products on slow platforms or with rigid tech partners can delay time-to-market, reduce agility, and hurt revenue—especially in the fast-paced gaming industry.

From indie teams to AAA studios, developers face shared challenges: maintaining stability, handling traffic spikes, minimizing latency, and scaling globally on a budget. Limited resources make it even harder for small to mid-sized studios to optimize for cloud-native environments, often resulting in delayed launches and inconsistent user experiences.

This eGuide explores key infrastructure considerations and proven strategies to help studios of all sizes build and scale real-time multiplayer games efficiently.

Delivering real-time, connected experiences remains a complex engineering challenge.





## From split-screen to scale: The multiplayer revolution

Multiplayer gaming has advanced significantly since the early days of split-screen consoles and LAN parties. In the early 2000s, titles like Counterstrike and Diablo II introduced gamers to the thrill of online connectivity, laying the groundwork for the always-connected experiences we take for granted today.

Fast forward to the present, and we're immersed in expansive MMOs, high-stakes PvP arenas and genre-defining battle royales such as Fortnite and Apex Legends, which are all powered by complex, real-time multiplayer systems.

Multiplayer formats now span a wide spectrum. Co-op play focuses on teamwork, while PvP emphasises competition. MMOs host thousands of concurrent users, and asynchronous play enables interaction without real-time connection.

“

*As the industry has matured, so too have player expectations. Gamers now demand always-available, low-latency experiences, regardless of platform or geography. A lag spike or outage is no longer just frustrating; it's a deal breaker.*

*From matchmaking and synchronisation to load balancing and orchestration, a robust, scalable infrastructure is key.”*

Cloud Infrastructure Product Manager -  
**Ingenuity Cloud**

## Scaling or failing: Why infrastructure makes or breaks game launches

Scalability isn't just a technical benchmark, it's the cornerstone of player satisfaction, retention and brand reputation. When a game isn't equipped to handle a surge in traffic or maintain consistent performance, there are consequences: players leave, negative reviews flood social media, and long-term brand trust takes a hit.

High-profile launches like SimCity (2013) and Diablo III's infamous "Error 37" highlight how poor scalability can overshadow even the most hotly anticipated releases. Despite achieving massive hype, these games struggled to meet demand at launch, leading to frustrated players and widespread backlash.

In contrast, the successful launches of Fortnite, Among Us and Roblox demonstrate how a solid backend system and proactive scaling can support viral success.

Efficient scalability doesn't just reduce the risk of downtime; it mitigates support costs, drives monetisation, and can retain and strengthen online gaming communities.



## Challenges of multiplayer gaming

While the demand is stronger than ever, delivering seamless, secure, and cost-effective multiplayer experiences presents unique challenges. Whether launching a new title or managing a growing player base, overlooking foundational issues can lead to spiralling expense, unstable gameplay, and churn.

Below we break down some of the most pressing challenges faced when developing multiplayer infrastructure and provide practical insights for tackling each one.



### Cost efficiency

As player numbers grow, so do the demands on infrastructure and – if left unchecked – costs.

One common pitfall is relying on always-on servers, which can waste resources during off-peak hours. Autoscaling helps by allocating resource usage based on demand, ensuring expenses are kept in check without sacrificing performance.

Effective cost control also requires continuous monitoring, right-sizing deployments and

choosing tools that are suited to your traffic patterns. According to a GDC survey, 34% of indie developers cite cloud compute costs for multiplayer, backend, and live ops as a major challenge. Without optimisation, expenses can quickly eat into profits, especially for games with unpredictable traffic spikes.

**34%**

of indie studios cite cloud compute costs as a major challenge



### Poor QA

Multiplayer games introduce quality assurance complexities that go far beyond single-player titles. Bugs related to latency, matchmaking and cross-region play are difficult to catch and even harder to reproduce.

Testing under real-world load is essential, yet 61% of game studios struggle to replicate production environments in the cloud, often defaulting to local testing instead. The result? Bugs that slip through the net and surface only after launch.

To combat this, it is important to invest in robust staging environments that mimic live conditions. Use chaos testing and continuous QA to simulate network variability and concurrency.



### Skyrocketing Egress Cost

Data egress (the cost of transferring data across cloud regions or between providers) can be a hidden drain on budgets. In some cases, egress fees can represent up to 80% of total cloud costs for live multiplayer services.

These costs typically hit indie developers the hardest, but even AAA studios can see millions lost annually. Consequently, smart architectural decisions are crucial. Strategies like edge caching, regional processing and compressing data transfers could help to significantly reduce overhead costs without degrading gameplay.



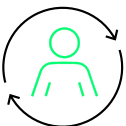
## Security Risk

Multiplayer games face increased attack surfaces. DDoS attacks, cheating, and breaches can disrupt services and be incredibly detrimental to player trust. This is particularly concerning for indie studios, with 37% reporting being vulnerable to security threats due to limited resources.

Meanwhile, competitive and live-service games attract targeted attacks due to their high visibility and financial stakes. A strong security position includes WAFs, anti-cheat systems and encrypted, secure transport protocols. Many studios now also rely on private backbones or providers with built-in DDoS mitigation networks.

**37%** OF INDIE STUDIOS FACE SECURITY RISKS DUE TO LIMITED RESOURCES

**20%**  
rise in churn with 100ms increase in latency



### Higher player churn

In a world of countless free-to-play alternatives, even minor technical issues can cause players to uninstall. Downtime, lag, and poor matchmaking are all cited as contributing factors, with a 100ms increase in latency driving churn up by 20%.

As expected, the stakes are the highest at launch. A poor first impression can hurt your long-term retention and revenue prospects.



### The solution?

Proactive monitoring, global orchestration, and responsive live ops. Optimising server locations and workloads across time zones and peak periods is complex but vital. Player loyalty depends on consistently smooth experiences and failing to deliver risks losing your audience before they even have a chance to connect.



## Unpacking the infrastructure stack: serverless vs. virtual vs. bare metal

---

When choosing the right infrastructure for your game whether you're an indie team or scaling toward AAA it's worth understanding how serverless, VPS, and bare metal compare. Each has its strengths depending on your stage, budget, and performance needs.

Serverless is ideal for smaller backend tasks like leaderboards, matchmaking, or handling webhooks. It's event-driven, scales automatically, and you only pay when your code runs. But it's not built for persistent game sessions or low-latency multiplayer—cold starts and short runtimes can make that tricky.

Virtual Private Servers (VPS) are the go-to for hosting dedicated game servers, especially in early stages. You get more control over the environment, consistent pricing, and decent performance. They're great for modded servers (think Minecraft or Valheim), and you can scale vertically if needed. However, you're still sharing hardware with others, and scaling large workloads takes more effort.

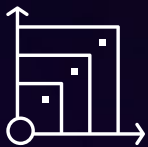
Bare Metal gives you the raw power. There's no virtualisation layer, so performance is predictable and fast, perfect for esports, real-time MMOs, or games where milliseconds matter. You get full control of the hardware, including CPU, GPU, and storage. The trade-off? It's more expensive, takes longer to provision, and needs proper infrastructure know-how.



# Rethinking the Cloud Fabric for improved multiplayer experience

Designing with scalability from the start could help to avoid issues that derail even the most promising game.

Whether you are planning for thousands or millions of concurrent users, early architectural choices have a lasting impact. Incorporating stateless design, modular components and feature flag systems gives your game room to grow, adapt and thrive under pressure, without compromising performance or player experience.



## Design for scalability

Scalability must be factored in from the start, not added on when demand reaches breaking point.

Even in pre-production, laying the groundwork for a scalable game architecture can prevent costly rewrites, technical debt, and player dissatisfaction later down the line.

Stateless architecture avoids tying sessions to servers, making horizontal scaling easier and reducing points of failure. Combined with modular code and feature flags to support flexible updates, live ops, and fast iteration.



## Choose the right infrastructure stack

Choosing the right infrastructure stack means balancing flexibility, cost, and performance. Consider server orchestration tools which offer consistency and autoscaling, and adopt managed services where they add value.

Opt for agility where possible. Embracing cloud-native architecture from the outset will support smoother scale-ups when demand strikes.

Always align your infrastructure stack with your player geography, concurrency goals and long-term growth plans.

## Handling security, networking and performance at scale

Treat performance, networking, and security as a single, unified strategy. All three are key to producing and delivering a reliable, scalable multiplayer experience.

Use regional distribution and edge computing to reduce latency. Likewise, build in security from day one: harden APIs, rate-limit endpoints, encrypt all traffic and enforce secure authentication. For live-

service and competitive games, add anti-cheat systems with real-time detection and rollback capabilities.

Networking resilience is just as important. DDoS mitigation, traffic throttling and secure transport layers protect your backend and – most importantly – your players.



## Building for growth from day one

From architecture and QA to deployment and post-launch support, scalable systems allow you to grow without compromising the player experience. The most successful multiplayer games blend brilliant gameplay with smart technical planning. They don't wait for problems to appear; they build with flexibility and growth in mind from the very start.

If you're a game developer, think beyond the launch. Invest in infrastructure that grows with you, observability that flags issues early, and tools that keep you agile.

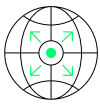
Great gameplay gets players through the door. Scalable infrastructure keeps them coming back.



# Build for what's next with Ingenuity Cloud

With global infrastructure support, teams can scale worldwide without performance trade-offs. We believe in simplifying the shift from local servers to the cloud—freeing developers to focus on building great games, not managing infrastructure.

Powered by a global network of low-latency data centers and top-tier partners like NVIDIA, our infrastructure is built to deliver the ultra-fast, seamless gameplay today's players expect—no lag, just high-performance gaming.



## Presence in over 30 datacentres worldwide

Whether you're serving a local audience or scaling internationally, we help you ensure your infrastructure is responsive, resilient, and ready to meet players wherever they are.



## 24/7 Uptime and protection

Players should not experience latency but reliability wherever they are — which is why we're committed to delivering network reliability around the clock.



## Guiding your tech teams through change and scalability requirements

Technical advisory and professional services that come with our enterprise services or delivered on request depending on your transformation requirements.



## Secure backbone to help you navigate public cloud

Private secure traffic and DDoS protection.



## Transparent billing options and visibility into consumption at scale

We provide clear, predictable billing structures designed to give you full visibility into consumption and costs.

Trusted by the best AAAs and indie studios

[Request case studies >](#)



## Sources and References

---

- 2023 State of the Game Industry report - GDC
- The 2023 State of Game Development & Design Report: Trends to Know - Perforce
- Navigating the High Cost of AI Compute - Andreessen Horowitz
- Gaming Respawned - Akami
- 2022 Unity Gaming Report - Unity



---

**THG / INGENUITY**  
Cloud Services

 [www.ingenuitycloudservices.com](http://www.ingenuitycloudservices.com)

 [cloud@thgintenuity.com](mailto:cloud@thgintenuity.com)

 [thg-ingenuity-cloud-services](https://www.linkedin.com/company/thg-ingenuity-cloud-services)