



Developing Global Multiplayer Games Using Cloud Spanner

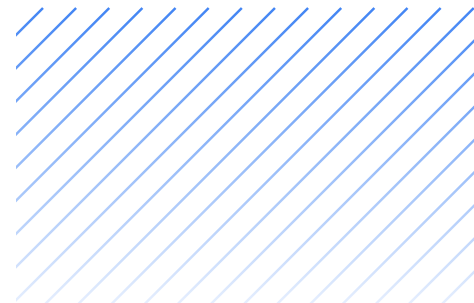
Yoojeong Choi and Paul Hyung Yuel Kim

 Google Cloud



Contents

Gaming industry Introduction	3
Common challenges of making a successful game	5
Gaming workload components	6
Database options for game platform services	9
How Spanner addresses the architectural complexity	10
Scalability	10
Integration with analytics	11
High availability	11
External strong consistency	11
Productivity	12
Security	12
Recommended Spanner adoption process	13
Enablement	13
Schema design	14
Generate randomized value for primary key	14
Define interleaved tables across frequently joined tables	14
Conversion & migration	15
Load test & pre-warming	15
Game launch	17
Use case examples	18
References	19
Other useful links	19






Gaming industry introduction

Like any industry, the gaming industry also consists of multiple parties. At the broadest level, entities involved in gaming are 1. Game Developer (aka Game Studio) 2. Game Publisher 3. First-Party 4. Third-Party and 5. Second-Party.

Game Studio is a company that designs and creates game titles. Much like film, games can be developed by small teams (Indie games) or by hundreds of people spread across international geography with multiple studios with budgets comparable to that of Hollywood films (AAA games).

Game Publishers provide services to Game Studio such as funding, marketing, distribution, public relations, and more. Publishers have long been mandatory to release games due to the high entry barrier but modern digital distributions and additional ways of funding (ex. Crowdfunding) is blurring the line between Game Studios and Publishers.


First-Party publishers are companies with deep pockets such as Microsoft for Xbox and Sony for Playstation. **Third-party** companies develop and publish games without being tied to a particular platform such as Xbox. Activision Blizzard, Electronic Arts, Square Enix, Capcom, and Ubisoft are the well known third-parties in the western market. With the Rise of PC and Mobile gaming, new mega Game Developers/Publishers have emerged in the Asian market such as NCSOFT, Nexon, Netmarble, Krafton, Tencent Games, and more. Lastly, **second-party** is an informal term frequently used to refer to third-party developers with platform exclusive commitments/contracts.



Traditionally, publishers held a stronger position as they had the budget, and market pipelines. With the advent of digital distribution and the blurring of platforms (Console vs general consumer devices) with the **rise of Android/iOS and Play/App Store**, Studios are gaining more power with the increased importance of mega games such as AAA titles. Adding to this complexity is **IP (intellectual property)** deals (Disney, Marvel Studios, etc.) added to the game title series.

New trends such as **5G/Edge Computing on Gaming** and **Game Streaming** (Stadia, Xbox Cloud Gaming, Netflix Gaming) are expected to further shift the relationship between various entities in the gaming industry.

From the computing requirements perspective, gaming genres (Action, Action-adventure, Role-playing/MMORPG, Simulation, Strategy, Sports, Puzzle, etc.) can be considered from 3 dimensions. Light vs Heavy compute bound, global networking for an engaging multiplayer experience with global players, and scalability. Most of the AAA games require DGS (dedicated game server) for providing a consistent world state for immersive gaming experience and are therefore highly compute bound. Gaming titles often take years to make and a ton of investment. Game players expect global community co-play experience (multiplay) as a mandatory for most of the modern game titles and therefore DGS is the single most important bottleneck from the scalability perspective. The success of the game depends on how well it can scale from thousands to millions of players in the first few weeks if it becomes a hit. To summarize, it's important to make sure DGS computing needs are met with high performance machines, global network connectivity is satisfied with low latency from major cities where the players are located, and achieve game scalability through flexible architecture designs as well as securing adequate capacity for both organic and sudden demand using Cloud provider infrastructure such as GCP.





Common challenges of making a successful game

There are many challenges involved with launching a successful game in the current entertainment industry landscape.

Players expect not only compelling games but a myriad of online features such as friend lists, leaderboards, periodic quests, multiple multiplayer modes, seasonal content add-ons, tournaments, e-sports, and much much more. Studios invest enormous resources into developing games, often for more than 2 years for AAA grade games and capital ranging in the two to three digit millions. To prevent re-inventing the wheel for every title, common components are aggregated into a separate Game Platform Services, often unique combinations per gaming companies. There is an active effort within the industry to commoditize such services. Some notable OSS solutions are [Agones](#), [Open Match](#), [Quilkin](#), and [Open Saves](#).

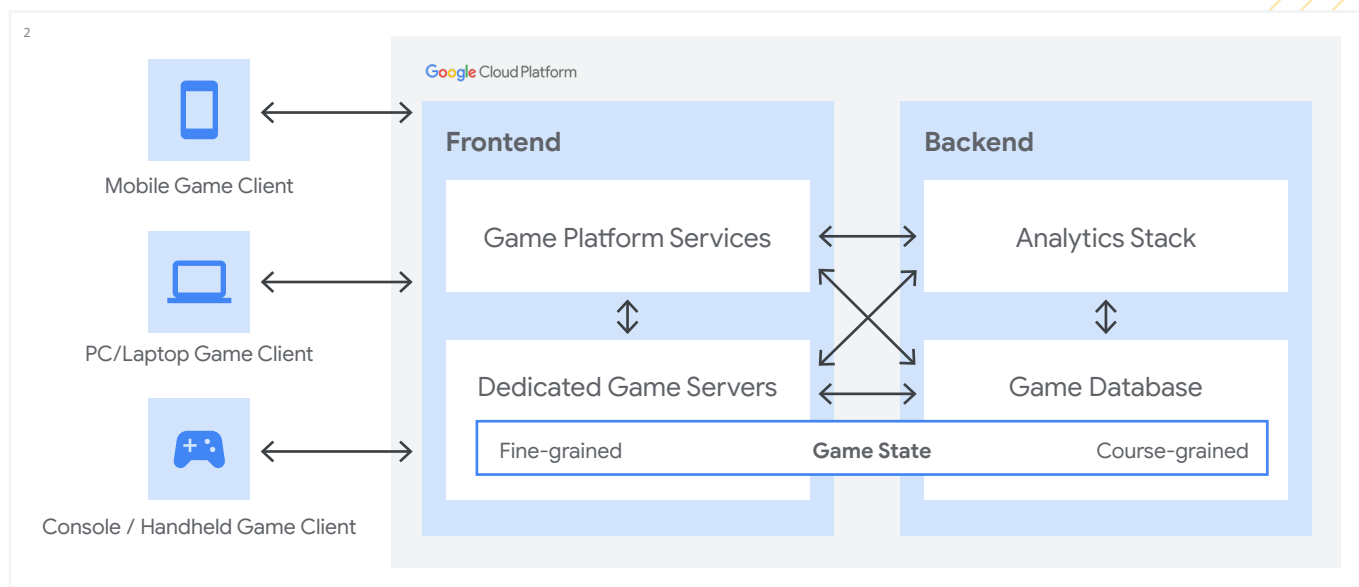
All that investment and time relies on the successful game launch as any hiccups during the initial opening of the game are often non-recoverable as the players don't come back and the hype of the anticipation quickly fades away. The ability to scale from handling thousands to millions of players on the day of the launch is what Google Cloud can help to achieve with unprecedented ease via solutions such as [Cloud Spanner](#) and scale at which Google Cloud operates [globally](#).

Successful game launch requires elements such as Play Store ads with favorable rankings & ratings, reviews on popular games channels for promoting and generating excitement, just like the film industry. In order to provide personalized marketing and increase player retention, Game Analytics solutions are an essential part of any significant gaming workloads. Data driven marketing helps create and keep the momentum of the game launch.

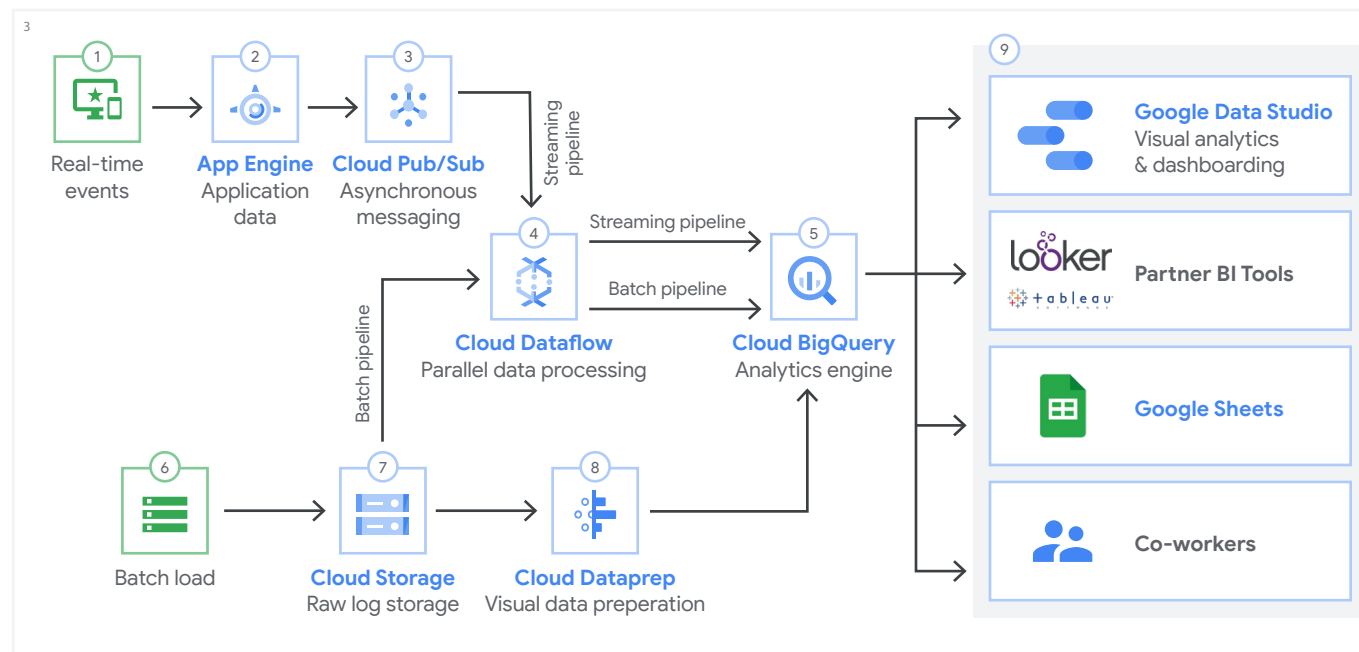
Let's take a closer look into each of those component stacks and their requirements & challenges.

¹ [https://en.wikipedia.org/wiki/AAA_\(video_game_industry\)](https://en.wikipedia.org/wiki/AAA_(video_game_industry))

Gaming workload components



Player Profile, Leaderboards, Game Chat and other services are provided from Game Platform Services while Game state is maintained on Dedicated Game Servers and Database.



Game Analytics stacks are often designed with both streaming and batch processing mode to provide personalized gaming experiences in semi-realtime (ex. Marketing events such as discounted game items) as well as game player trends for the game developers to continuously evolve the game for better gaming experiences.

² <https://cloud.google.com/architecture/cloud-game-infrastructure>

³ <https://cloud.google.com/architecture/mobile-gaming-analysis-telemetry>

To support the various success elements for a game launch, it requires server backends. Some examples of the server backends are DGS (dedicated game servers) and Platform Services such as Matchmaking Servers, Chat and Messaging Servers, Leaderboard Services, Analytics for detecting bad actors as well as creating personalized services, and much more. If the game launch goes viral as everyone hopes, the backend servers need to be scaled seamlessly to 100 or 1000 times the initial size in a matter of hours.

Player Profiles Services: Players expect cross-device experience where they can freely move between the devices they own with a single consistent view of data including saved games, profiles, and various other data. Global single source of truth for player data is needed that can be accessed at a reasonable latency such as milliseconds range. Both RDBMS and NoSQL are good candidates for this as long as the network fabric can support the latency requirements. Spanner is a great solution that can meet the consistency requirements as well as the scalability requirements for all game sizes with a strong security foundation.

Micro Payments Model: Game business model has made a significant shift in the last few years as they started supporting free-to-play model where players can download and play the game for free with the ability to make IAP (In App Purchases) such as weapons, outfits, and game items by purchasing game currency such as coins. This requires a highly consistent datastore that supports transactions with full ACID properties that can also scale heavily should the game succeed at the initial launch. Again, Spanner is a great match for this component.

Analytics: To support the Micro Payments Model, game companies need to maximize the long-tail revenue to realize the full revenue potential. This requires collecting and analyzing a large number of metrics about player behavior patterns, game items, purchase history, favorite game contents and more. Analytics will provide insights on how to keep the game relevant by creating personalized game experiences. Streaming data pipelines that can process the semi-realtime EDW requirements are needed to realize the required insights at the various levels mentioned above. Pub/Sub⁵, Dataflow⁶, Dataprep⁷, BigQuery⁸, and Looker⁹ are often used for dataflow, EDW, and intelligence solutions.

⁴ <https://cloud.google.com/spanner>

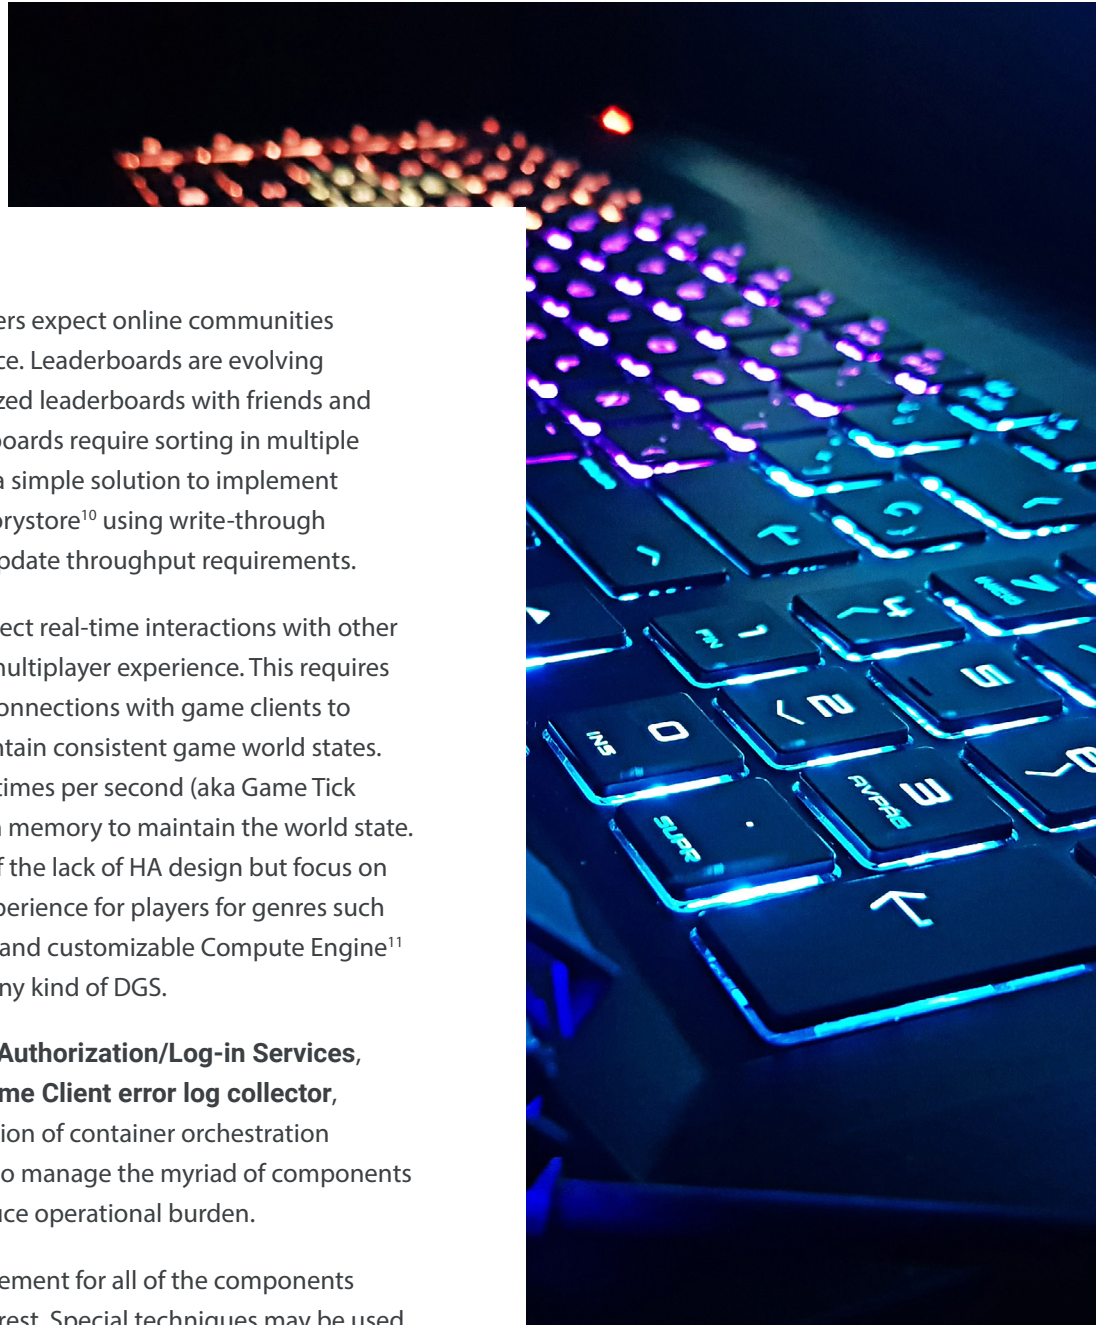
⁵ <https://cloud.google.com/pubsub>

⁶ <https://cloud.google.com/dataflow>

⁷ <https://cloud.google.com/dataprep>

⁸ <https://cloud.google.com/bigquery>

⁹ <https://cloud.google.com/looker>



Leaderboards and Rankings: Players expect online communities for competitive game play experience. Leaderboards are evolving from a single global list to personalized leaderboards with friends and clan rankings. Therefore, the leaderboards require sorting in multiple keys with high performance. This is a simple solution to implement on memory DB such as Cloud Memorystore¹⁰ using write-through architecture for attaining the high update throughput requirements.

Game-State Authority: Players expect real-time interactions with other players with low latency for online multiplayer experience. This requires DGS to host synchronous network connections with game clients to collect game events in order to maintain consistent game world states. Since the update happens multiple times per second (aka Game Tick Rate), highly accessed data is kept in memory to maintain the world state. Game Studios understand the risk of the lack of HA design but focus on performance to provide the best experience for players for genres such as MMORPG. Google offers a secure and customizable Compute Engine¹¹ that can meet the requirements of any kind of DGS.

Other notable components include **Authorization/Log-in Services, In-Game notification systems, Game Client error log collector, Matchmaking Services**, etc. Adoption of container orchestration systems such as GKE¹² is on the rise to manage the myriad of components through a common platform to reduce operational burden.

Also, **security** is the common requirement for all of the components such as encryption in-transit and at-rest. Special techniques may be used such as adding authentication tokens to each 'application level' message packet, byte offset matching (DPI¹³-lite) for network layer security, and more. Such techniques require implementation at both the game client and backend. Backend can have multi-tier architecture to offload the security processing such as proxy based solutions.

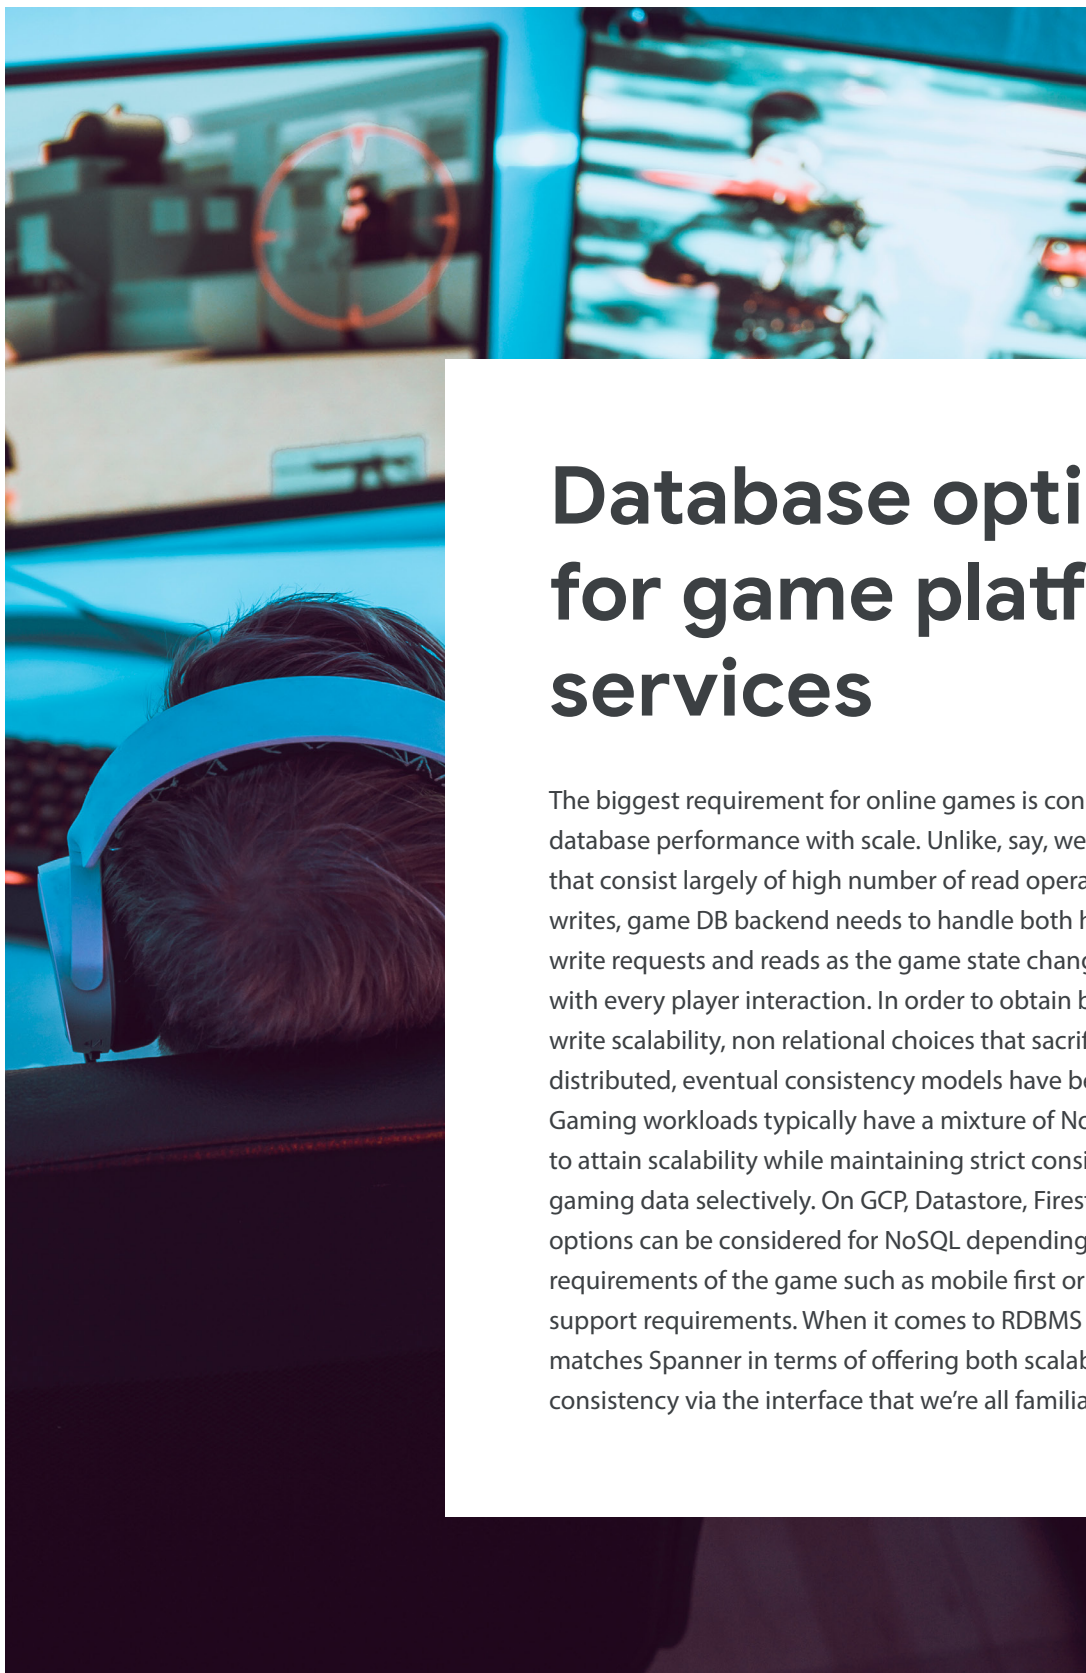


¹⁰<https://cloud.google.com/memorystore>

¹¹<https://cloud.google.com/compute>

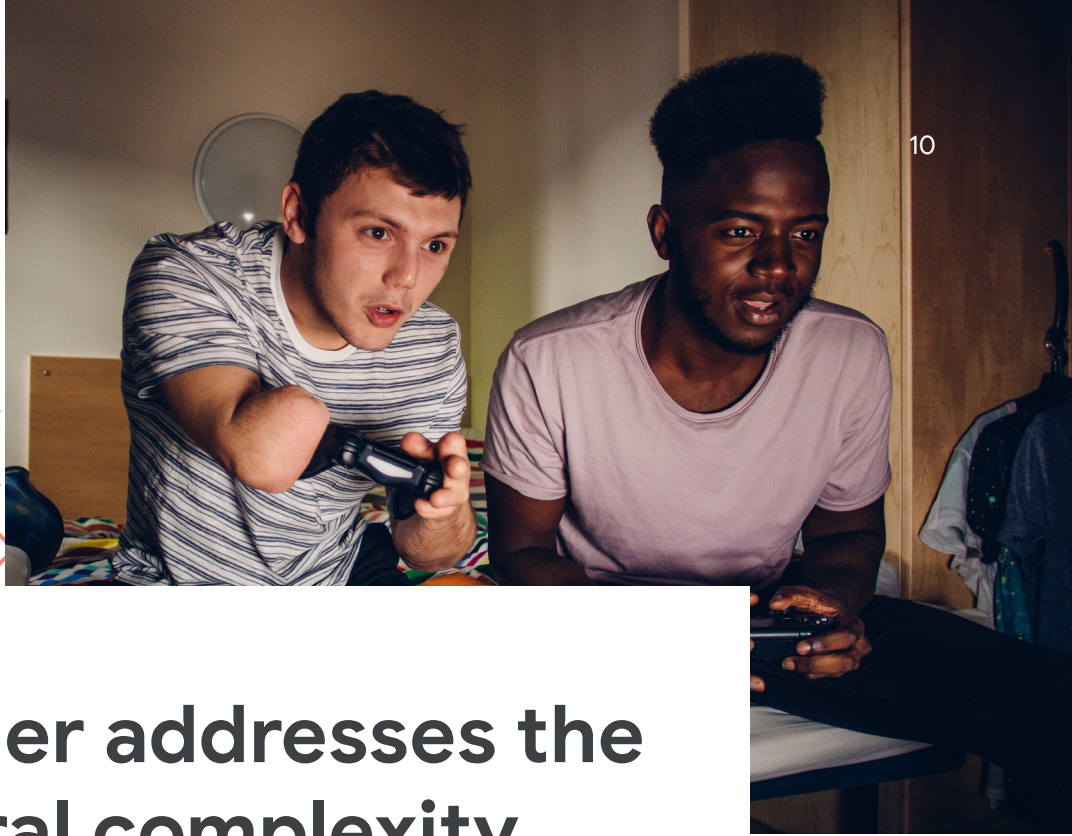
¹²<https://cloud.google.com/kubernetes-engine>

¹³https://en.wikipedia.org/wiki/Deep_packet_inspection



Database options for game platform services

The biggest requirement for online games is consistent predictable database performance with scale. Unlike, say, web applications that consist largely of high number of read operations vs few writes, game DB backend needs to handle both high number of write requests and reads as the game state changes constantly with every player interaction. In order to obtain both read and write scalability, non relational choices that sacrifice ACID over distributed, eventual consistency models have been embraced. Gaming workloads typically have a mixture of NoSQL and RDBMS to attain scalability while maintaining strict consistency on gaming data selectively. On GCP, Datastore, Firestore, and Bigtable options can be considered for NoSQL depending on the unique requirements of the game such as mobile first or cross platform support requirements. When it comes to RDBMS on GCP, nothing matches Spanner in terms of offering both scalability and global consistency via the interface that we're all familiar with, SQL.



How Spanner addresses the architectural complexity

NoSQL typically lacks the features on transactions and advanced querying. To overcome this limitation, dispersing data over both NoSQL and RDBMS can be considered. However, this adds complexity to maintaining a single source of truth. Spanner offers a unique opportunity to simplify this as it's the only option that offers both scalability, consistency, security, as well as high availability from a single database solution.

Scalability

Games are notoriously difficult to predict with regards to traffic demand. Depending on the success of the game, scaling requirements can easily cross into millions of concurrent users in a matter of days or even hours. Sharded MySQL databases are predominantly used for database scaling use cases where transactions are needed, increasing operational expense due to the manual nature of sharding MySQL databases. The goal of the Spanner development was to create a data storage service for those applications that have complex, evolving schemas, or those that want strong consistency with wide area replication requirements. Spanner is able to scale to arbitrarily large database sizes and operations per second, supporting well known Google applications such as Gmail, GCS and AdWords. In addition, regular/predictable fluctuations in traffic can be satisfied by scaling the spanner nodes via api calls. Un-anticipated traffic demand can be addressed via [Autoscaler tool for Cloud Spanner](#), an open source tool that can scale up or down spanner nodes based on cpu utilization metrics and help control costs. Spanner decouples compute from data storage which makes it possible to scale the pool of processing resources separately from the underlying storage. This means that horizontal upscaling is possible for achieving higher performance on dimensions such as operations per second for both reads and writes. With Spanner, customers don't need to concern themselves with read replicas and sharding anymore and consider scalability over read, write, and compute dimensions separately. For example, compute power can scale separately from storage for compute intensive operations such as random number generation. Storage can also grow separately without requiring a matching increase in compute resources. Other scaling concerns, such as maintaining low latency for requests originating from multiple continents, multi regional Spanner instances can satisfy the requirements.

Integration with analytics

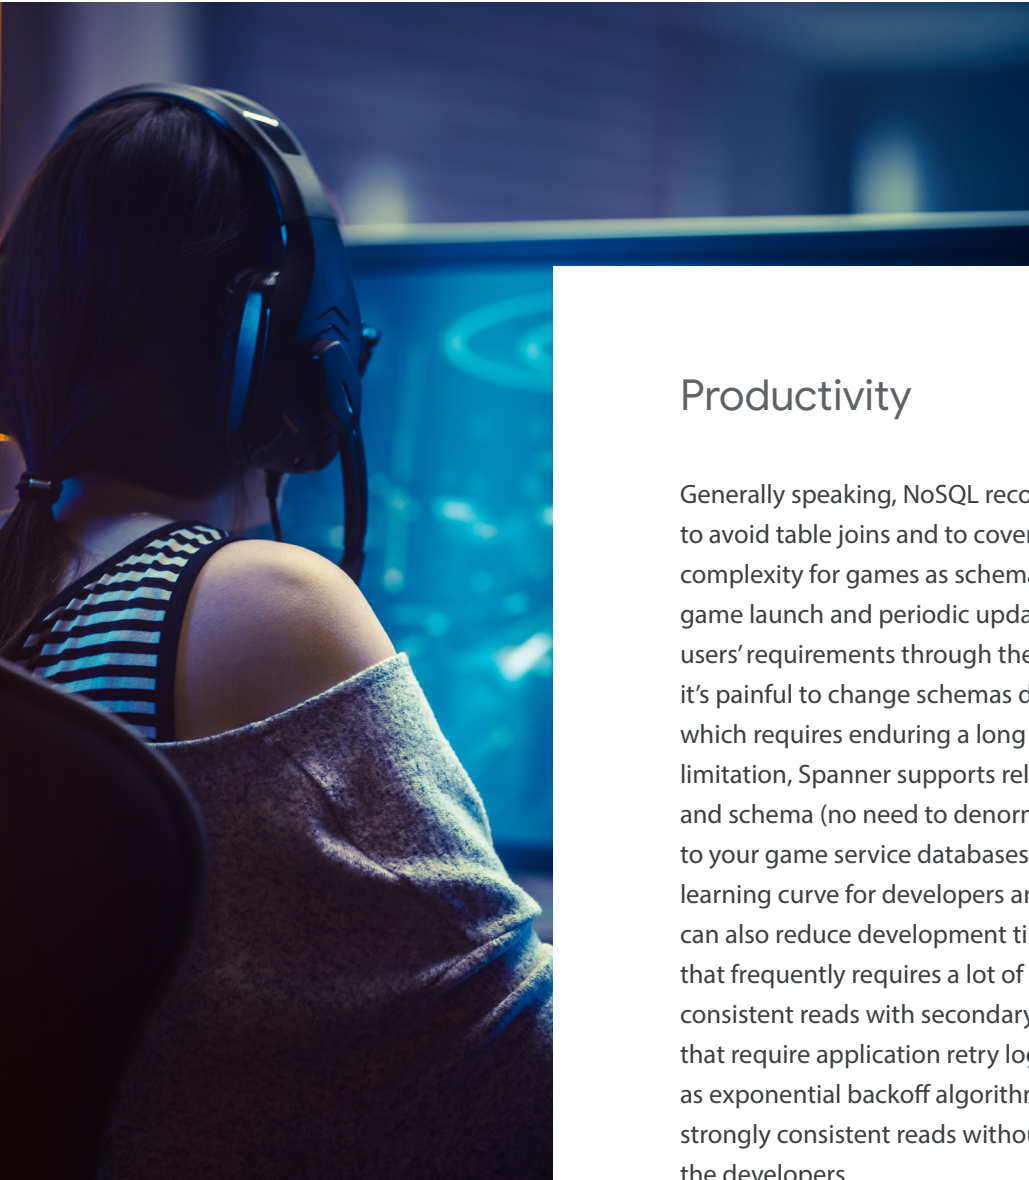
Throughout the full game life cycle after launch, data engineers want to build and maintain data pipelines as well as a matching analytic system in order to obtain various KPIs and insights. ETLing or real time streaming ingestion is done by integrating several solutions ([Dataflow](#), [Data fusion](#), [Datastream](#) etc.). Such solutions can get complex very quickly depending on the number and variety of data sources and relevant business logics. BigQuery federation query from Spanner can simplify these pipelines to ingest dataset into BigQuery and enable easy exploration before building complete pipelines in real time manner.

High availability

Unplanned game downtimes are the single most dangerous threat to destroying the longevity of game titles. That's why game companies seek highly available backend databases to minimize game service interruption in case of unplanned failures. Cloud Spanner delivers industry-leading 99.999% availability for multi-regional instances—10x less downtime than four nines—and provides transparent, synchronous replication across both regional and multi-regional configurations. This multi-region configuration can also meet Disaster Recovery requirements from regulation or compliance perspectives. In addition, Spanner is a fully-managed service with automatic no-downtime patches and upgrades, even schema changes are applied to Spanner without downtime.

External strong consistency

Most games need to provide Leaderboards that should return consistent ranking at any given time across the game players all over the world. The periodic rank update operation can be simplified by delegating the relevant query to Spanner and keeping the results on memory DB such as Memystore for fast lookups. Spanner can provide accurate rank query results as it can always serve strongly consistent reads based on the unique power of TrueTime that doesn't need any multi-phase consensus protocols or additional locks. Strongly consistent reads are default which means that there's no additional required work from the game developers to achieve this. Of course, Cloud Spanner also provides '[global bounded-staleness reads](#)' with TrueTime when required.



Productivity

Generally speaking, NoSQL recommends schema denormalization to avoid table joins and to cover diverse access patterns. This adds complexity for games as schema changes are inevitable after the game launch and periodic updates are mandatory to meet the end users' requirements through the entire game life cycle. In that case, it's painful to change schemas due to data redistribution operation which requires enduring a long game downtime. Contrary to this limitation, Spanner supports relational semantics like ANSI SQL and schema (no need to denormalize) and enables easy updates to your game service databases online. ANSI SQL can shorten the learning curve for developers and DBAs. In addition, ORM support can also reduce development time. Another common constraint that frequently requires a lot of work from developers are strongly consistent reads with secondary indexes in distributed databases that require application retry logic with cushioning techniques such as exponential backoff algorithms. Spanner supports read-on-write, strongly consistent reads without any additional work/code from the developers.

Security

Spanner has compliance certifications which game services might need depending on the regional jurisdiction, like PCI, SoC compliance, and FedRamp. In addition, [VPC-SC support](#), Audit Logging including not only admin jobs but also user activities (DML, DDL, even Query) and Access transparency which is where GCP really differentiates itself that if a GCP person or SRE has to touch your data for whatever reason, it can be tracked and logged which can increase your game database's security level even further.



Recommended Spanner adoption process



Enablement

Most engineers in the gaming industry are still much more familiar with relational databases like MySQL, PostgreSQL and SQL Server and so forth over NoSQL technologies. Even though Spanner supports relational semantics, the first and most important step for accelerated Spanner adoption is to get acquainted with the differences compared to traditional relational DB with emphasis on the distributed architecture. It is recommended that the development team start evaluating Spanner at the earliest stages of the game development lifecycle to avoid investing in undifferentiated work. Cross team collaboration between the DBA and the dev team needs to happen early on to prevent re-inventing the wheel.

Schema design

Most gaming workload data flows are based on individual player activity, which means that data access paths are based on game end user id (e.g. player id). Hence, the most important thing for increasing performance is choosing the right primary keys to prevent hotspots.

Generate randomized value for primary key

Sample Go code to generate randomized INT64 value

```
func RandomInt64() int64 {
    val, _ := rand.Int(rand.Reader, big.NewInt(int64(math.MaxInt64)))
    return val.Int64()
}
```

Sample Node.js code to generate version 4 UUID

```
import { v4 as uuidv4 } from 'uuid';
uuidv4(); // '9b1deb4d-3b7d-4bad-9bdd-2b0d7b3dcb6d'
```

In addition, each game user related information needs to be clustered as much as possible to query the required data at once with only a small I/O overhead. For instance, player profile, guild and achievement information are mostly queried together.

Define interleaved tables across frequently joined tables

Interleaved tables	Sample query
<ul style="list-style-type: none"> + Parent Table + Child Table 1 <ul style="list-style-type: none"> + Grandchild Table 1 + Grandchild Table 2 + Grandchild Table 3 + Child Table 2 <p>→ Key: player_id</p>	<pre>SELECT FROM `Parent Table` a JOIN `Child Table 1` ui ON a.player_id = ui.player_id JOIN `Grandchild Table 1` us ON a.player_id = us.player_id WHERE</pre>



Conversion & migration

The ANSI SQL and Schema support of Spanner allows customers to easily convert existing Relational DB schemas as well as relevant queries compared to efforts required for converting from relational to NoSQL. It is proven that several gaming customers completed this conversion in about 2~3 weeks. In addition, an evaluation and migration open source tool, [Harbourbridge](#) can accelerate the conversion from MySQL and PostgreSQL to Spanner. If zero downtime data migration is required for business continuity, serverless replication service, Datastream as well as 3rd party CDC and data stream solutions like Striim¹⁴ can enable you to minimize the downtime during the cutover.

Load test & pre-warming

Spanner is a distributed database that divides your data into chunks called [splits](#). Until meeting the target performance requirements such as throughput and latency, continuous load testing and optimizations need to be conducted. [Query statistics](#) give insights into which queries put heavy load on your database. After figuring out the major slow queries, visualized [query execution plans](#) indicate which execution step takes a long time and consumes much CPU and which steps returns many rows. (Figure1. Sample query execution plan) This information enables you to tune slow queries intuitively with ease. Additionally, [Transaction statistics](#) and [Lock statistics](#) help you find out the cause of the slow running transactions regarding the write workloads.

To secure the minimum required throughput on the launch day, Spanner needs to be in a warm state with enough splits already balanced across all of the nodes via [pre-warming the database](#).



¹⁴ <https://www.striim.com/>

(Please refer to the [Tuning a query using the query plan visualizer](#) for more detail)

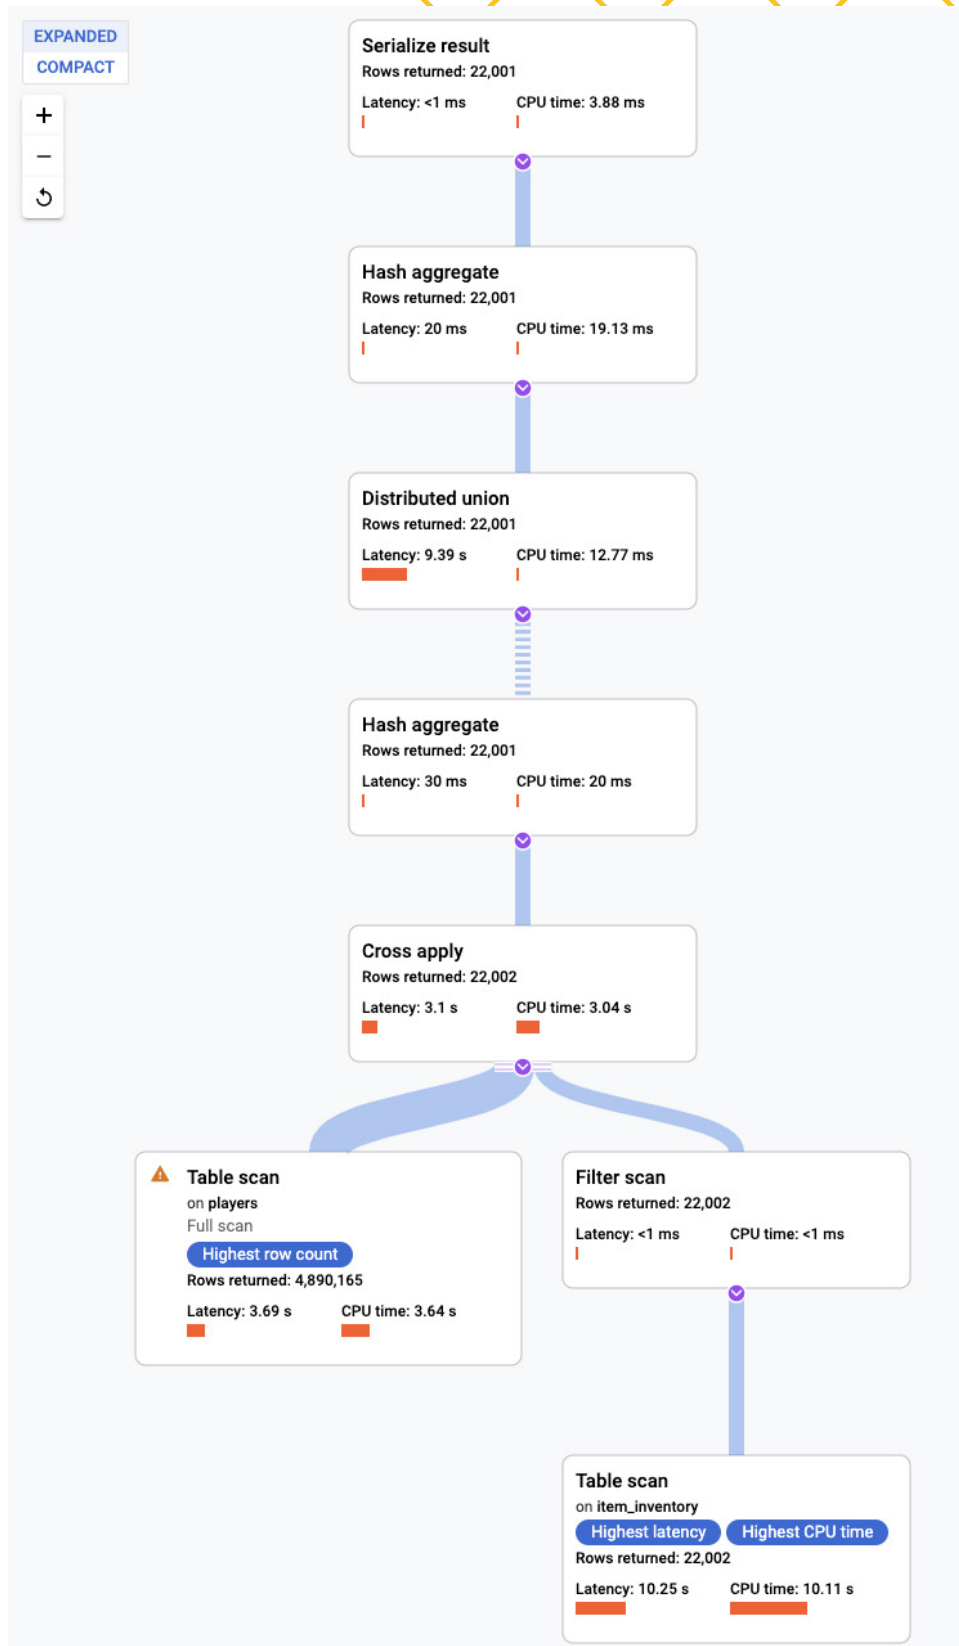


Figure 1. Sample query execution plan

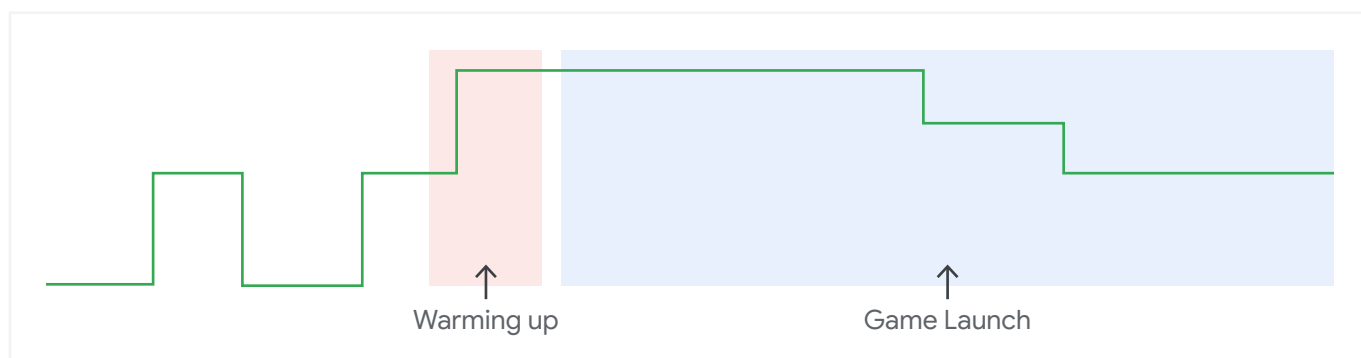
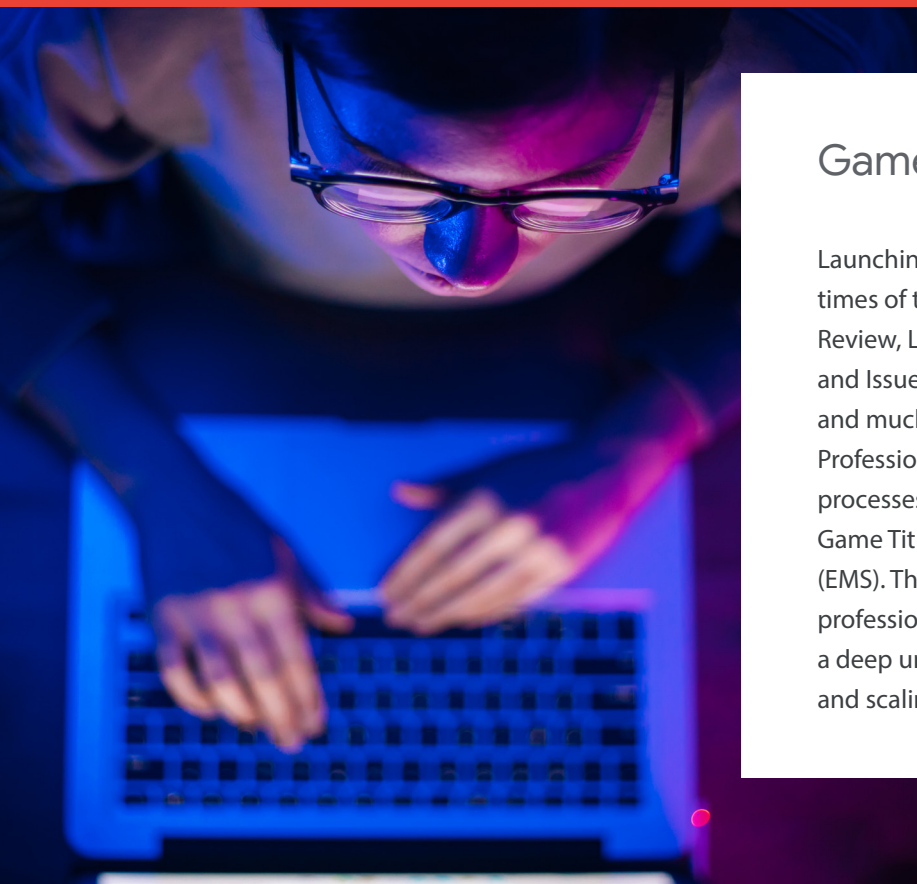


Figure 2. Flexible Spanner Node Count through game life cycle

This graph is from a real world case with a big game launch through pre-warming. After the peak traffic on the launch day, the number of nodes could be decreased for optimization as the traffic from the opening hype went down.



Game launch

Launching new games requires specific activities at specific times of the game development stages such as Architecture Review, Load Testing, Capacity & Quota assessment, Risks and Issues tracking, War Room, escalation procedures, and much more to ensure a smooth game success. GCP Professional Services Organization (PSO) can help these processes through specifically designed programs such as Game Title Launch Assist and Event Management Service (EMS). These programs are designed to partner the right GCP professional resources such as Consultants and SREs to gain a deep understanding of the game and provide architectural and scaling guidance for the launch event.

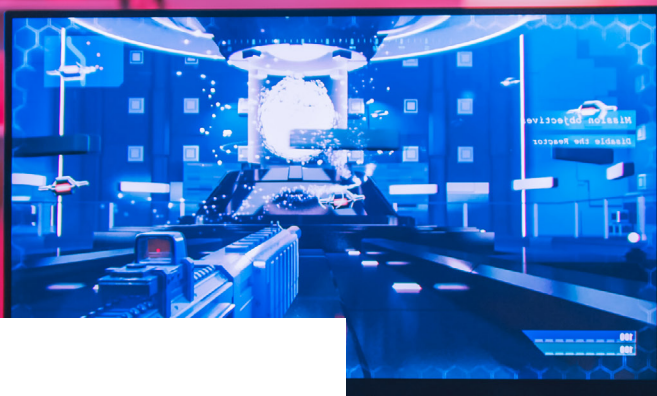
Use case examples

Use case requirements for gaming workload	How Spanner addresses these requirements
Online Read/Write Scalability	Spanner scales out/in without downtime to deal with unpredictable game workload, not only Read but also Write.
Zero downtime for continuous game service	Online Schema change and no maintenance downtime keep game DB alive.
Capability to cover diverse access patterns	Strong consistent secondary indexes support any access pattern without additional apps. code as opposed to other NoSQL DB with eventual consistency.
Development productivity	ORM (Hibernate/Spring Data/ Django (Beta)) support, OSS JDBC driver and local emulator help increase programming productivity. In addition, Operation suites and OpenCensus integration enable easy monitoring and introspection.
Minimize RTO, RPO	Spanner Point-in-time Recovery can recover old versions of data from any point-in-time within retention (1 hour~7 days) immediately without bringing up the old version whole DB.
Audit Logging to meet regulatory	Spanner Audit logging supports any changes on game DB including admin activity as well as Data Read/Write. Cloud Logging can sink Audit logs into BigQuery for easy access, cost-effective long-term archiving and analysis.
Unlimited scale	Limitless scale from GBs to PBs; Trillions of rows and millions of columns; Linear scalability; Separately scale compute and storage;
Business-critical reliability	Availability SLA of 99.99% with replication across multiple zones and 99.999% with replication across regions; Automatic failover in case of a zonal failure; Proven track record of powering core Google services
Easy to manage	Fully managed service



References

- Using Memorystore for Redis as a game leaderboard [[Tutorial](#), [Github](#)]
- Cloud Spanner Emulator [[Blog](#), [Github](#)]
- Spanner Client libraries [[Java \(Tutorial\)](#), [Go \(Tutorial\)](#), [Node.js \(Tutorial\)](#), [Python \(Tutorial\)](#), [PHP \(Tutorial\)](#), [Ruby \(Tutorial\)](#), [C# \(Tutorial\)](#), [C++ \(Tutorial\)](#)]
- Spanner JDBC Drivers [[Google's OSS \(Tutorial\) - Recommended](#)], [Simba \(Tutorial\)](#)]
- ORMs [[Spring Data](#), [Hibernate ORM](#), [Django](#)]



Other useful links

[Spanner Schema design best practices](#)

[Cloud Spanner Ecosystem - An open source community for Cloud Spanner](#)

[Best Practices for using Spanner as a gaming database](#)

[How Pokémon Go scales to millions of requests with Cloud Spanner - Video](#)

[How Vimeo uses Cloud Spanner - Video](#)

