



## Exploring Services Architecture at Bungie

Michael Williams

# Hi Everyone!

- **I'm Michael Williams**
- **This was built by a ton of talented folks**
- **I'll be in the chat!**



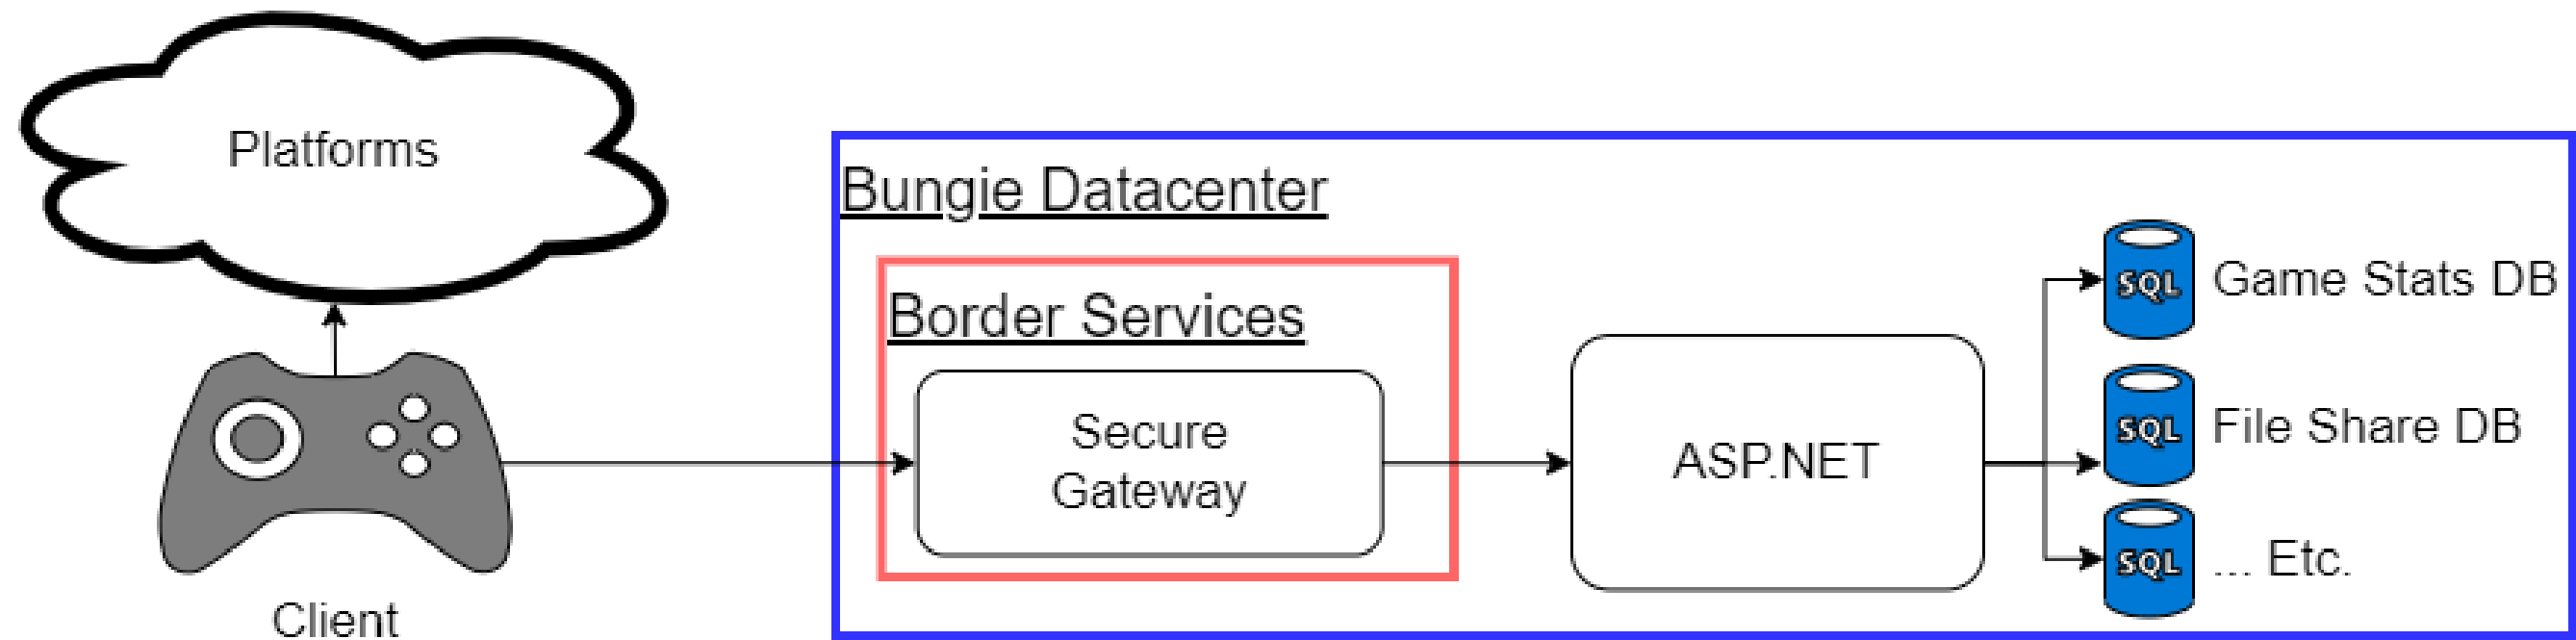




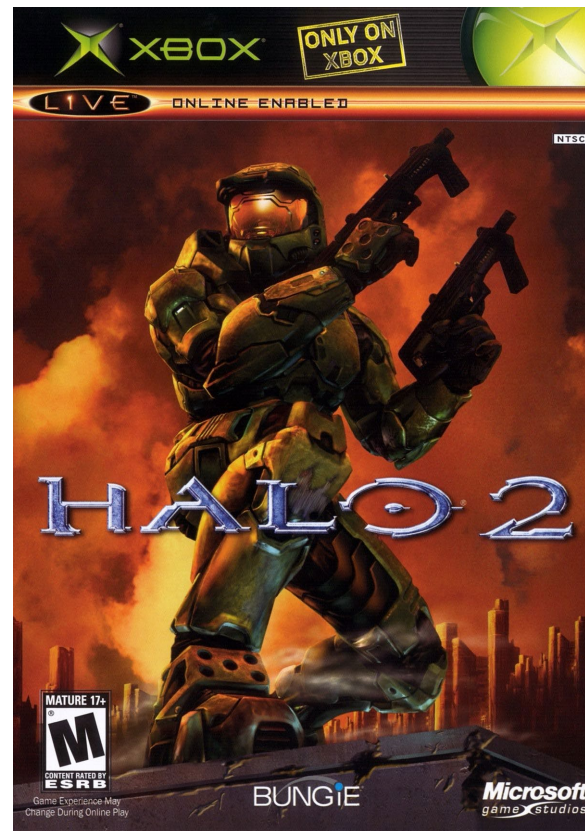
# A Brief History of Services

- **Halo Services**
  - Xbox Live handled Matchmaking, Friends, Session state, and (most) Presence
  - Bungie Services
    - Bungie.net Stats
    - File Share (screenshots, saved films, edited maps)
    - Cheat detection
    - Playlist settings
    - Player server-to-client messaging

# Halo Service Technology



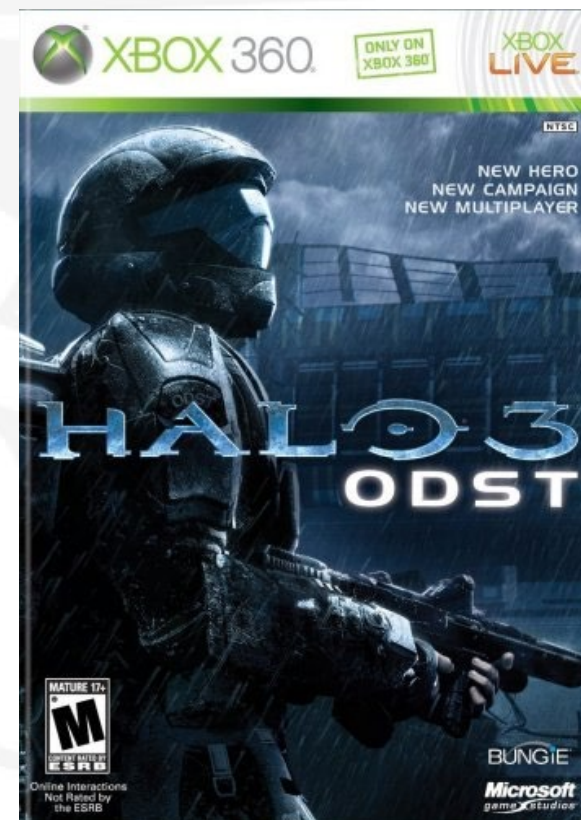
# Service Engineers per-game



.5



2



3



4

Note: Excludes Bungie.net Engineers and Gameplay Networking Engineers

# Pivot to Destiny

- **Always Online game**
- **Loot-based gameplay means larger player data**
- **Multi-platform means high concurrency**
- **A complex hybrid hosting plan**
- **A “single shard” universe**

# Pragmatic Choices

- **Stick with Windows servers to host Destiny**
- **Continue using C# to write services**
- **Microsoft SQL as the Backing store**



# Additional Choices

- **Destiny services aren't highly latency sensitive**
  - We can build out a single Datacenter to control complexity!
- **Grow the team to meet the challenge**
  - 4->16 Service Engineers at the time of Destiny 1 Launch

# Success!!

Home › Originals › A Thank You to Sony, Microsoft and Bungie for a smooth Destiny...

Originals

## A Thank You to Sony, Microsoft and Bungie for a smooth Destiny launch

By **Matt Liebl** - September 11, 2014

## Destiny 2 enjoys a fairly smooth launch

There are some issues being reported, but for the most part, it's good news.

2017-09-07 11:49 • [Bengt Lemne](#) •

## ...Mostly

## Kicks and queues: Destiny's Rise of Iron expansion is having a bumpy rollout

Posted 5 years ago by [Chris Carter](#)

33

EDITORS' PICK | Jan 29, 2020, 08:27am EST | 6,141 views

### 'Destiny 2' Restores Missing Currency In Its First Ever Rollback, Losing An Hour Of Game Progress



# Destiny Today

- **Nearly 40+ Services**
- **18 unique SQL Databases**
  - Not counting shards!
- **9 Redis caches**
- **Plus plenty of other tech**
  - Kibana,
  - Elastic Search / Graphana,
  - Redis,
  - Etc

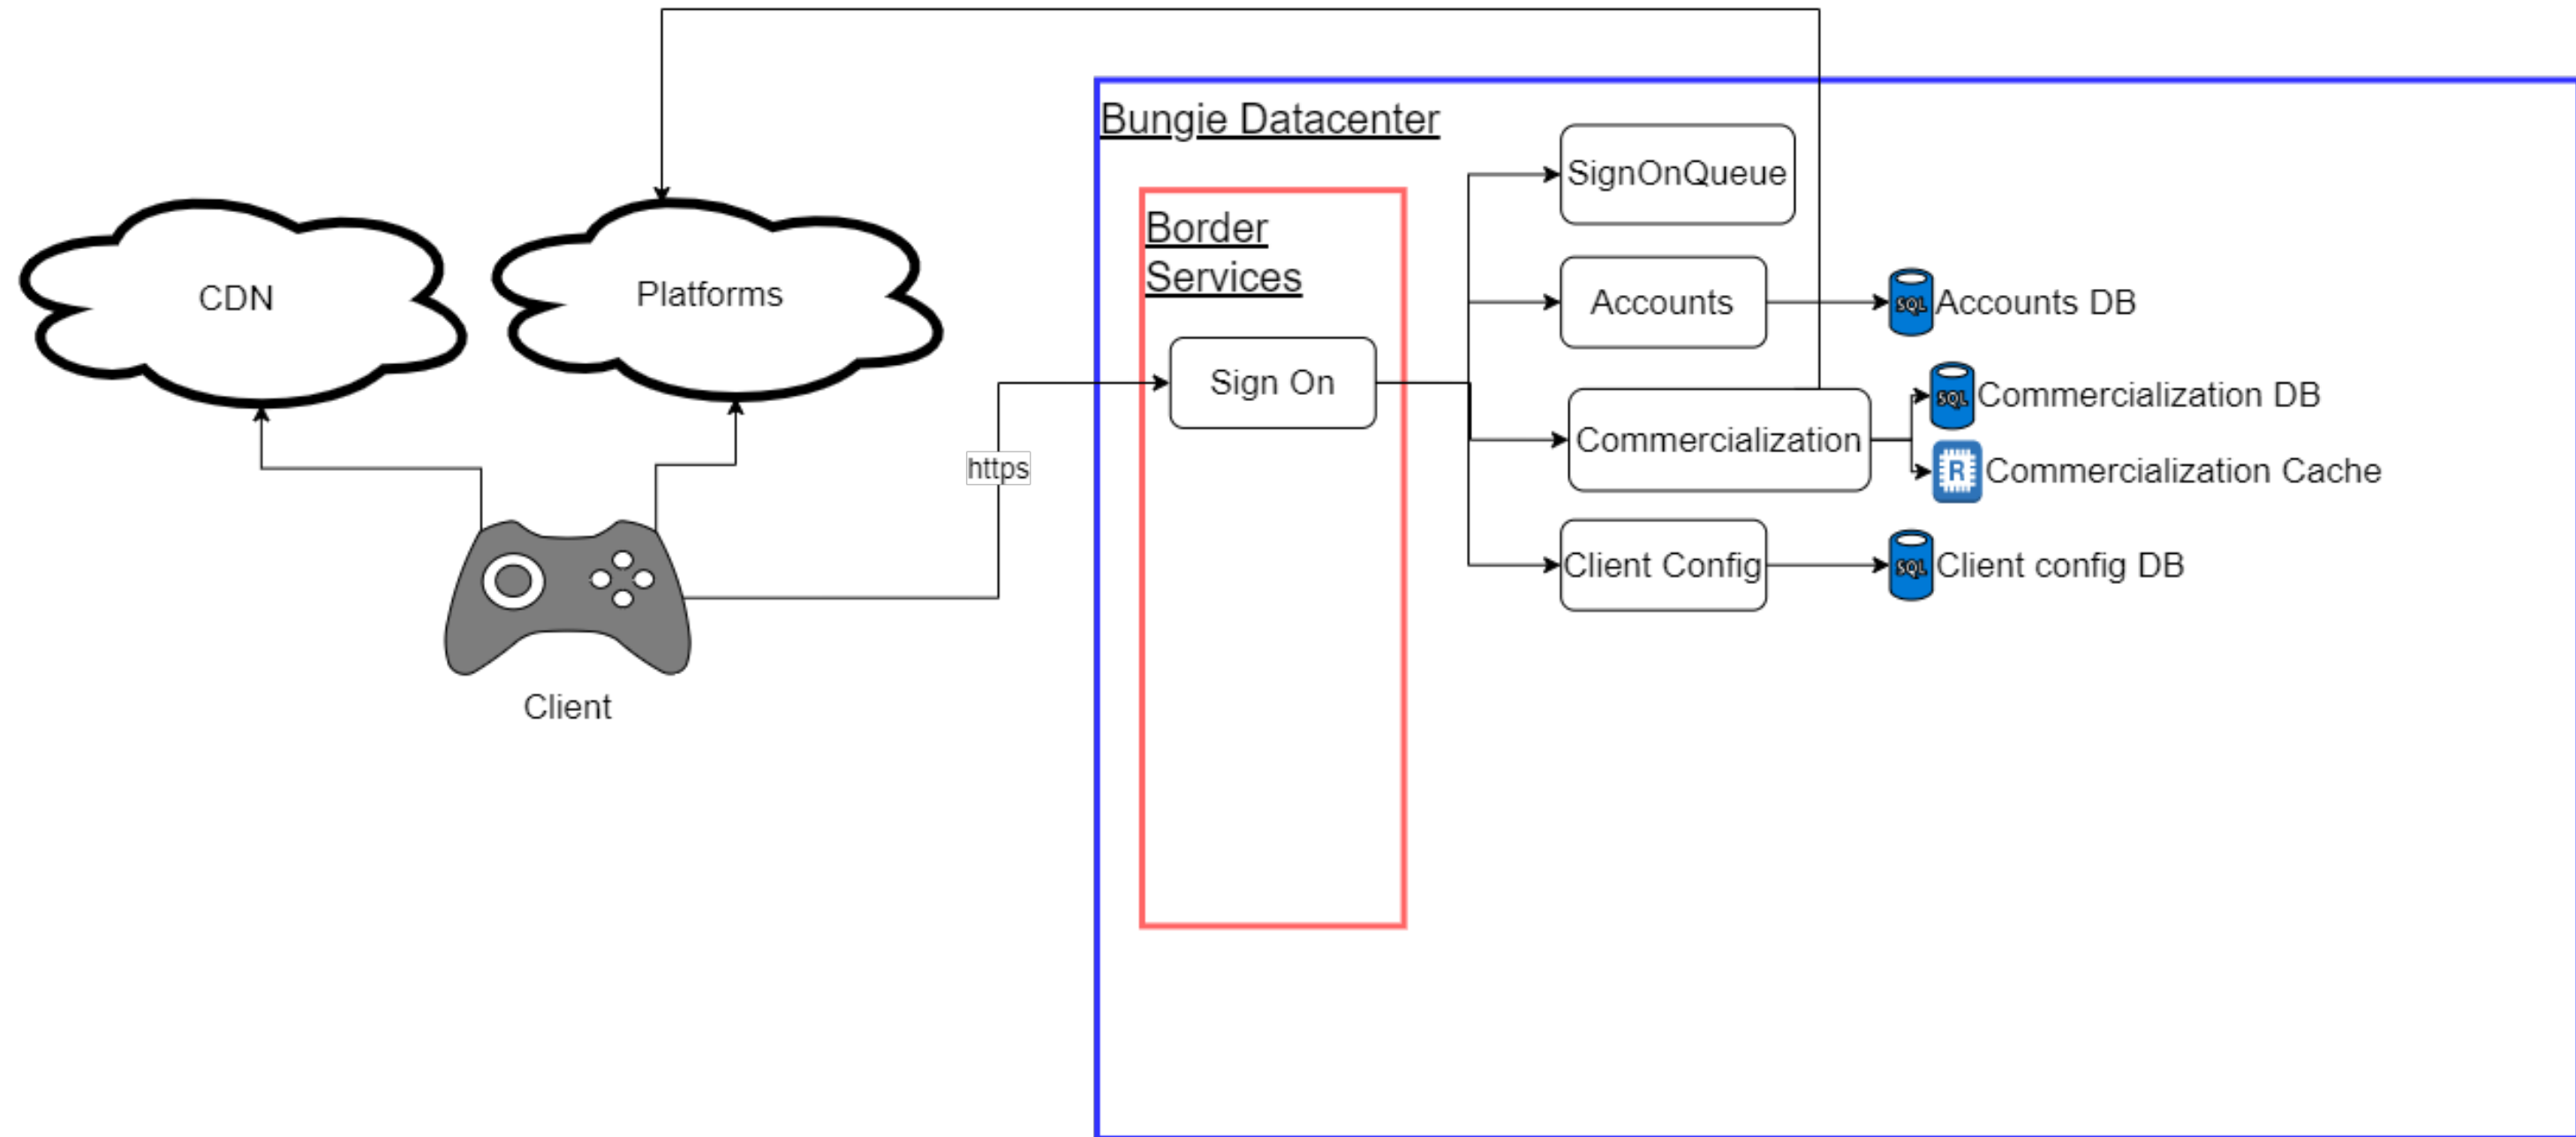


# Why can't I hold all these Microservices?

SignOn	(Queue Pub-Sub)	IPM Locator	Invites
SignOn Queue	Relay Conversation Manager	Mission Control	Chat
Accounts	Nat Relay	(Log Processor)	(Matchmaking)
Commercialization	Clans	(Website Gateway)	PlayerPrivacy
STUN	Activity Host	(Mission Control Gateway)	MonitoringProxy
Client Config	Activity Host Proxy	(Perf Counter Gatherer)	Leaderboards
Service Config	Bubble Host	AWS Elastic Relay	Presence
(Claims)	Bubble Host Proxy	GraphiteS3Relay	Prototype Service
BAP (Front door)	IPM (Intra Process Monitor)	SMS Verification	
World Server	IPM Agent	Friends	

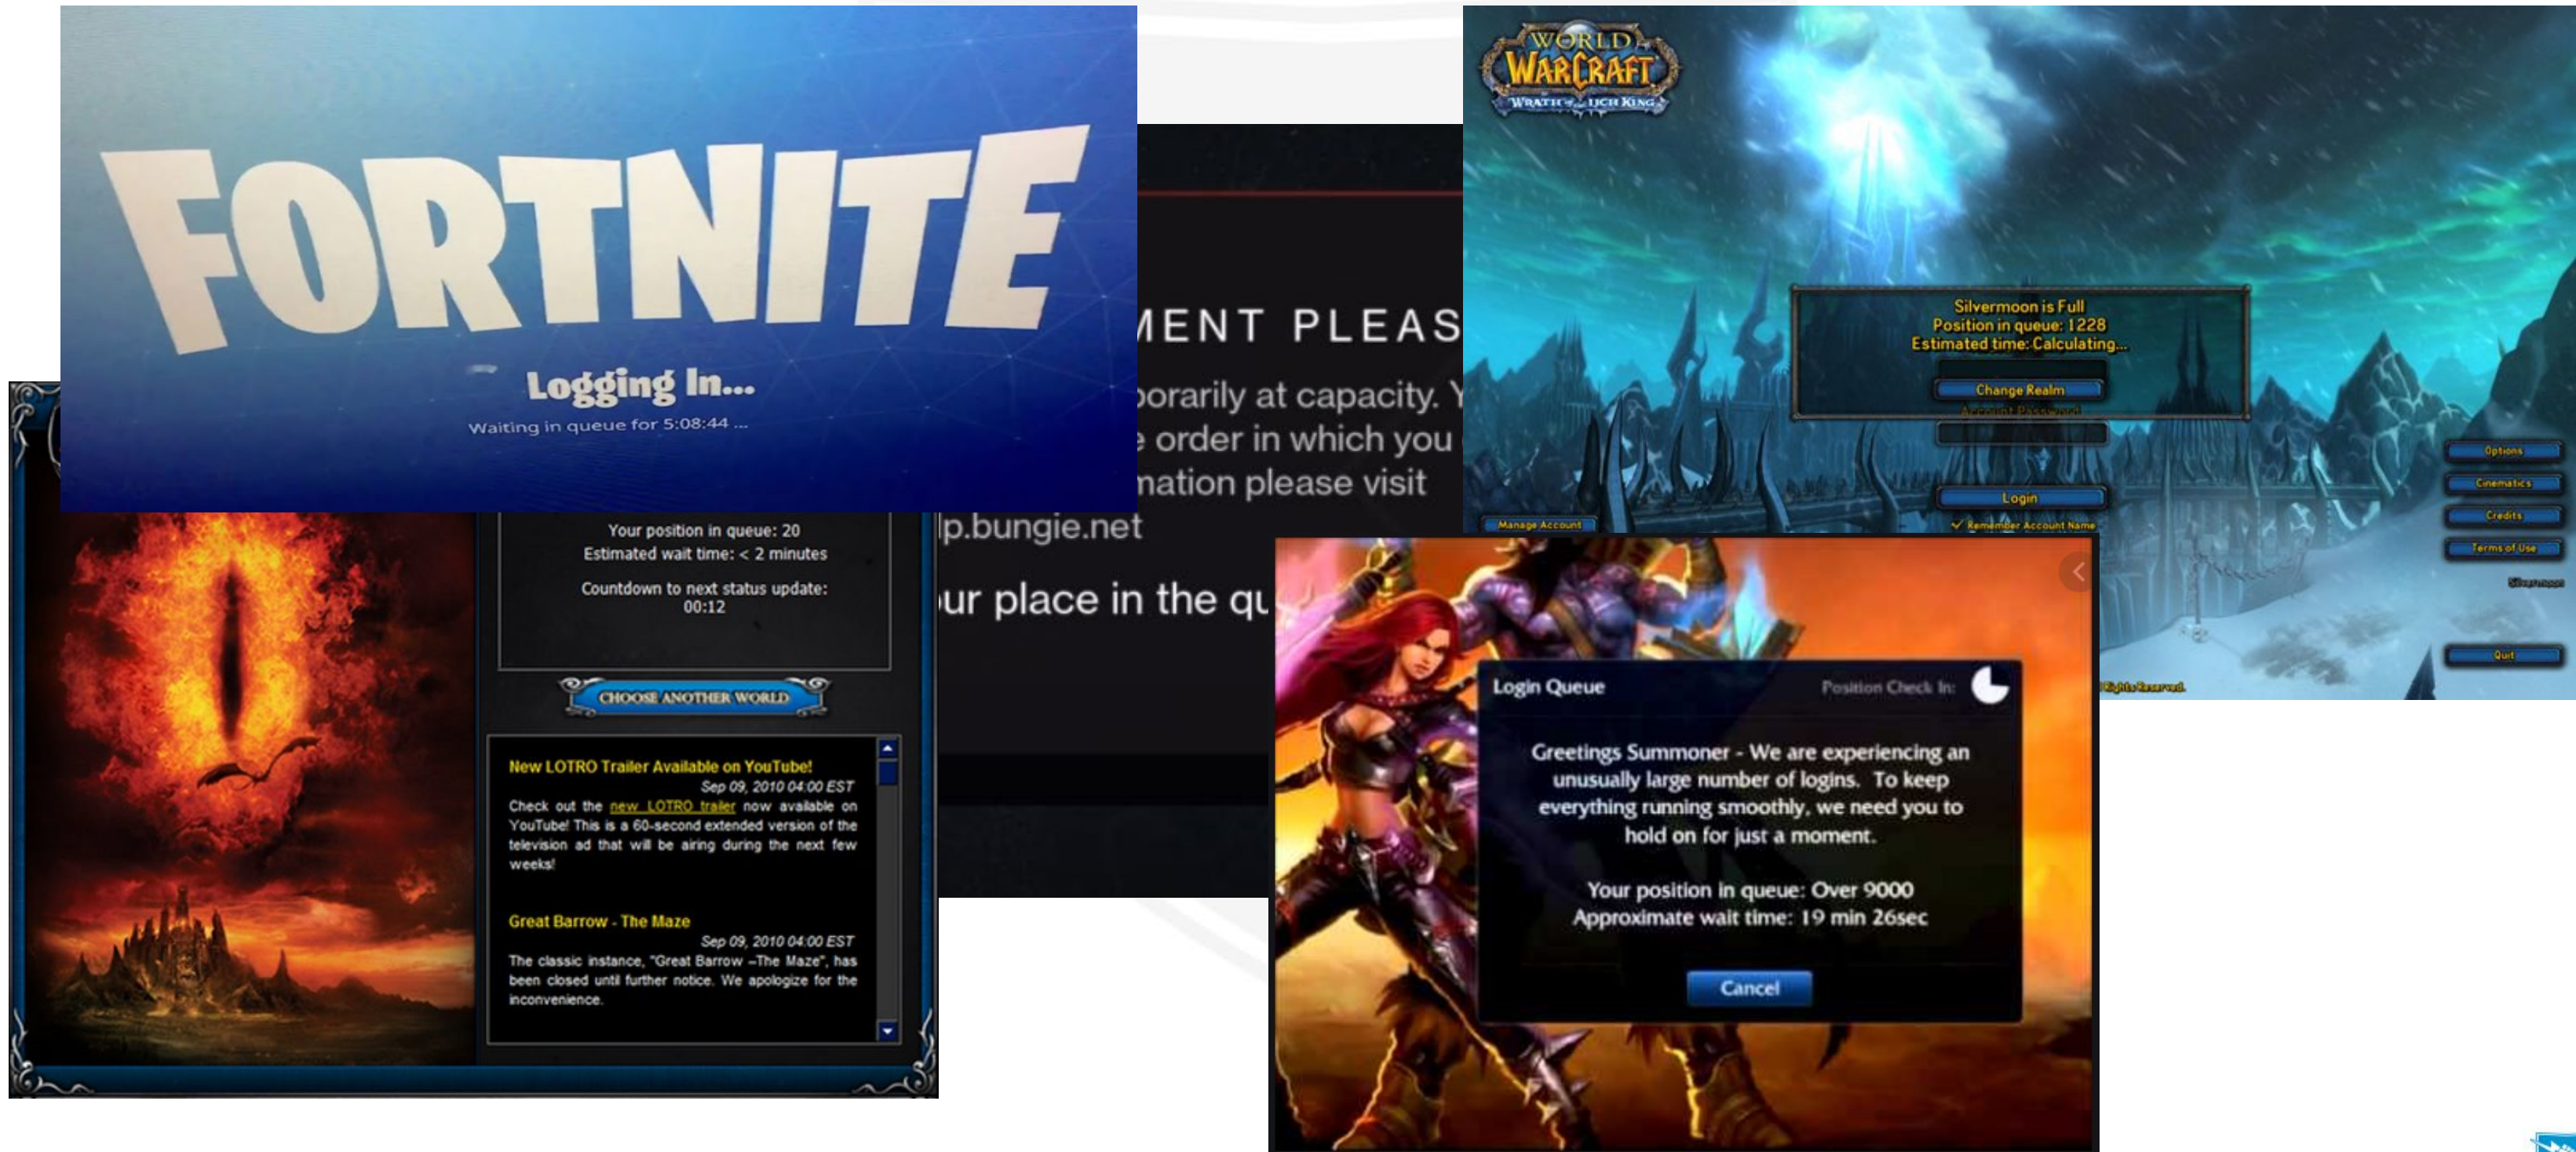


# Bootflow





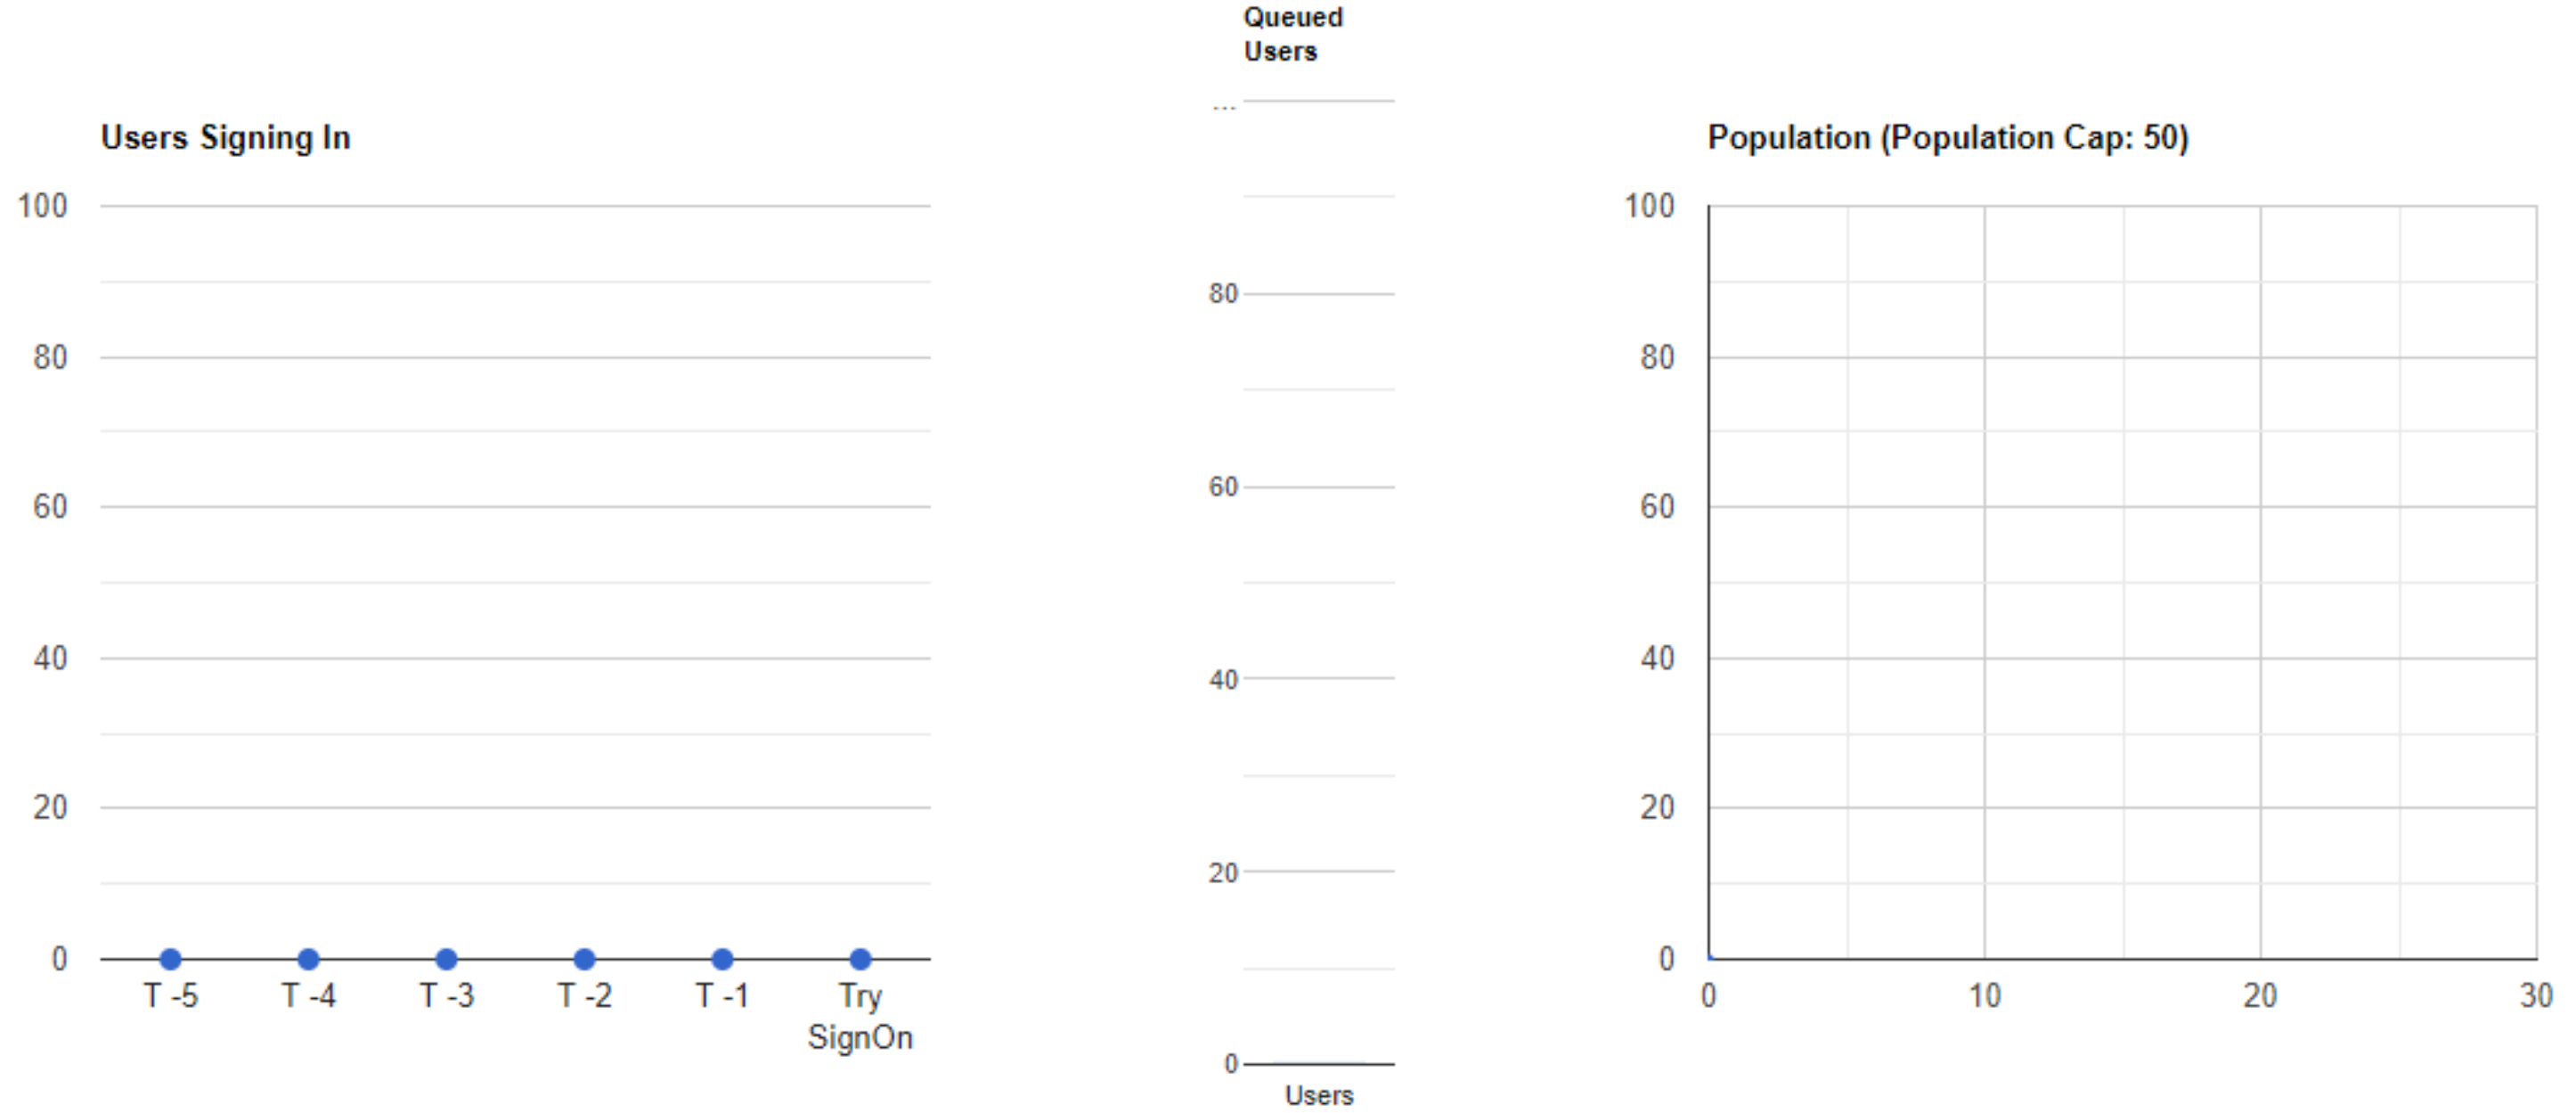
# Sign On Queues and Throttles



# Risks to Online Service Uptime

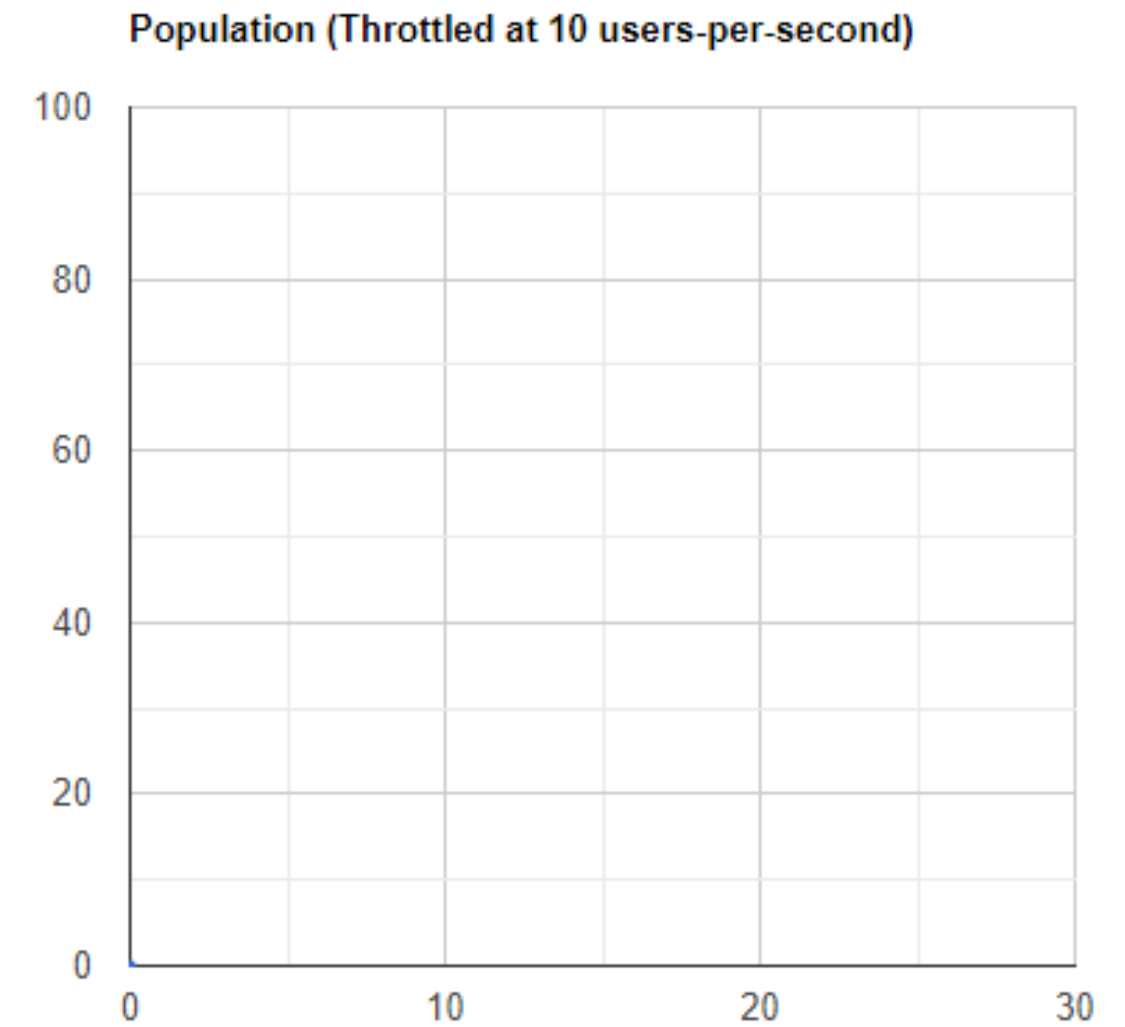
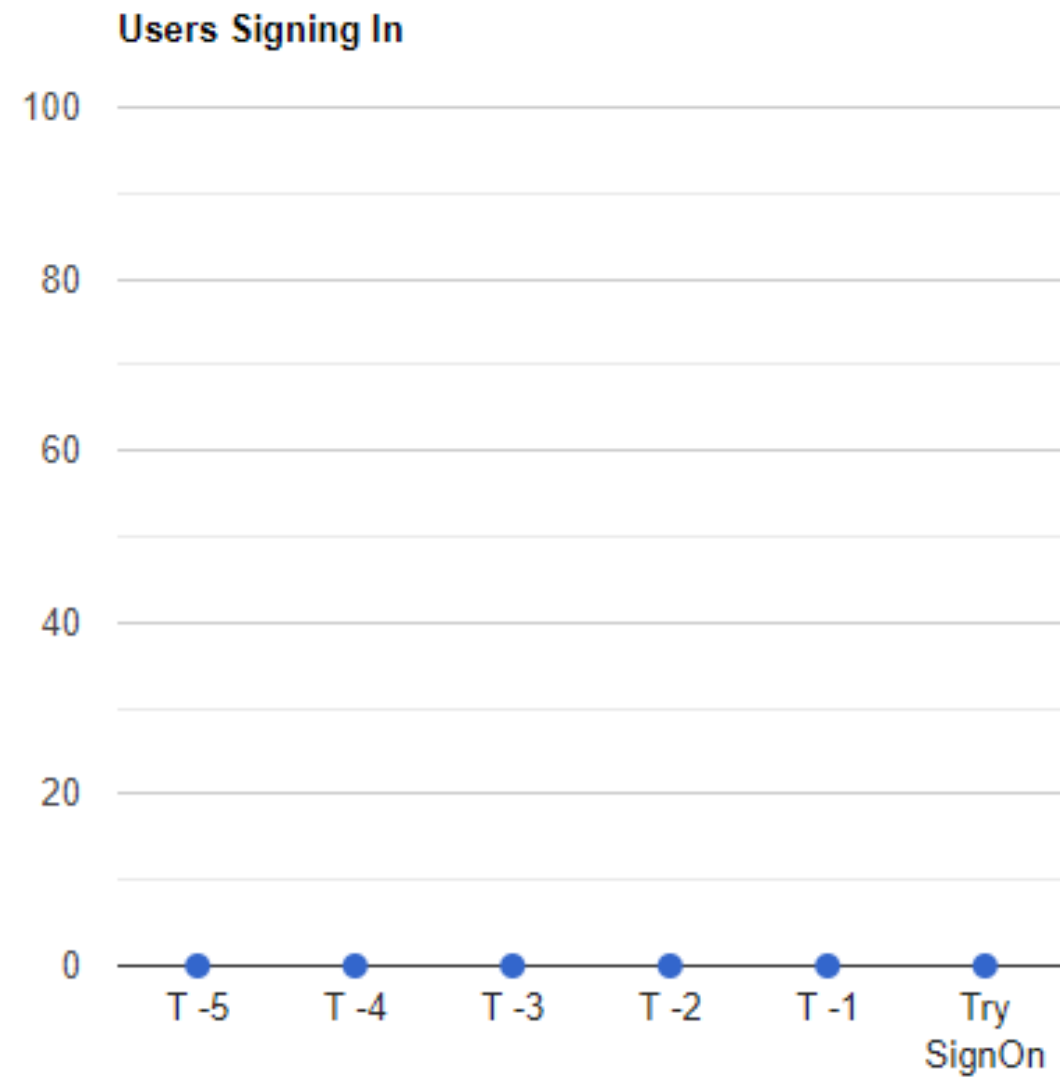
- You have a maximum capacity
- "Thundering Herds" of players
- Hardware failures
- You will likely want planned downtime

# Sign On Queue



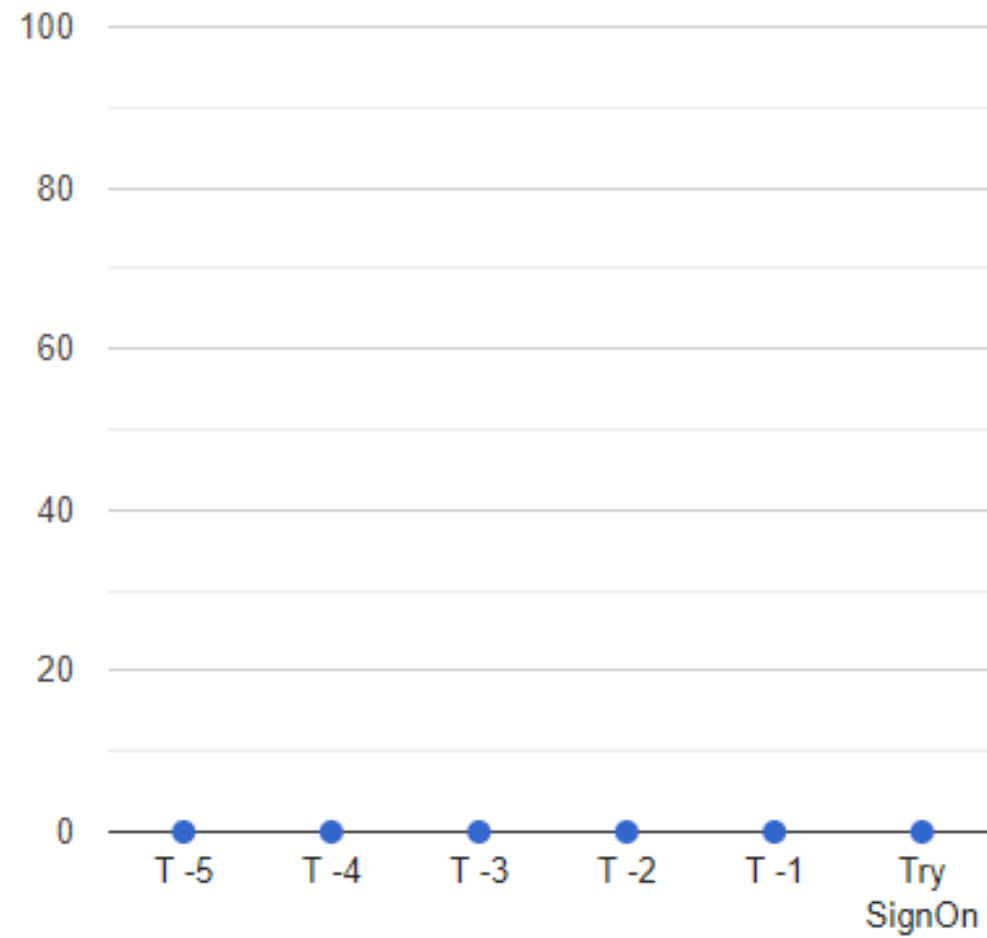


# Sign On Throttle



# Combined!

Users Signing In



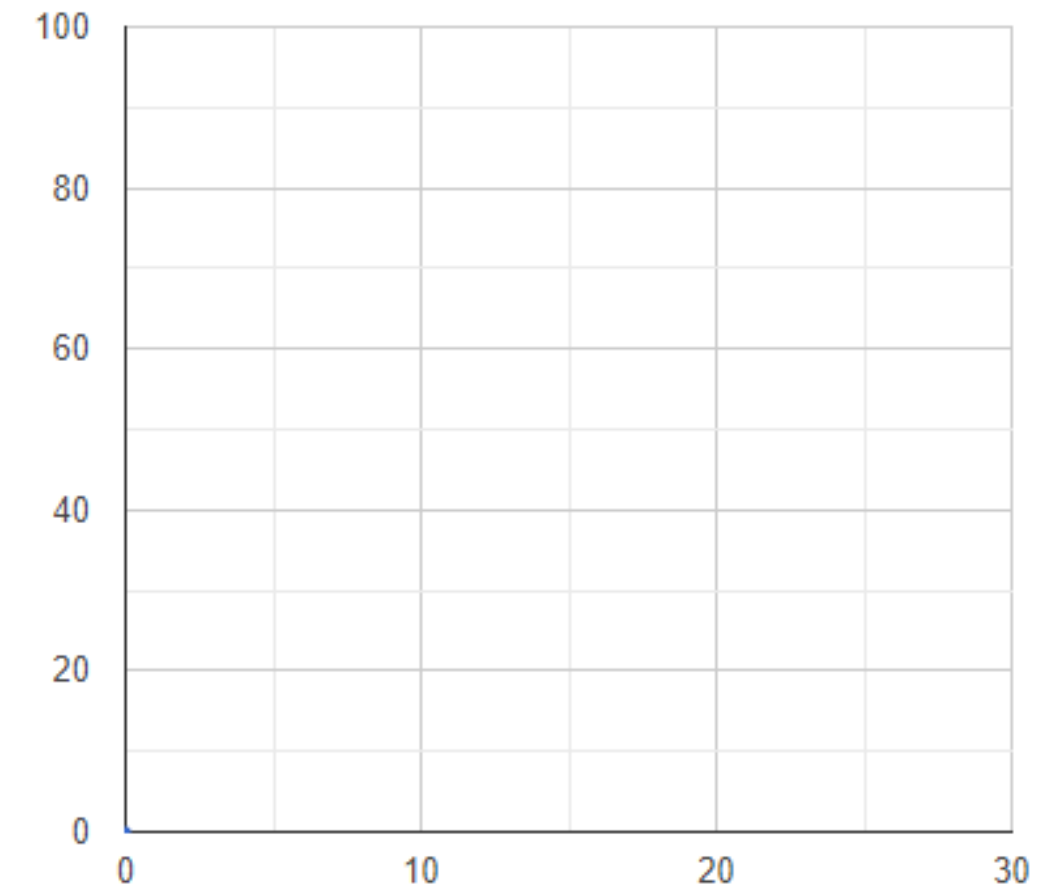
Throttled Users



Queued Users



Population (Population Cap: 50, Throttle: 10 users/sec)



# Queue & Throttle Response Opportunities

- **Responses can control client behavior!**
  - Give the client a token to provide on retry
  - Can send custom messages for the client to display



# Destiny Queue Response

Current Position in Queue	
Message to display to user	
Next time to retry	
Remaining Time Estimate (currently unused)	
Queue Token	Ticket ID ("real" queue position)
	Original entry time
	Last known queue state.

# Destiny Queue Response

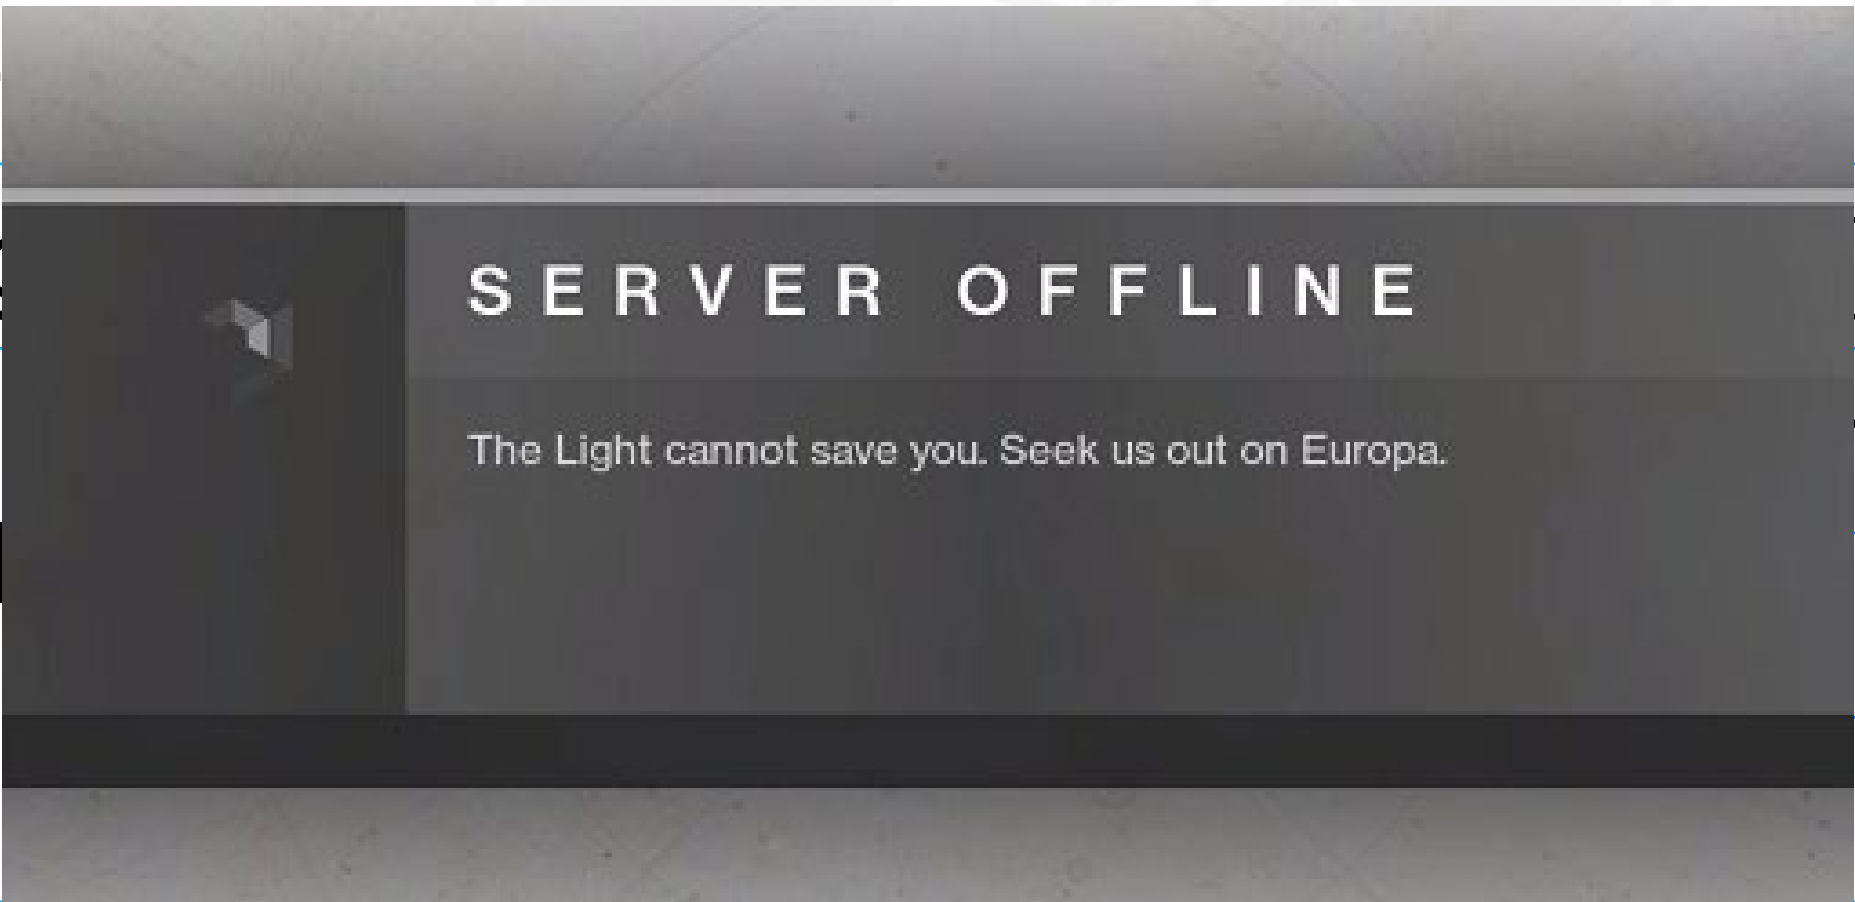
Current Position in Queue

Message to display to user

Next time

Remaining

Queue Tol



)  
tion)

# Destiny Queue Response

Current Position in Queue	
Message to display to user	
Next time to retry	
Remaining Time Estimate (currently unused)	
Queue Token	Ticket ID ("real" queue position)
	Original entry time
	Last known queue state.

# Best Practices

- **Limit your queue & throttle dependencies**
- **Put the queue and throttle in front of everything, including authentication**
- **Have an allow-list for test accounts**

# Best Practices (Continued!)

- **Use the queue as the main gating for your game**
- **Default population cap to 80% of your known capacity**
- **If an issue is happening:**
  1. Clamp the population cap to 0
  2. Let people drain from the game (or kick them)
  3. Then slowly ratchet the population cap back up



# Key Takeaway #1

- **Sign on Queues and Throttles are one of your first and best tools to handle and prevent Service issues**

# Into Orbit!



# BAP Server (Bungie Access Protocol)

- **Stateful gateway to Destiny's service layer**
- **Communication on TCP**
- **Securely encrypted using the token provided from the Sign On service.**

# WorldServer

- **Character data is worked on in-memory**
- **One of our few stateful services**
- **A given server can handle ~5000 accounts**
- **Hosts C++ game logic DLLs that run on player data**

# Claims Service ("ClaimZ")

- **Acts a simple routing service to the stateful Worldserver**
- **Redis Backed**

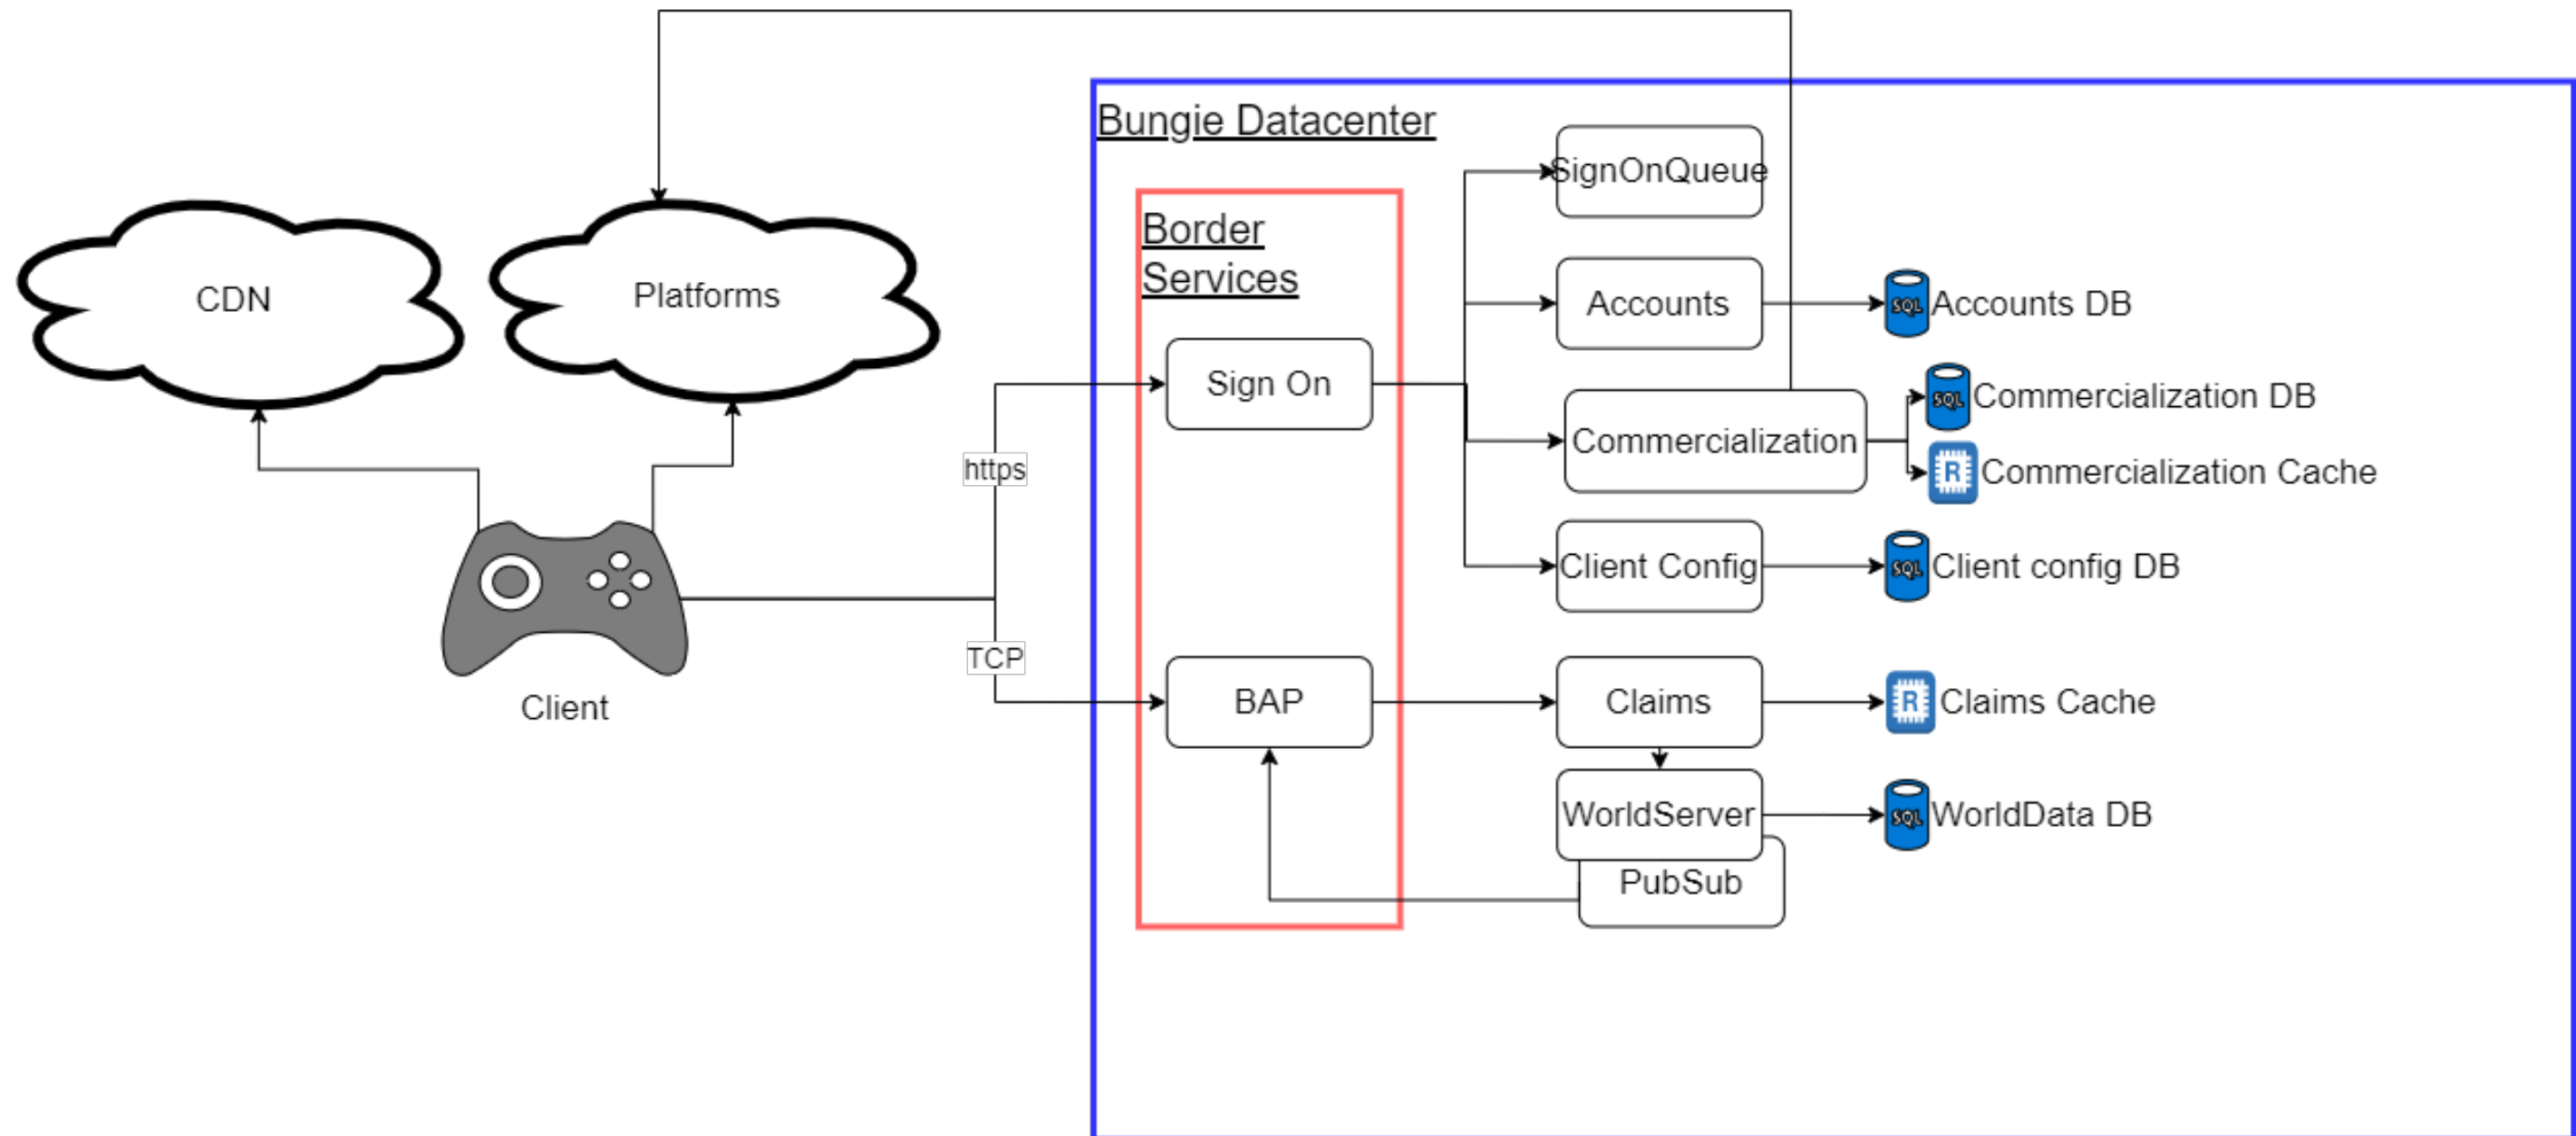




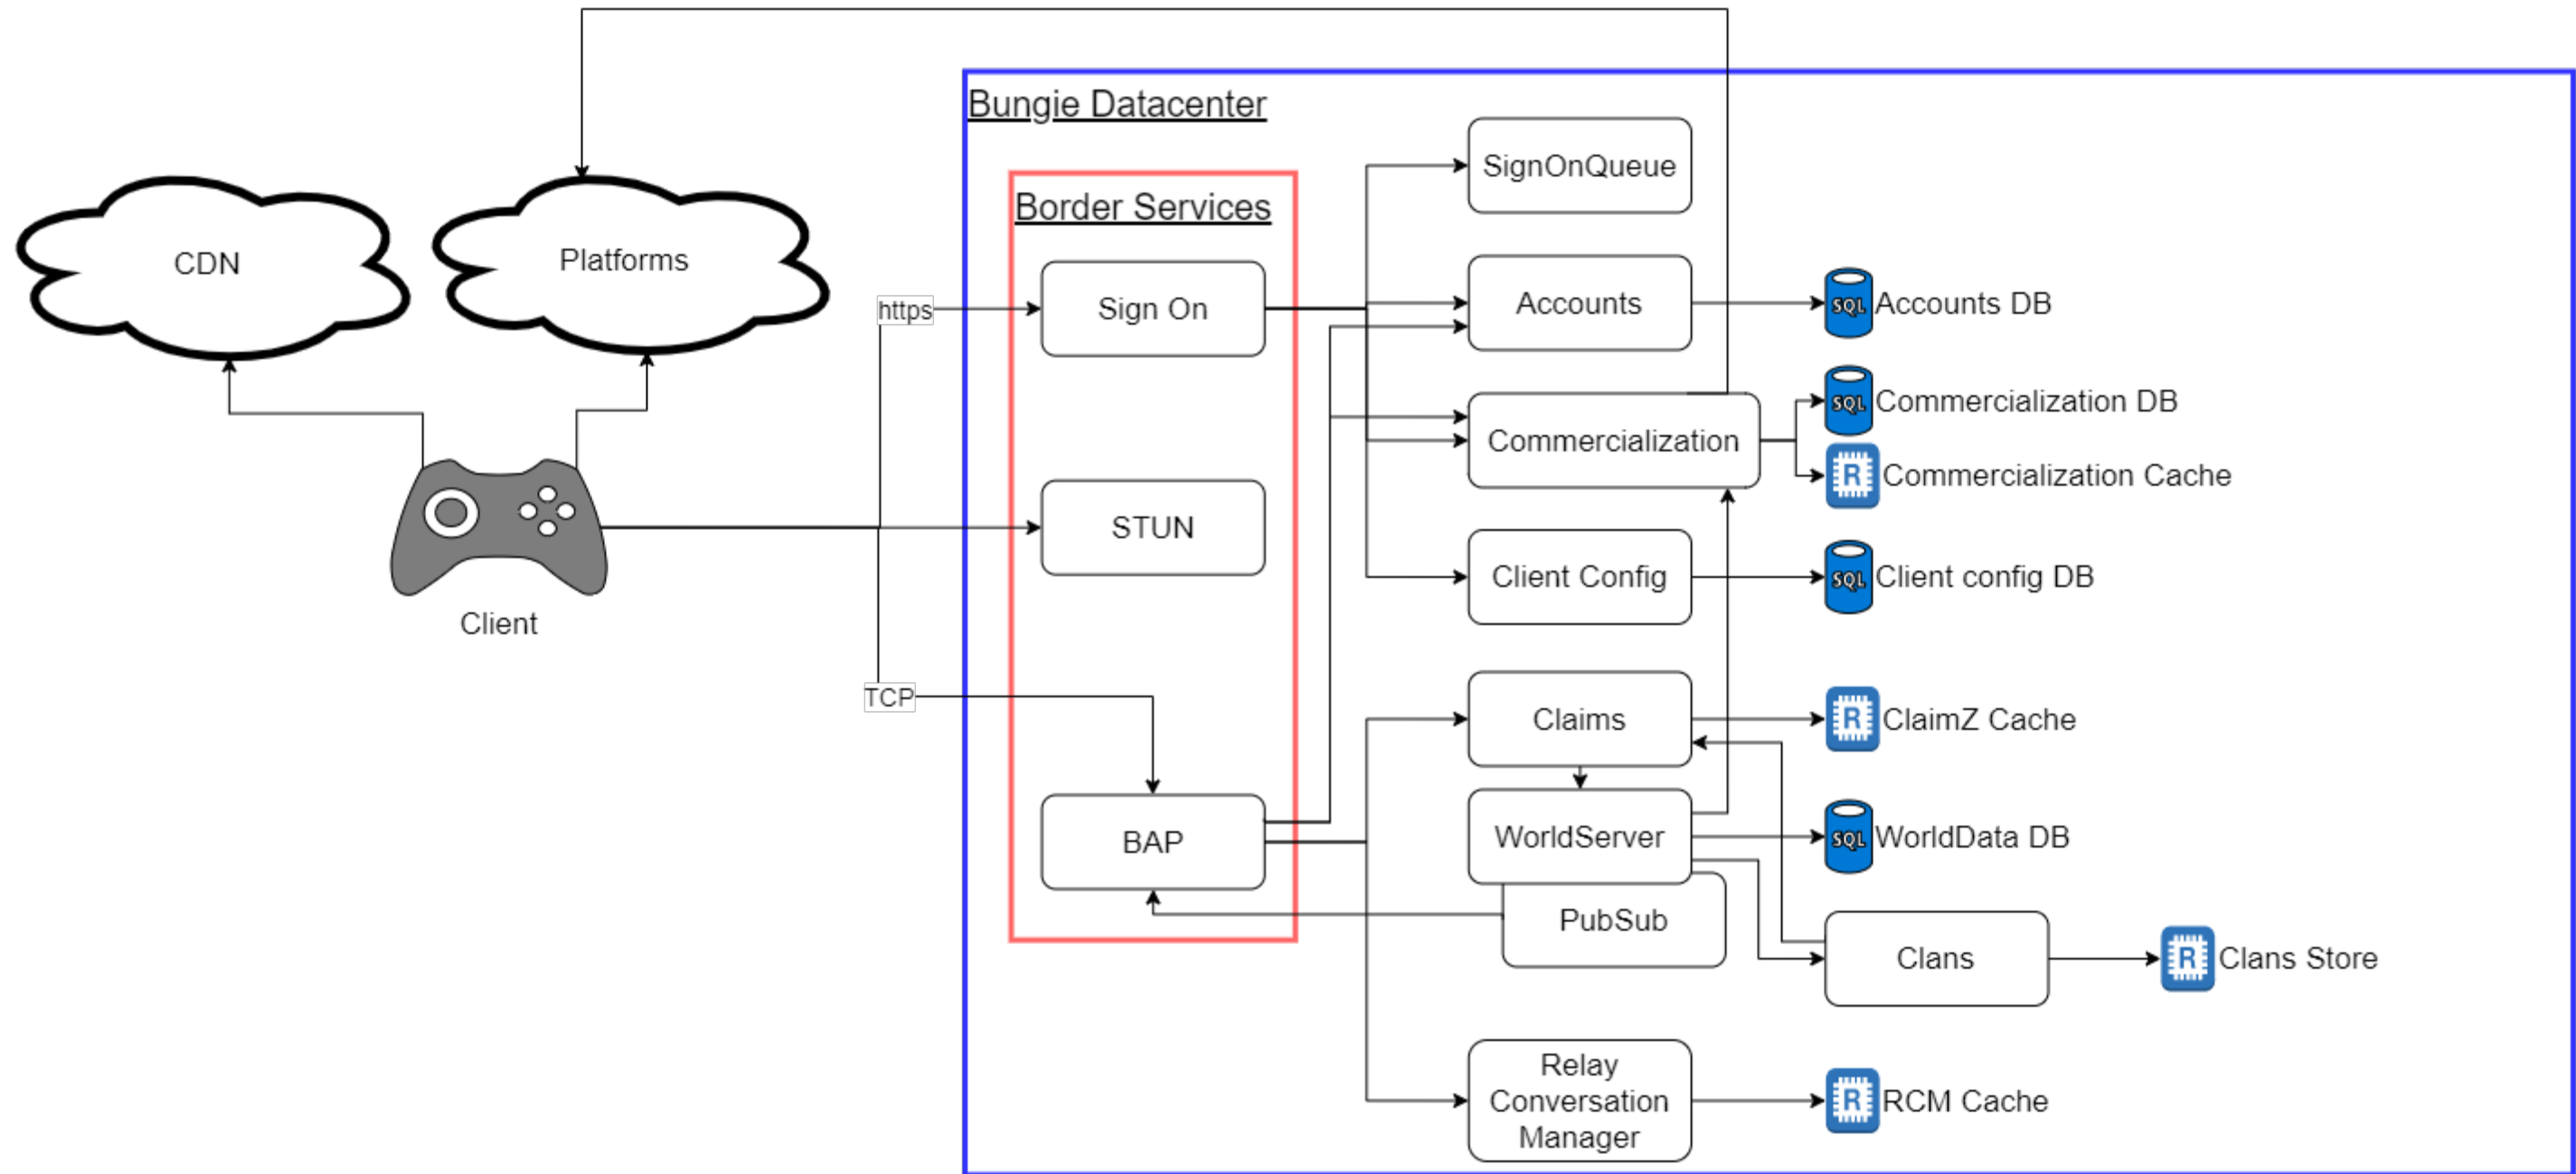
# Character PubSub Service ("QueueZ")

- **A subcomponent of WorldServer**
- **PubSub system with in-order differential updates**
- **Can subscribe to different levels of detail on any character**
- **A given client is subscribed to itself, party members, friends, clanmates**
- **Uses ZeroMQ**

# Into Orbit!



# Into Orbit!



# Character Storage





# Character Storage - Schema

AccountCharacters	
AccountId	CharacterId
1	10
1	11

AccountCharacterAttributes		
CharacterId	CharacterAttributeld	CharacterAttributeValue
10	101	1111
10	102	2222
11	101	1111
11	102	4444

CharacterItems	
CharacterId	ItemId
1	10
1	10

CharacterItemAttributes		
ItemId	ItemAttributeld	itemAttributeValue
1000	201	1234
...	...	...

# Character Storage - Issues

- **Heavy duty loads**
  - 6000 account attributes
  - 4000 character attributes per-character
  - 25 attributes per item
  - 3 characters + 500 account items = 30,000+ rows!

# Character Storage – Issues (Continued)

- **Analytics queries were difficult**
- **SQL fixups were terrifying**
- **Low savings from incremental updates**

# Character Storage V2

- **A New blob-store model**
  - Accounts blob
  - Per-Character blobs
  - One blob for all items
- **Binary blobs encoded using protocol buffers**
- **Still stored in SQL**

# Character Storage V2 - Results

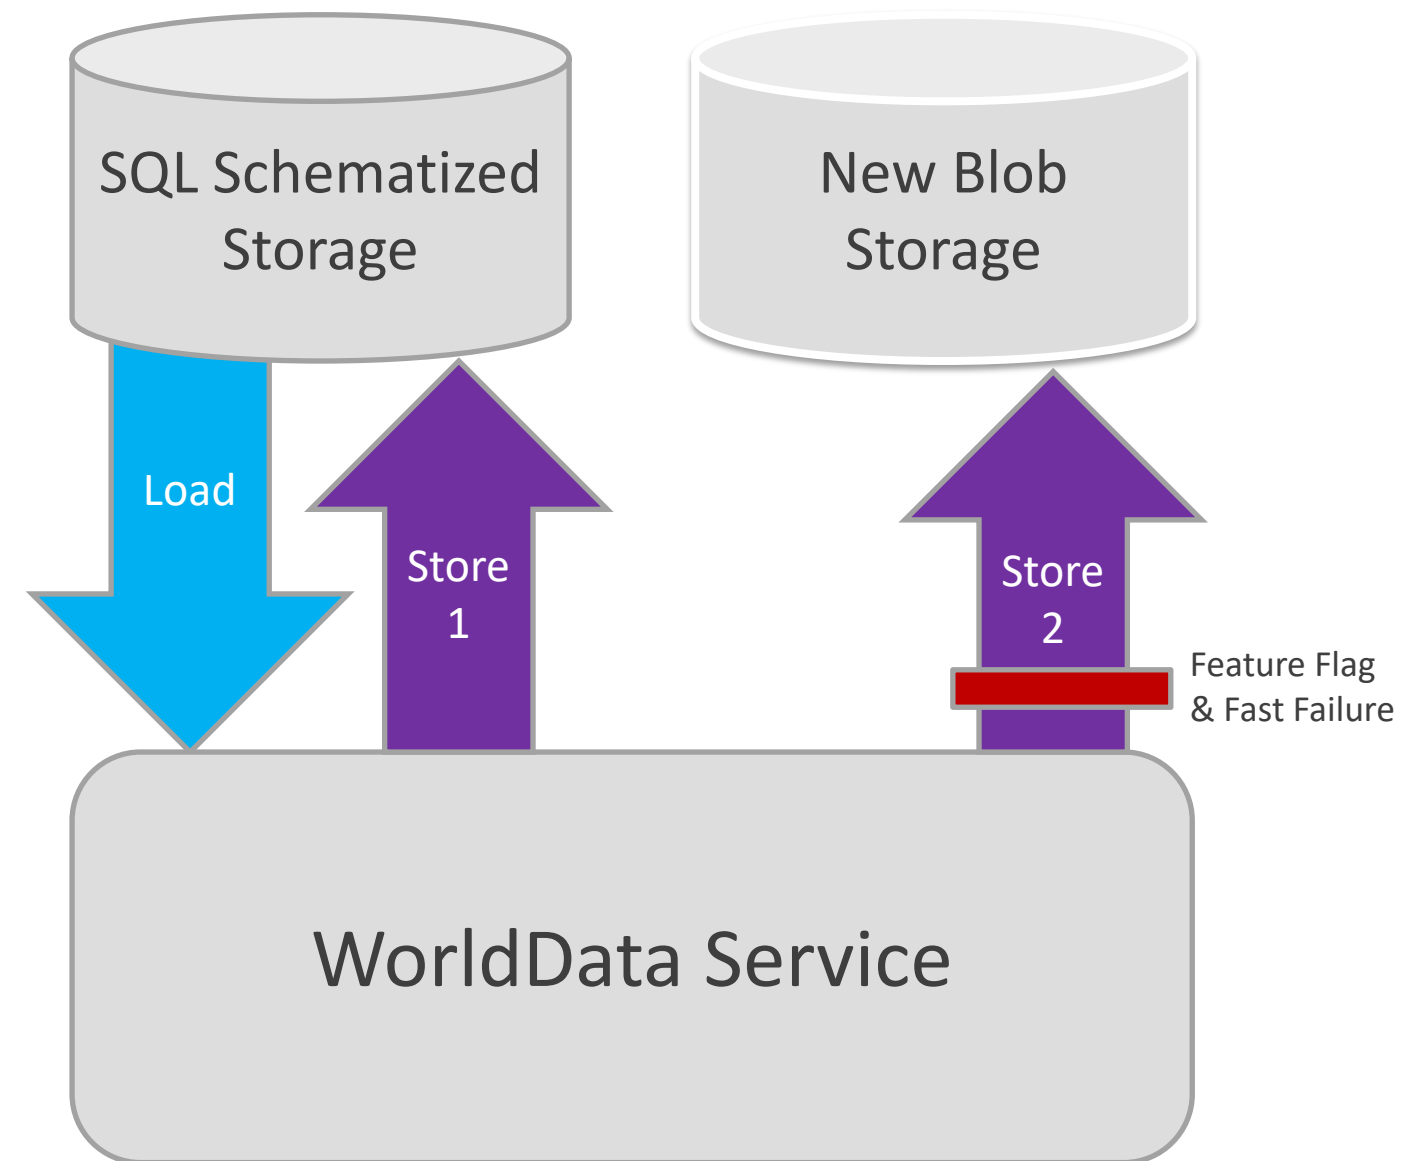
- **90% savings in space, and improved load times**
- **Easily stores 200,000+ “rows” for an account**
- **New debugging functionality made possible**
- **No major conversion downtime impact for users!**
  - **?!?**



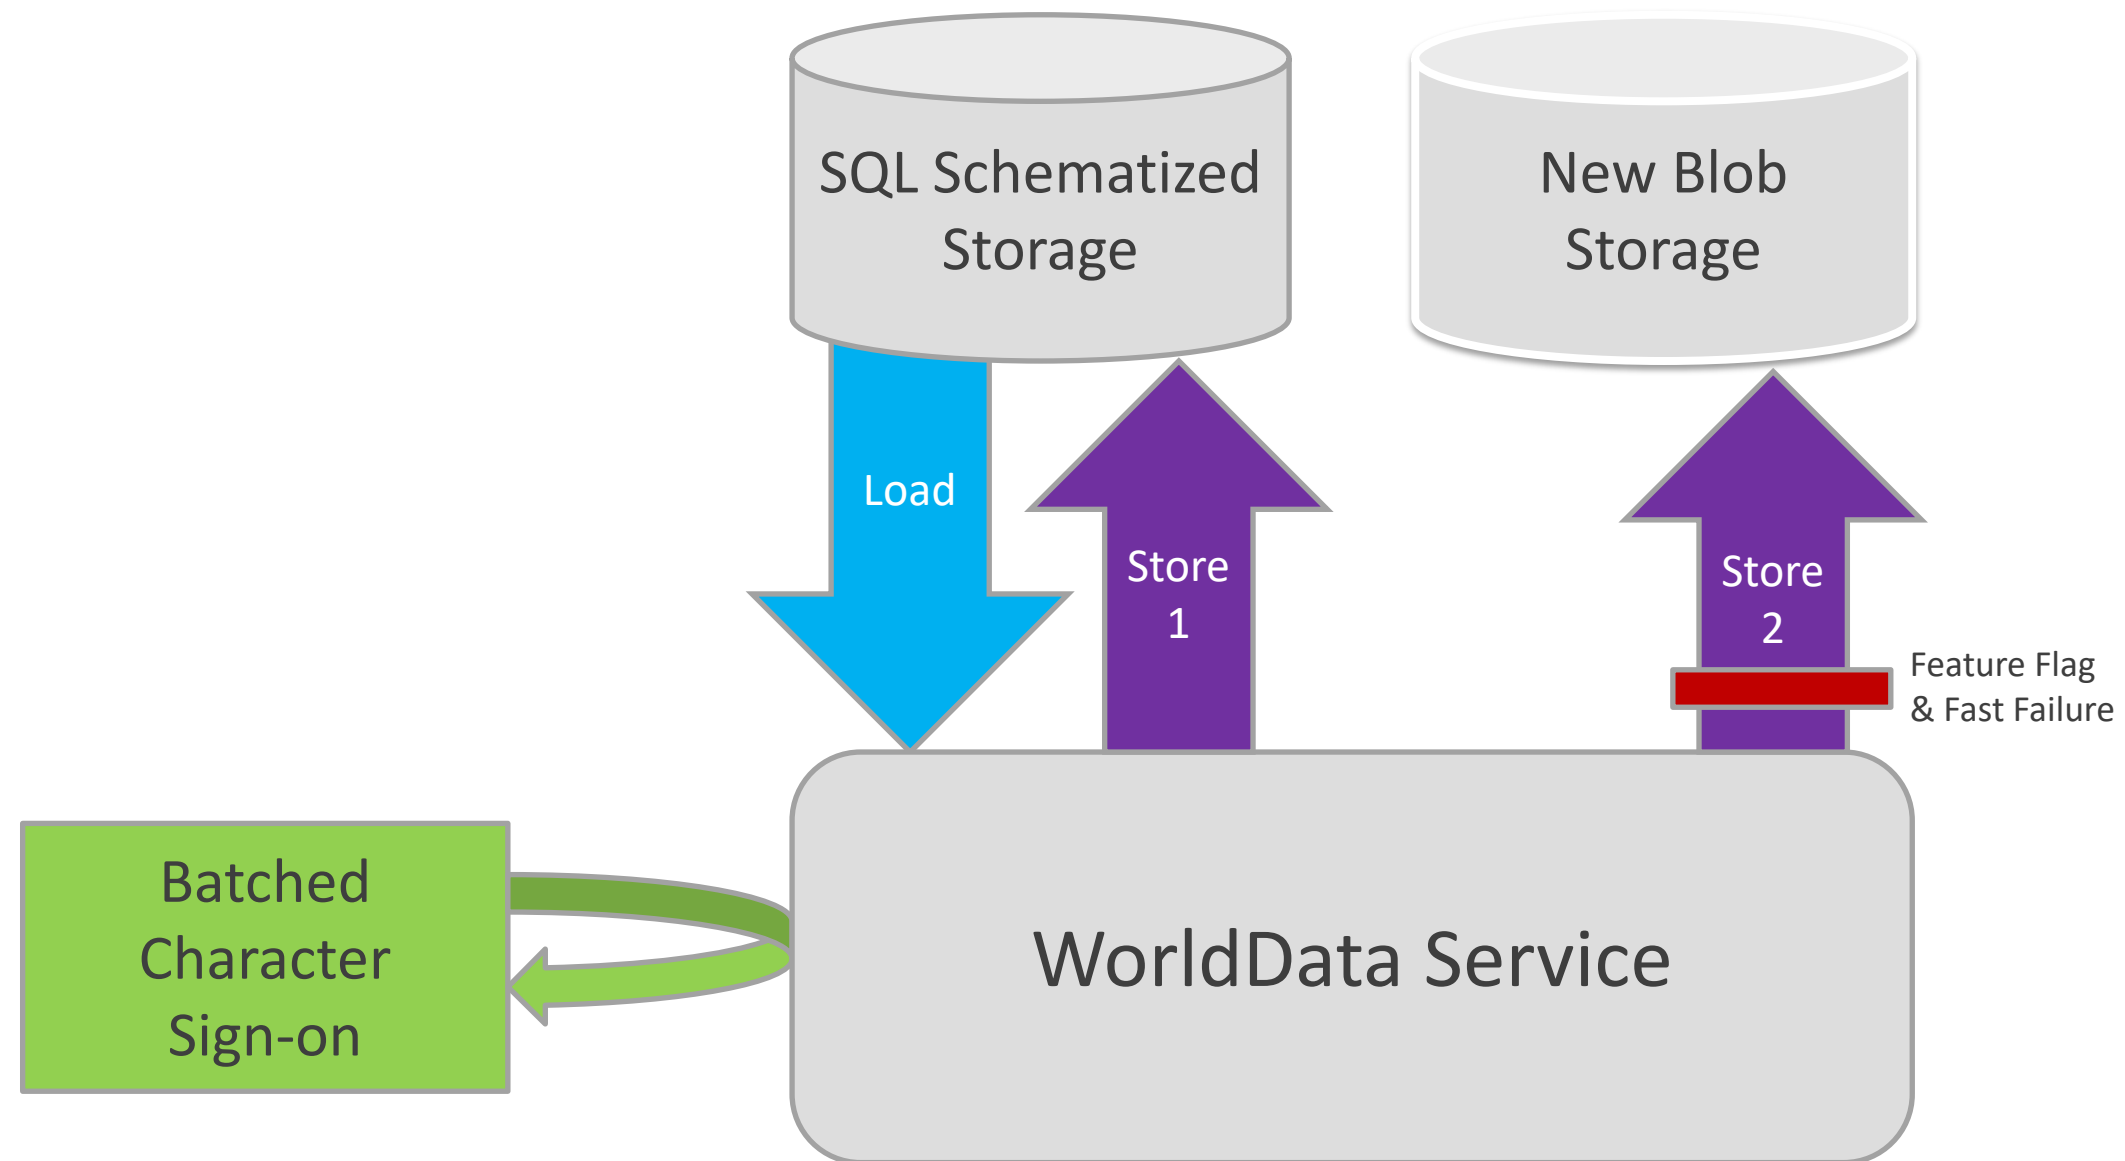
# Migration best practices



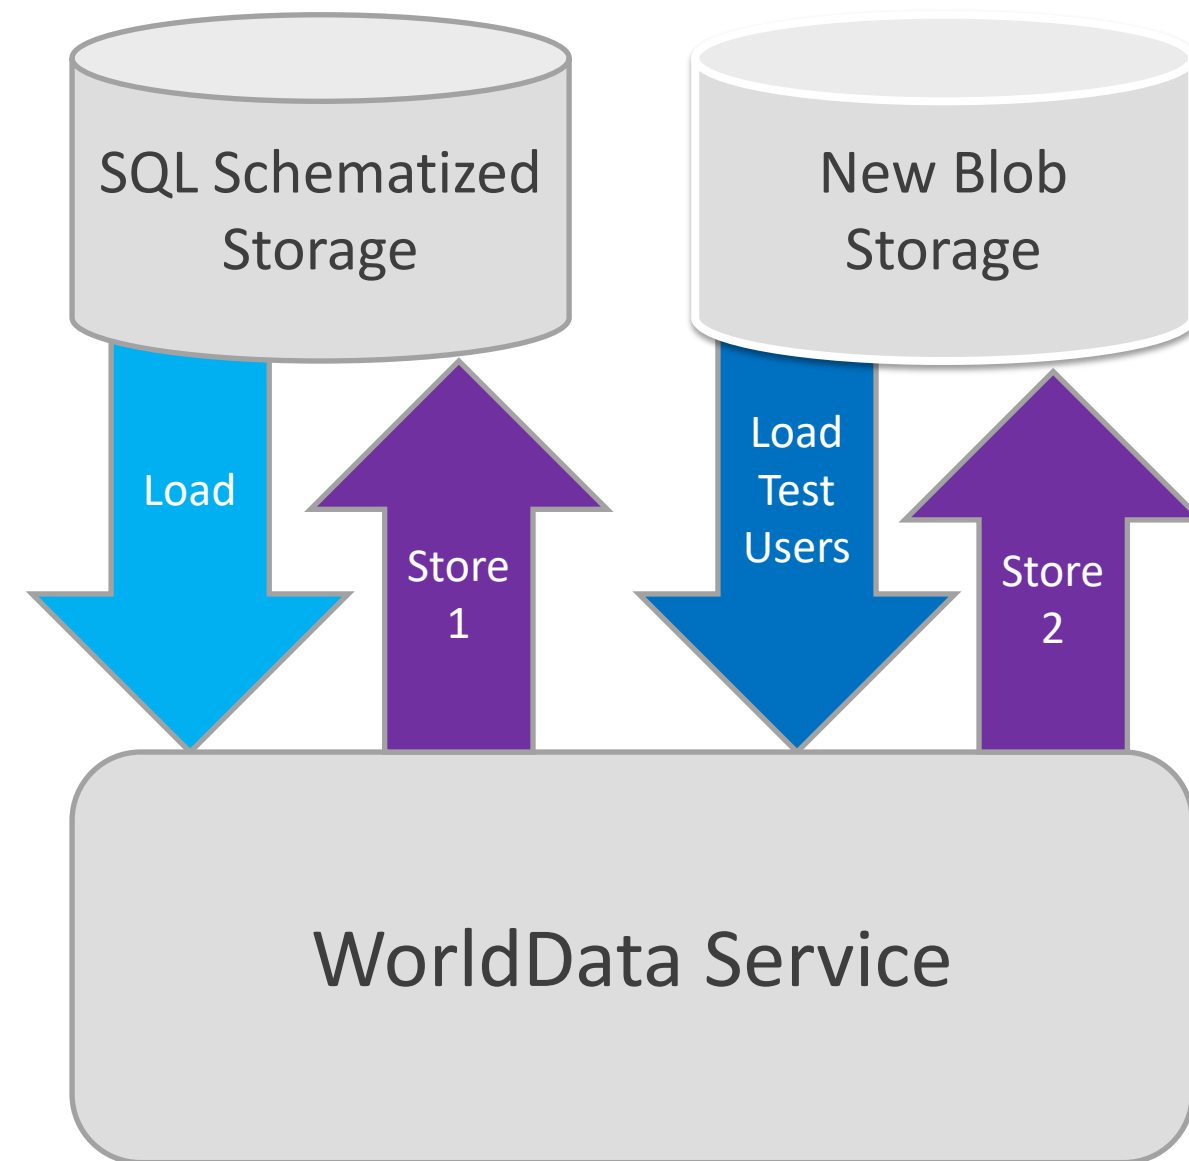
# Migration best practices – Step 1



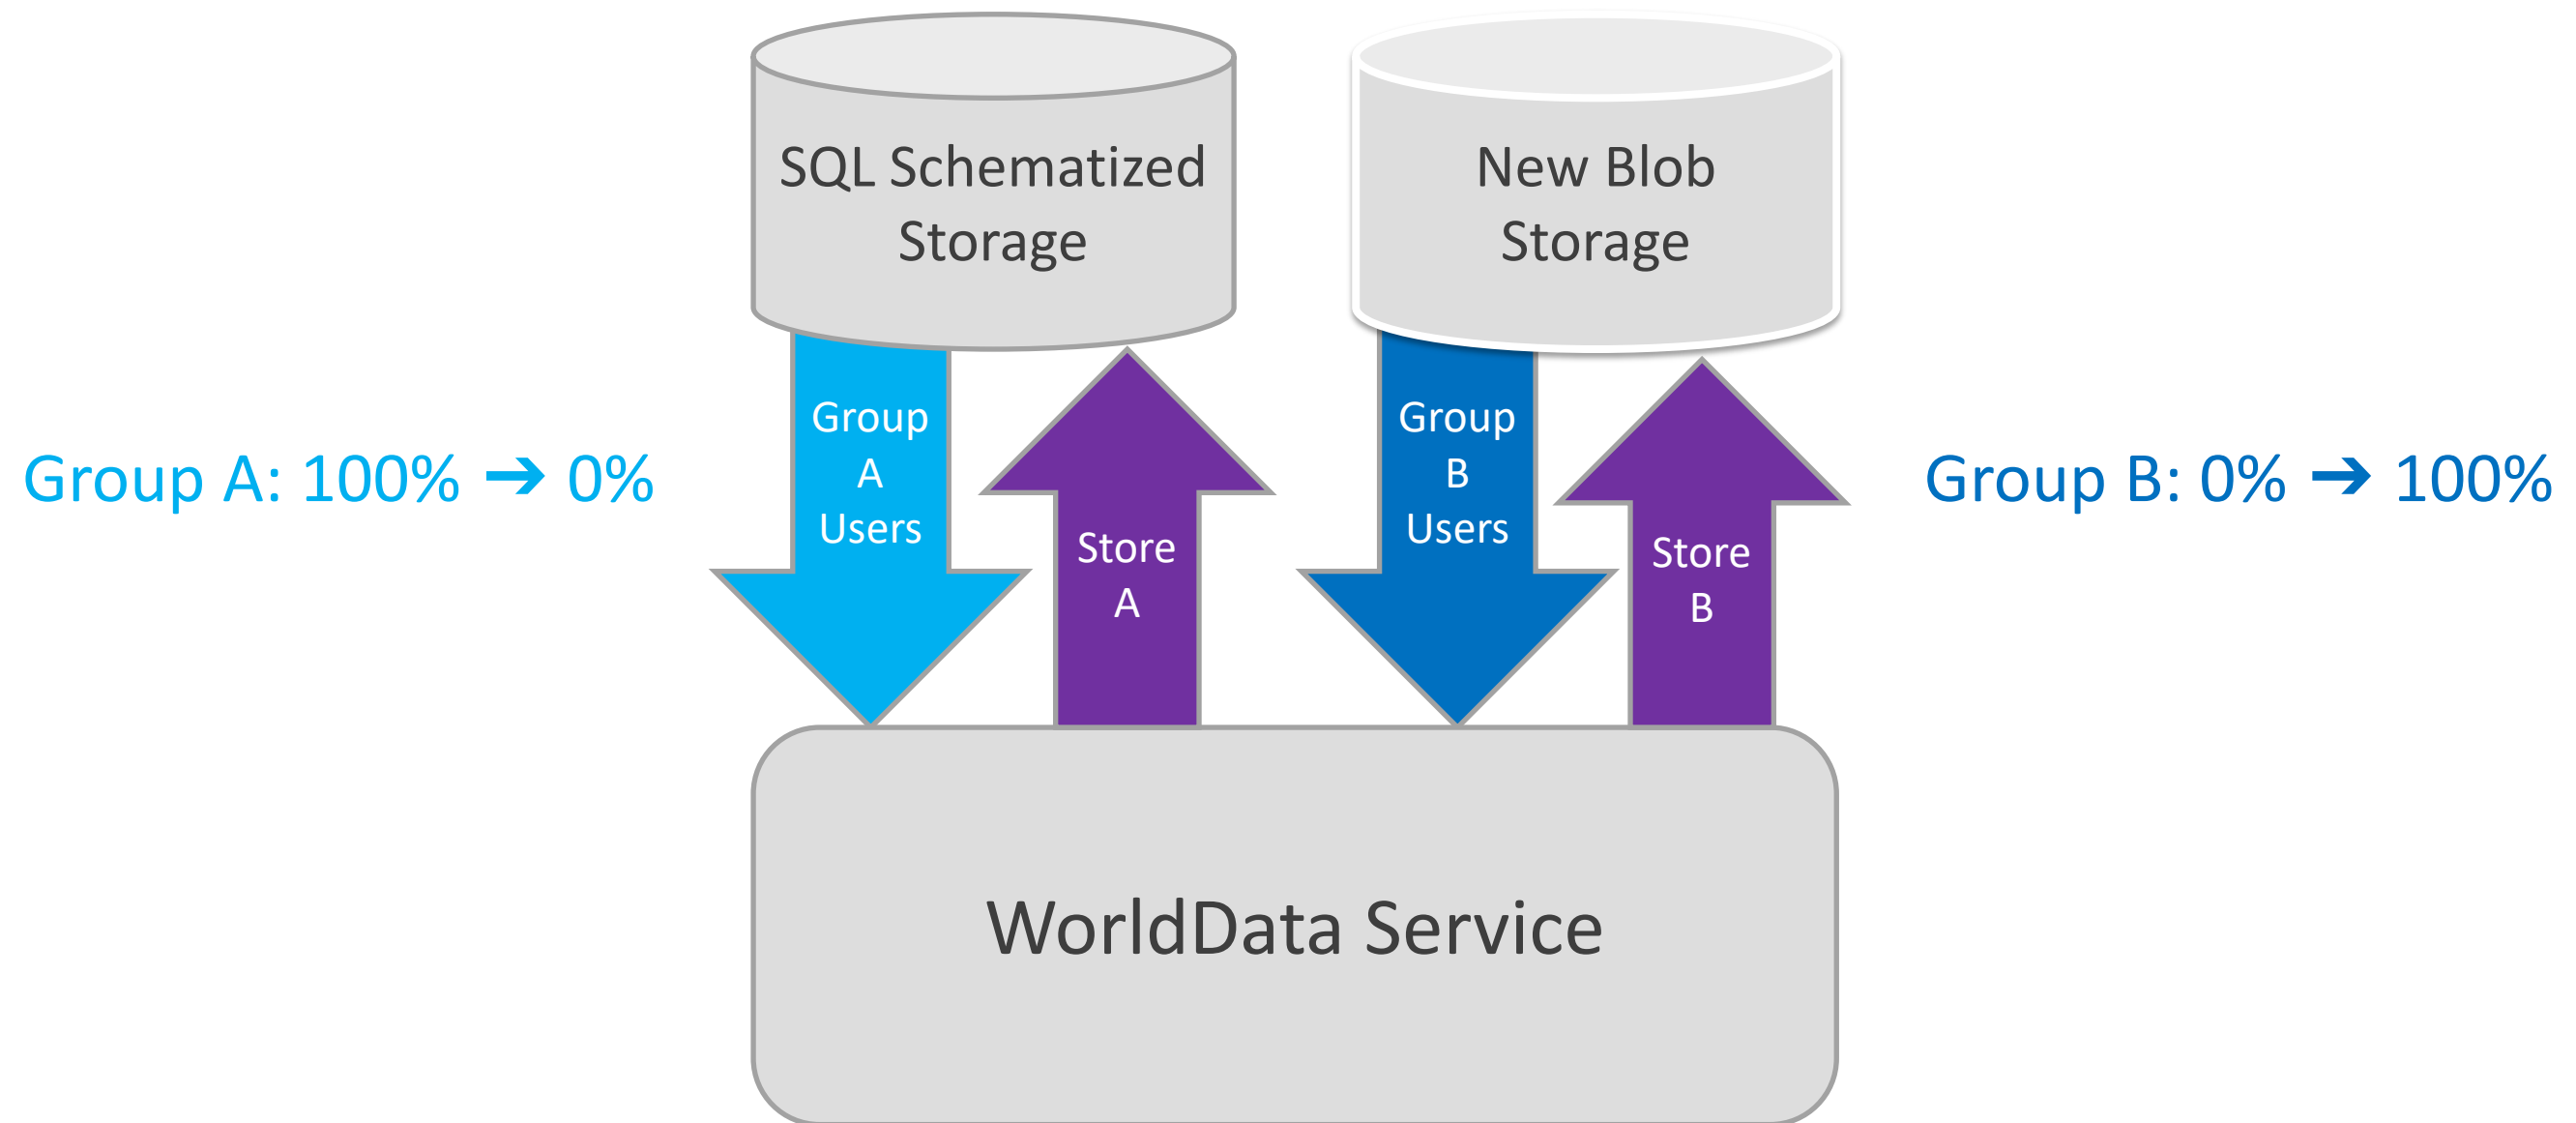
# Migration best practices – Step 2



# Migration best practices – Step 3

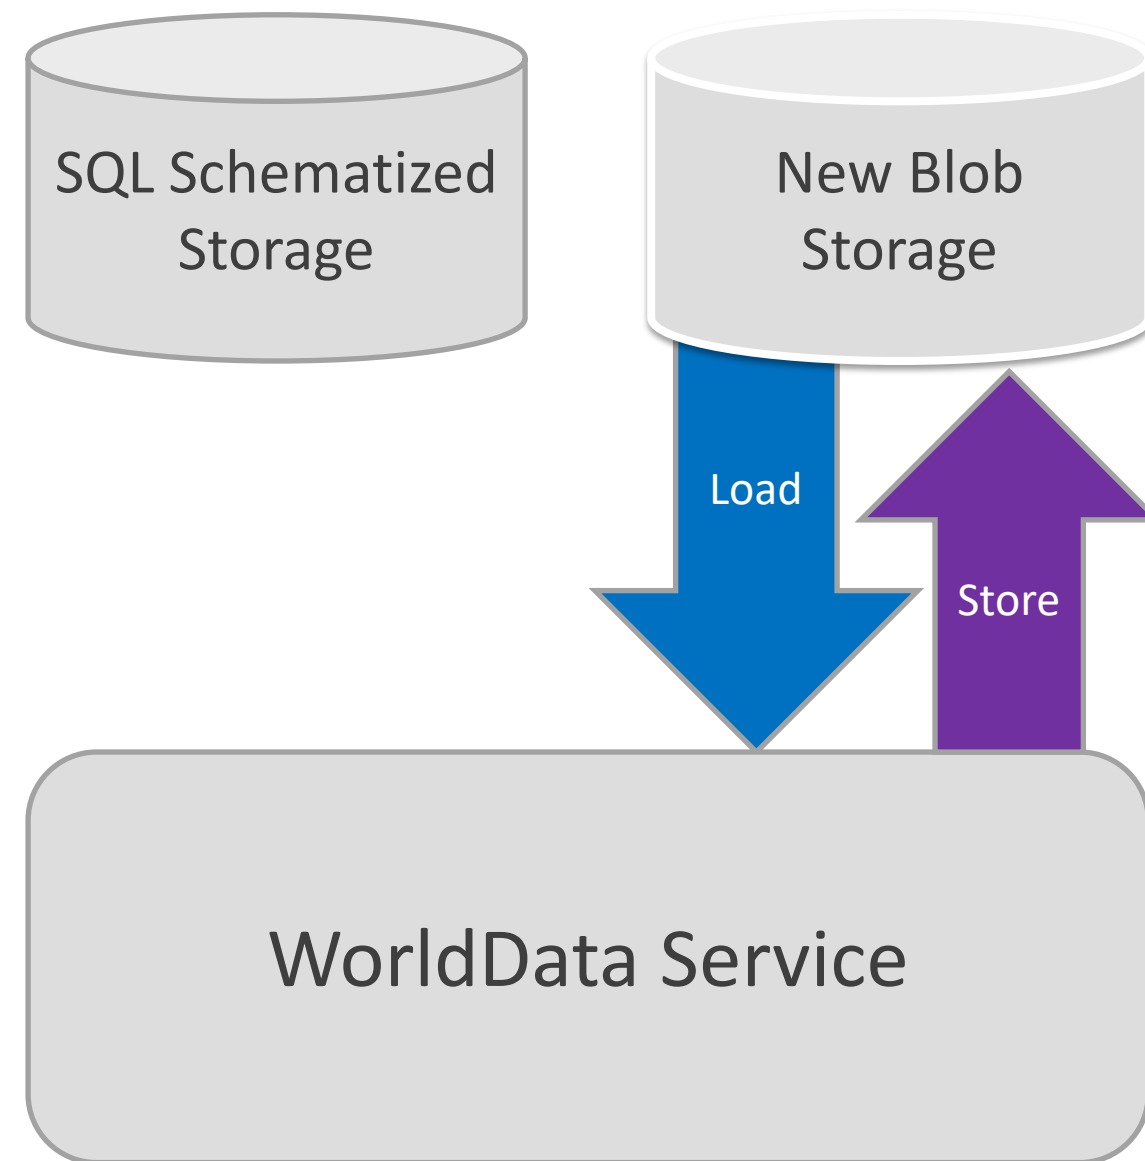


# Migration best practices – Step 4





# Migration best practices – Step 5



# Migration best practices - Generalized

- **Write to old and new systems**
- **Make sure all accounts have written to both**
- **Set test accounts to read from new**
- **Slowly ramp the population from 1% to 100%**
- **Turn off the old version**

## Key Takeaway #2

- **Whenever possible, launch new systems side-by-side with old systems, and slowly cut over**



# Data Reliability





# Data Reliability

- **Character data is very important to persist**
- **If a player disconnects, we persist their data**
- **But what about if the WorldServer crashes?**
- **Constant write-through super-expensive**
- **Change logging is super-complex**



# Data Reliability - Solution

- **Persist character data every 5 minutes.**
- **Do a bonus persist on "important" events**
  - Like getting an exotic item
- **If a WorldServer dies, player loses at most 5 minutes of progress.**

## Key Takeaway #3

- **Understand the real reliability requirements for your systems and be skeptical of 100% targets**

# Clans System

- **Clan data can have many simultaneous writes**
- **Used a stateless model with optimistic concurrency**



# Clans Optimistic Concurrency Model

- 1. Receive clan action from WorldServer**
- 2. Load clan data and clan data version**
- 3. Run action against clan data**
- 4. Attempt to persist.**
- 5. If the stored version is different from the persisted version, go back to step 2 (Retry N times)**

# Clans System Results

- **In practice, this worked really well**
- **Used Redis for the persistent store due to read/write rate capabilities**
- **Extra consideration – Different member content**
  - When do you version the clan data?



# Activity Hosts & Bubble Hosts





# Activity Hosts & Bubble Hosts

- **“Shared World Shooter: Destiny's Networked Mission Architecture”**
  - Justin Truman, GDC 2015
- **Act as the game script & physics hosts for the game**

# Activity Hosts & Bubble Hosts

- **4 services**
  - Activity Host (AH),
  - Activity Host Proxy (AHP)
  - Bubble Host (BH),
  - Bubble Host Proxy (BHP)
- **1 Proxy per machine, many Activity & Bubble hosts**

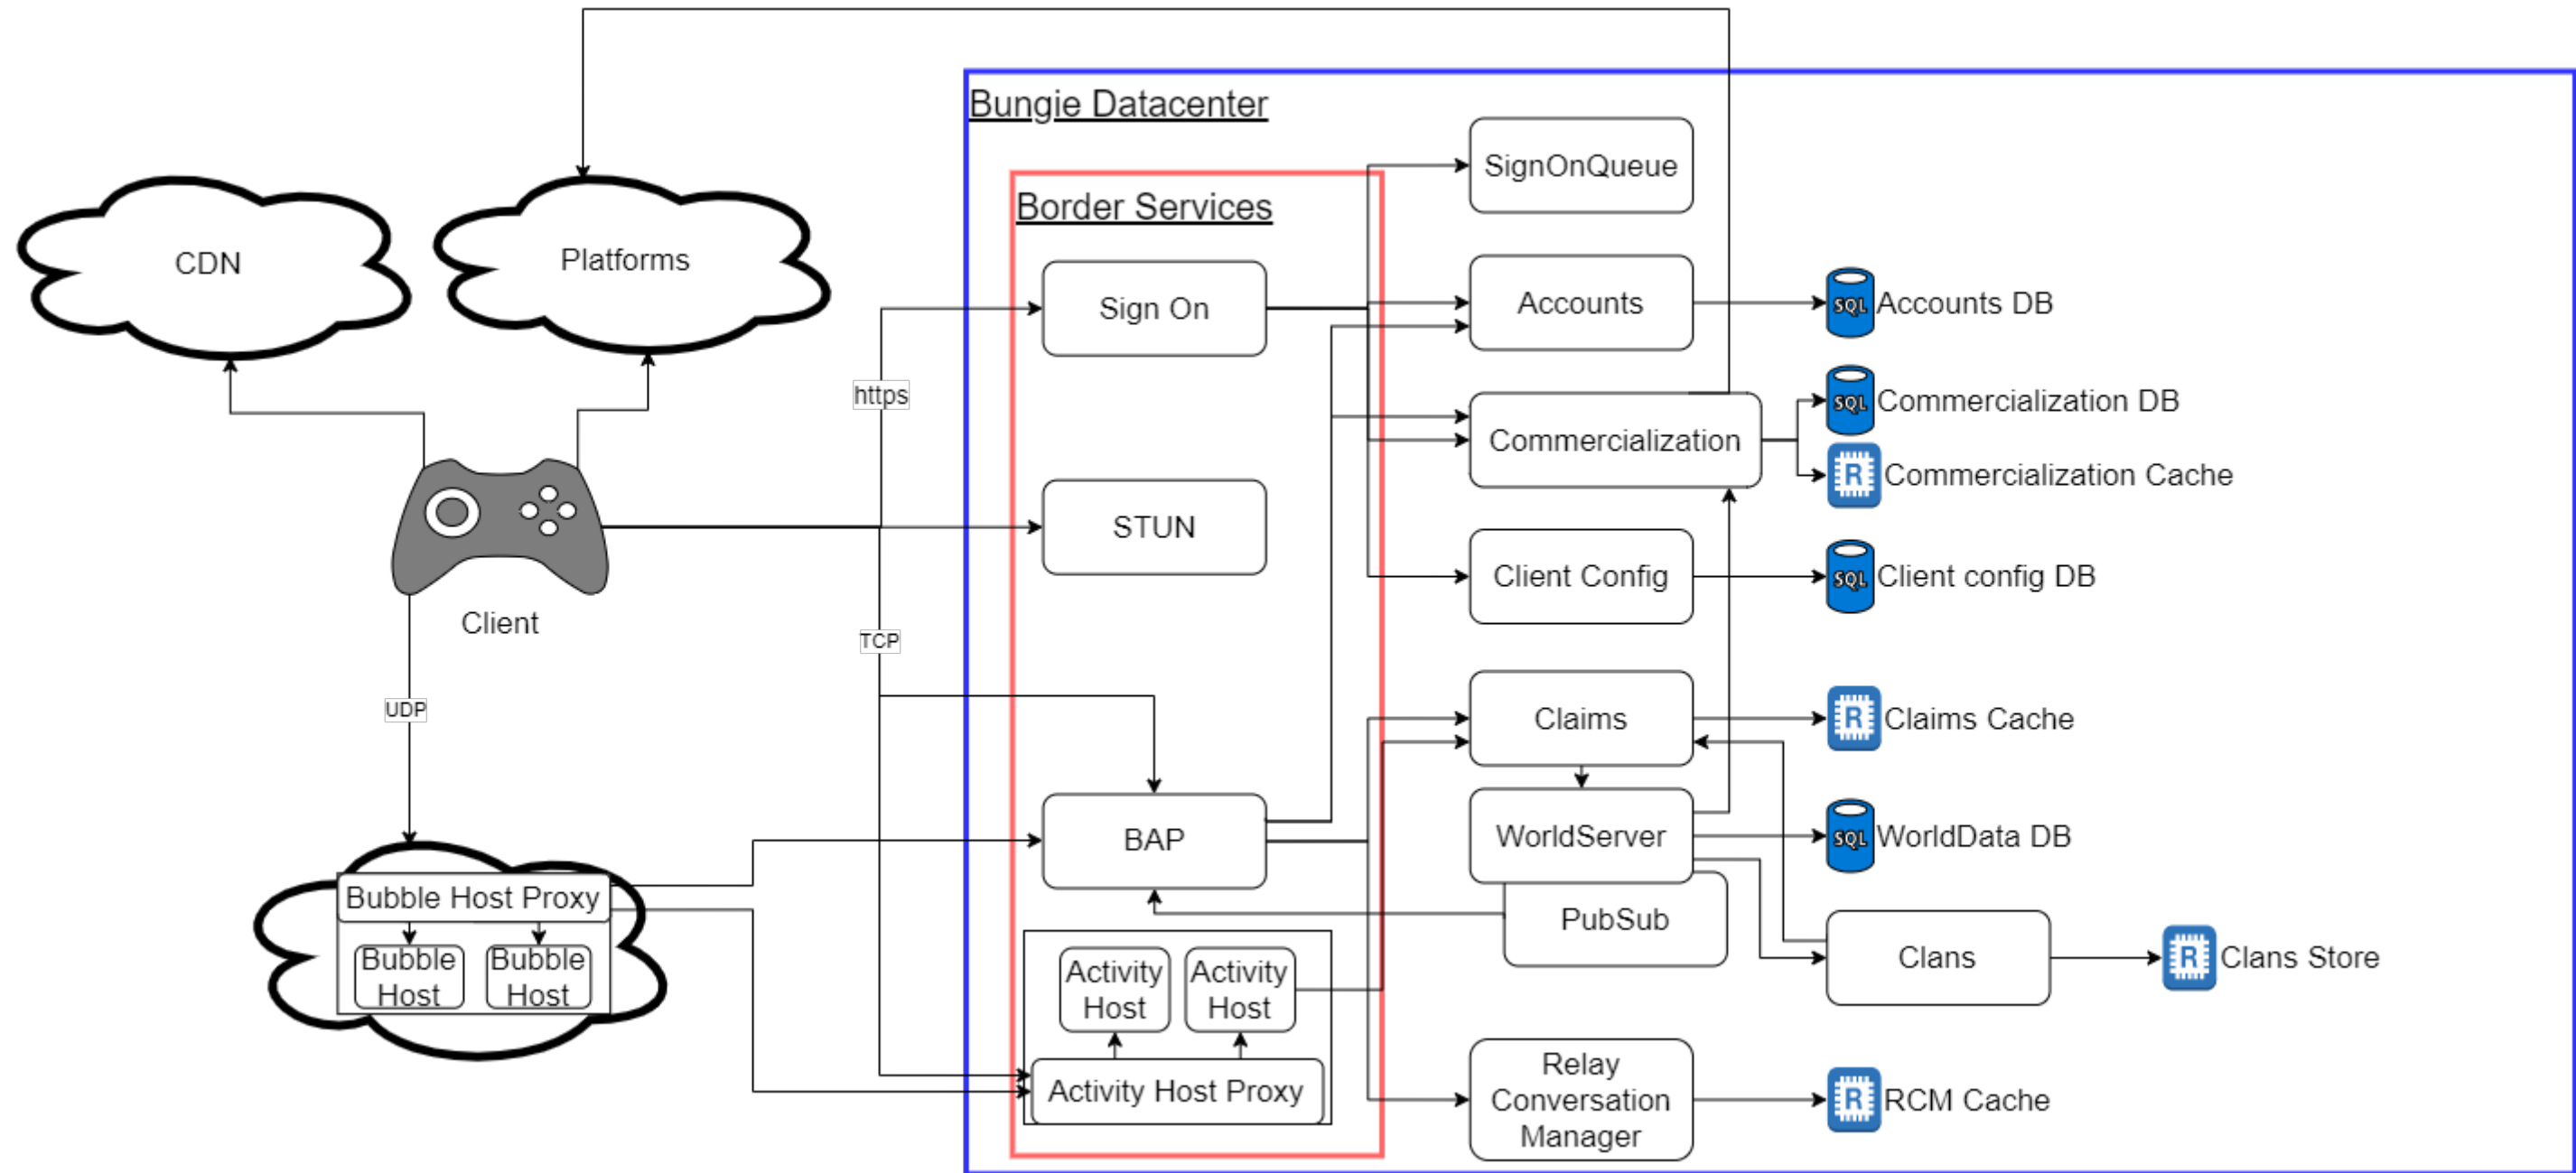
# Activity Hosts & Bubble Hosts

- **AH & BH are cut-down versions of the game client**
- **AHP & BHP are C# routing and management services**
- **1000s of Activity Hosts and Bubble Hosts per server**
- **AHP/BHP start up "Zombie" instances in advance**

# Activity Hosts & Bubble Hosts

- **Activity Host uses TCP.**
- **Bubble Host uses UDP**
- **Bonus Host Type: Group Activity Host**

# Destiny's Service Infrastructure



# Topic Grab Bag!

- **Stress Testing**
- **Cloud Usage**
- **User Error Reporting**
- **Service Settings**
- **Logging**





# Stress Testing

- **Major run each annual release**
- **Allocate a portion of our datacenter as a stress cluster**
- **Cloud hosted virtual clients**
- **Finds key ship-stopping issues every time**

# Stress Testing - Drawbacks

- **Expensive**
- **Labor intensive**
- **Heavy maintenance burden**
- **An Imperfect simulation**

# Stress Testing - Alternatives

- **Stress Test individual services as much as possible**
  - Doesn't catch multiple services contending for the same resource
- **Rely on the Queue to protect you**
- **Prefer soft-launching new features and systems**

# Cloud Usage

- **Destiny 1**
  - Peer-to-Peer networking allowed a single datacenter
  - NAT Relay had to be geolocated
- **Destiny 2**
  - Allowed Bubble Hosts to scale into the Cloud
  - Still not geolocated!
  - Used to handle population spikes for large launches

# User Error Reporting



**Protheon**  
@Protheon\_

the 4 horsemen of Bungie servers



11:56 AM · May 11, 2021 · Twitter for iPhone

266 Retweets 16 Quote Tweets 1,807 Likes

# Service Settings

- ## Definition

```
public readonly bool CrossSaveWarningEnabled = Setting(  
    true,           ↩ Default Setting  
    "Whether or not to give a one-time warning to users about cross save.", ↩ Text description of setting behavior  
    Behavior(Retrieving.PerRequest))    ↩ Describing how soon the setting will "take" after being changed (per call? after service restart?)  
    .Override(false, Service.Any, Environment.Onebox)  
    .Override(false, Service.Any, Environment.Shared); ↩ Optionally overrides the default setting per-environment
```

- ## Usage

```
if (SignOnServiceSettings.CurrentValue.CrossSaveWarningEnabled)  
{  
    // Send the Cross save warning!!  
}
```



# Service Settings

≡ Edit Setting:

Search:  .

## CrossSaveWarningEnabled

— Whether or not to give a one-time warning to users about cross save.

Not available on servers operated by public cloud providers.

Member of [Bungie.Server.SignOnService.SignOnServiceSettings](#)

from ...\Bungie.Server.SignOnService\Settings\SignOnServiceSettings.cs, line 923

» [Protect CrossSaveWarningEnabled from deployment changes.](#)

Service	Environment	Tag	Server	Value	
*	*	*	*	True	<input type="checkbox"/> <input type="checkbox"/>
*	Onebox	*	*	False	<input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/>
*	Shared	*	*	False	<input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/>

[Generate Export XML](#) » [Display Inline](#)

[Generate C# Code](#)

Services will pick up changes to this setting PerRequest

### Published Change History

[Expand All](#)

- Thu Sep 24 2020 11:09:45 GMT-0700 (Pacific Daylight Time) by DeployTool(UserName)

# Logging

```
_logger.Warning(sessionContext,  
    m => m("Invalid machine identifier: {0}",  
    Hexer.ToHex(extendedMachineIdentifier)));
```

- Session Context has caller information for filtering
- String Format only called if error level active
- Calculates a hash of the format string

# Logging

kibana

Discover

Visualize

Dashboard

Timelion

Canvas

Maps

Machine Learning

Infrastructure

Logs

APM

Uptime

Dev Tools

Monitoring

Management

Server Datamine ...

ServerName

Clear form

Server Datamine - Events Over Time

Count

Date per minute

Server Datamine - Events

1-50 of 98,032

Time	Priority	LogMessage	ServiceType	CategoryName	LoggerName	MachineName
May 28th 2021, 15:50:49.172	Status	Channel 29: From 127.0.0.1:56746: received Message Type=[StateUpdateRequest]	BubbleHostProxy	MessageHandlers	GameServerControlM essageHandler	

Table

JSON

View surrounding documents

View single document

BuildNumber

21

CategoryName

MessageHandlers

ContentVersion

15212.15212-0

Date

May 28th 2021, 15:50:49.172

Epid

998

FormatStringHash

65e4dde9

LineNumber

59

LogMessage

Channel 29: From 127.0.0.1:56746: received MessageType=[StateUpdateRequest]

LoggerName

GameServerControlMessageHandler

MachineName

MethodName

ProcessMessageAsync

Priority

Status

ProcessName

BubbleHostProxy:42500

ServerName

ServiceType

BubbleHostProxy

SessionId

4287529021173

ThreadId

195



## Key Takeaway #4

- **Invest in making it as easy as possible for your engineers to do the right thing**



# Not everything goes right





# Time for cowboy hats





# Game-Logic based data corruption

- **Incredibly dangerous, but also hard to catch**
- **Logic change that corrupts player or economy**
- **Example: Currency cap drastically reduced**
  - Responsible for Destiny's longest-ever unplanned downtime, as we executed a full game rollback

# Corruption Mitigations

- **Test with real player data**
- **Offline the game immediately!**
- **Have a rollback runbook**
- **Investigate options for fast-recovery**

# Data Store performance degradation

- **Data growth can cause big performance shifts**
- **New call patterns create new pressures**
- **One slow query can cause others to slow too**
- **Finding the real culprit can be nontrivial**

# Performance Degradation Mitigations

- **Sign On Queue and Throttle**
- **Disable optional systems**
- **Stress testing can catch many issues in advance**
- **Have instrumentation on the timing of every call**
- **Alert at dangerous thresholds**

# Retry-based death spiral

- **Systems will tend to retry when failures occur**
- **This can be automatic, or player driven**
- **Unexpected extra load from retries can cause systems to topple over... causing more retries**
- **Requests queuing can timeout before reaching the front of the queue... causing more retries**

# Death Spiral mitigations

- **Simulate failures during stress testing**
- **If Death-spiraling – Offline, and slow ramp population**
- **Where possible, avoid retries**
- **Prefer retry systems with backoff and Jitter**
- **Prefer short queues with rapid-rejection**



## Key Takeaway #5

- **Understand your failure space and have response runbooks ready**

# The Future of Services?





# The Future

- **Bungie Services not Destiny Services**
- **Support higher concurrencies**
- **Further leverage the cloud**
- **Increase iteration and deployment speeds**
- **Scale the team to meet the challenge**

# In summary

- 1. Sign on Queues and Throttles are one of your first and best tools to handle and prevent Service issues**
- 2. Whenever possible, launch new systems side-by-side with old systems, and slowly cut over**
- 3. Understand the real reliability requirements for your systems and be skeptical of 100% targets**
- 4. Invest in making it as easy as possible for your engineers to do the right thing**
- 5. Understand your failure space and have response runbooks ready**

# Thank You!!



<http://bungie.net/careers>

<http://www.linkedin.com/in/michael-williams-engineer>

Bonus Slides - Ignore





# Hardware failure

- **This can happen at any layer of your stack at any time**
  - Network gear
  - Hard drives
  - Key servers
  - ISPs
- **Can cause very unexpected states**
  - Network partitions
  - Significant latency increases
  - Intermittent request failures
- **Soft failures are often much worse than full hardware failures**
  - A slow hard drive is significantly worse than one that stops completely

# Hardware failure mitigations

- **Build out hardware redundantly where possible**
  - And regularly test the failover systems
- **Track metrics on hardware performance, not just simple health**
- **When possible, disable optional systems while working around a hardware failure.**

# Certificate issues

- **Your services will likely involve a huge number of certificates**
  - Cloud certs, signing certs, publisher certs, encryption certs, etc.
- **Any one of these expiring can immediately cause whole critical parts of your service to fail**
  - And certificates are rarely tied to optional elements of your services
- **Certificates expire all the time!**

# Certificate issue mitigations

- **Track your certificates**
- **Use multiple alert methods (including nag-mails) when certs are a few months out from expiration**
- **Update certificates early**
  - Because renewing certificates often requires partner conversations, a renewal cycle can become a big risk if you let it run down to the wire
- **Invest in certificate tracking solutions.**
- **Consider shorter renewals for certificates!**
  - A certificate that expires once every 5 years likely means the folks who installed it last time aren't in the same role today!