
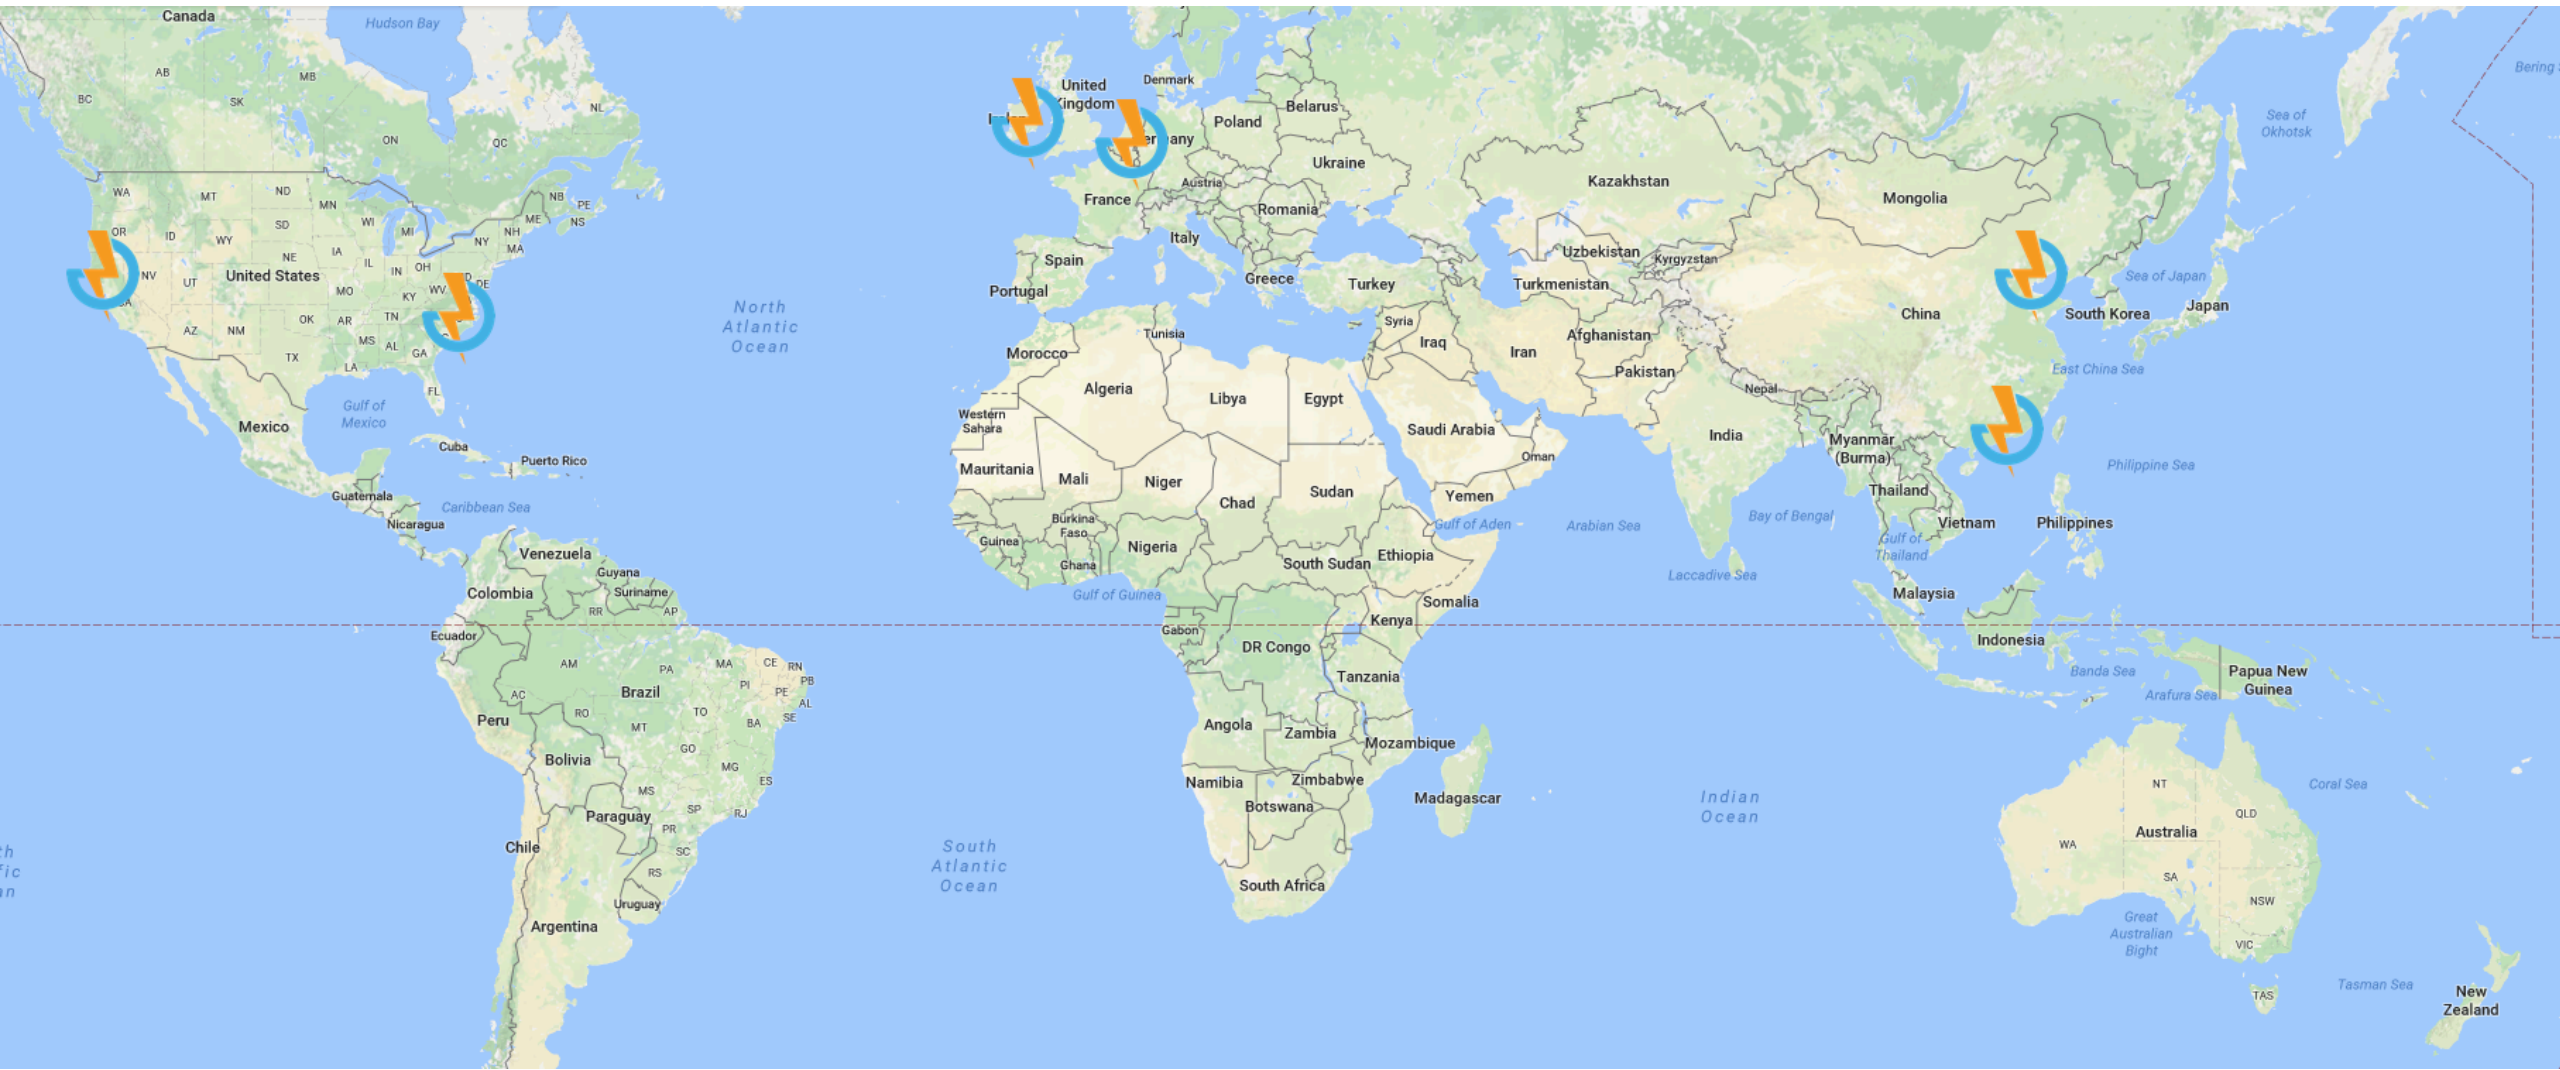


35 million




15 billion











Building Reliability In An Unreliable World



GAMESPARKS

GameSparks

Who?

Backend-as-a-Service provider for game developers

What?

All the server-side functionality a game needs

I see....



Failure – what is it?

“Failure is the state or condition of not meeting a desirable or intended objective, and may be viewed as the opposite of success”

<https://en.wikipedia.org/wiki/Failure>

Something that impacts customers

Something that impacts our service

Something that impacts our business

Failure – what causes it?

Provider issues

The Internet

Customers 😊

- Sudden change in load

- Bad code

- Bad data model

Attacks

Noisy neighbours

- “Strangers”

- “Family”

Human error

Failure – how to protect against it

Expect failure at every turn!

Stuff breaks – in ways you never imagine

People do dumb stuff

Minimise the Failure Domain

“section of a network that is negatively effected when a critical device or network service experiences problems”

“Smaller failure domains reduce the risk of disruption over a large section of a network, and eases the troubleshooting process.”

https://en.wikipedia.org/wiki/Failure_domain

GameSparks Failure Domains

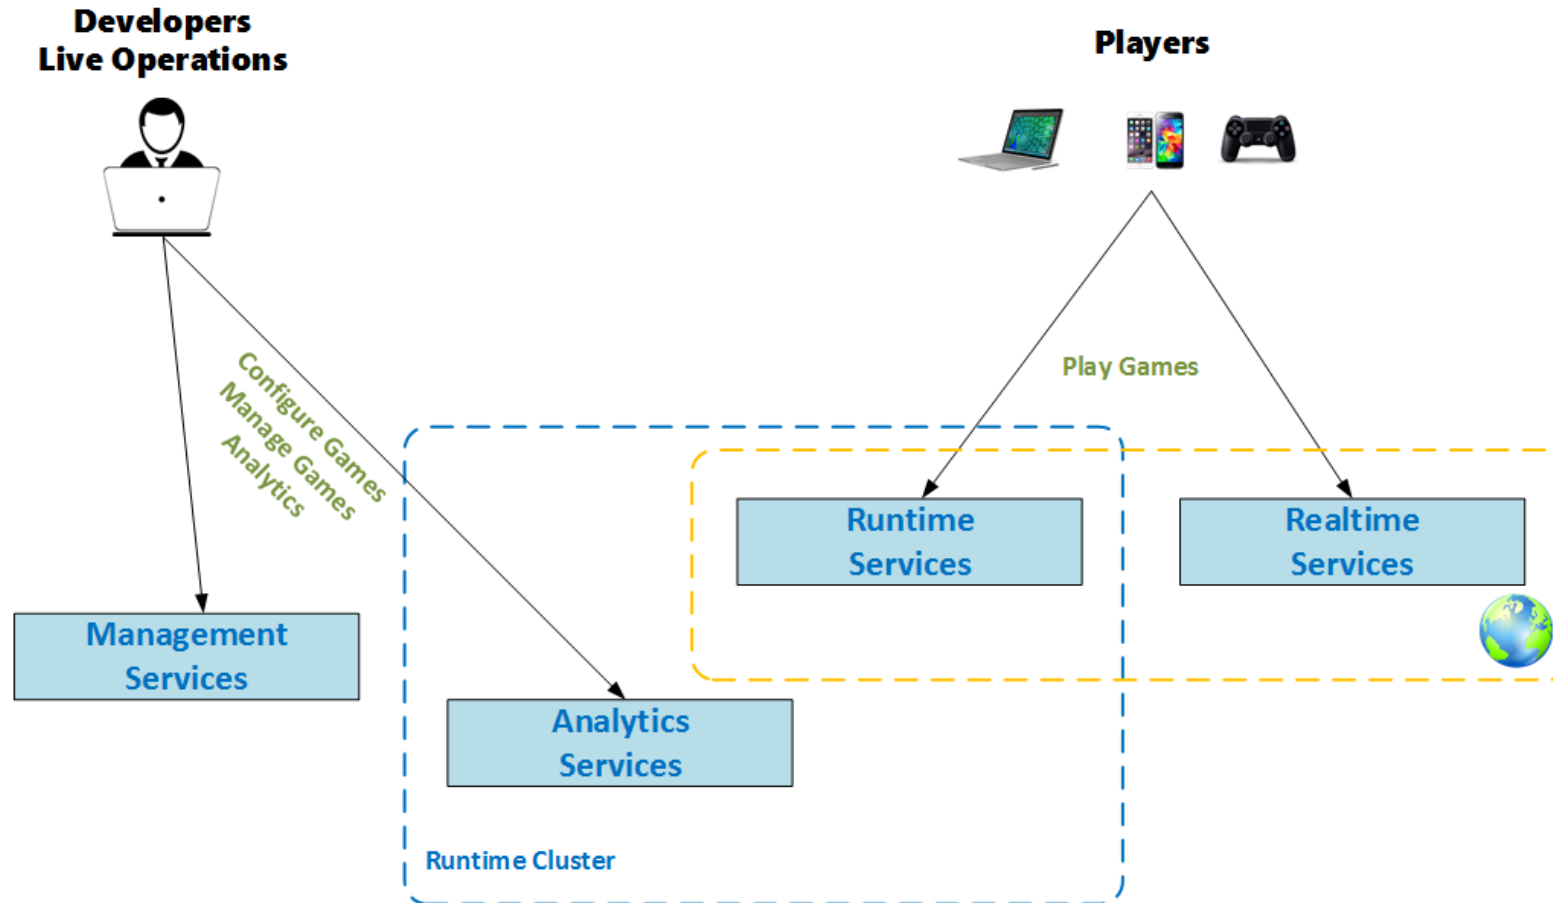
Platform Component

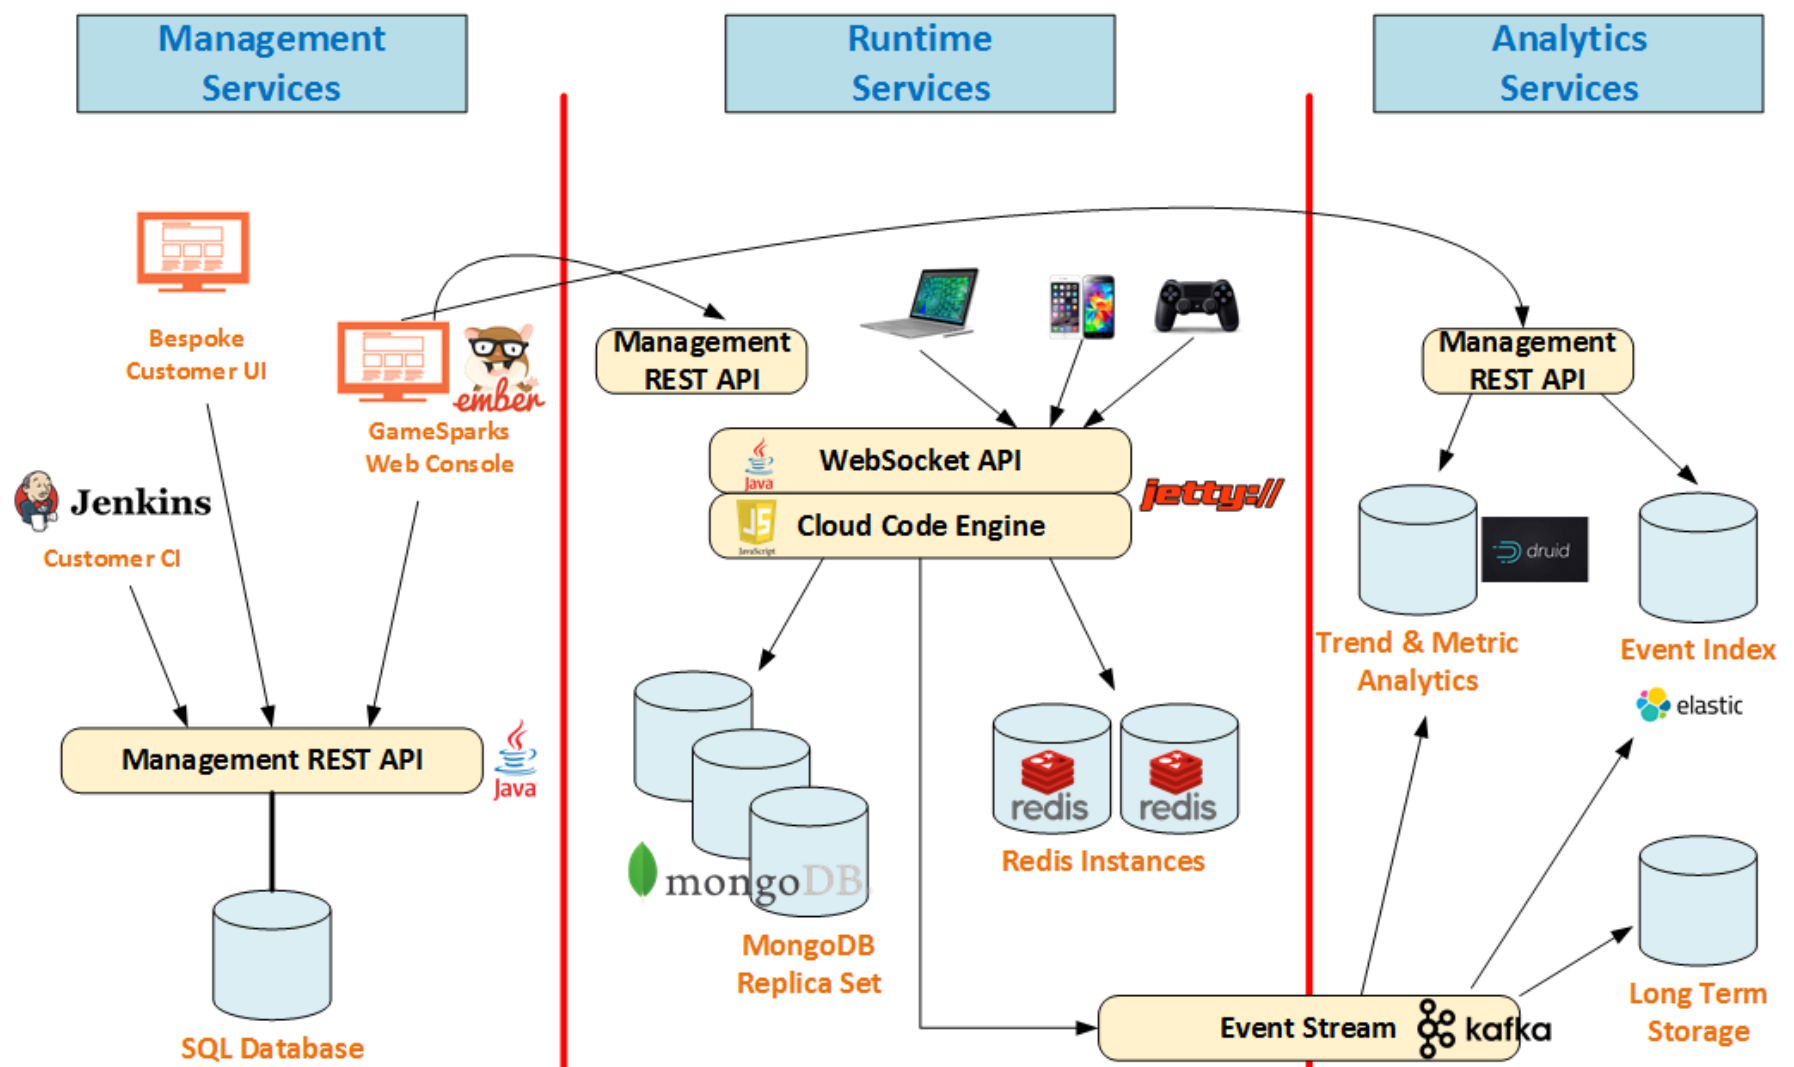
Component Deployment

Game

Technology Component

(Very) High-Level Architecture





Management Services

Runtime Services - API

| PartitionKey ^ | RowKey | Timestamp | apiKey | clusterName | filename | gameId | snapshotId | snapshotVersion | ts |
|----------------|--------|--------------------------|--------------|-------------|----------------------|--------|--------------------------|-----------------|------------|
| live | 298467 | 2016-07-19T15:11:49.497Z | c298467syBE8 | gsp-auw001 | 298467-snapshot-live | 298467 | 578e4290e4705a0583f7e2b7 | 1 | 1468941109 |
| live | 298500 | 2016-04-29T17:54:23.702Z | Y298500eGRAg | gsp-aeu001 | 298500-snapshot-live | 298500 | 57239fc62de9e962c6ef62dc | 1 | 1461952463 |
| live | 298504 | 2016-05-09T18:51:31.812Z | A298504y08ZF | gsp-aeu002 | 298504-snapshot-live | 298504 | 57237e372de9e962c6eef001 | 1 | 1462819891 |
| live | 298514 | 2016-04-30T12:48:49.552Z | x298514cl58m | gsp-auw001 | 298514-snapshot-live | 298514 | 5724a9a380e2e01df79a1527 | 1 | 1462020530 |
| live | 298524 | 2016-07-13T08:23:15.562Z | h298524PBBca | gsp-auw001 | 298524-snapshot-live | 298524 | 5785fa4ec5f88005c1498748 | 1 | 1468398195 |
| live | 298579 | 2016-07-17T13:51:56.100Z | l298579qkZf5 | gsp-aeu002 | 298579-snapshot-live | 298579 | 574434d6a5247706005750f4 | 1 | 1468763516 |
| live | 298590 | 2016-05-06T09:07:59.833Z | M298590z3b1T | gsp-acn001 | 298590-snapshot-live | 298590 | 572c5ec72de9e962c60a4eb1 | 1 | 1462525679 |
| live | 298593 | 2016-05-04T13:40:31.910Z | P298593s59c9 | gsp-aeu002 | 298593-snapshot-live | 298593 | 5729fbc22de9e962c6036ce6 | 1 | 1462369231 |
| live | 298594 | 2016-12-06T12:54:32.444Z | J298594x9MWs | gsp-aeu004 | 298594-snapshot-live | 298594 | 5846b3a2654d550a7749cc5f | 1 | 1481028872 |
| live | 298612 | 2016-05-05T14:26:26.663Z | p298612yry6K | gsp-auw001 | 298612-snapshot-live | 298612 | 572b57d52de9e962c60730bf | 1 | 1462458386 |
| live | 298621 | 2017-03-02T06:30:04.378Z | c298621Bgd3P | gsp-auw005 | 298621-snapshot-live | 298621 | 58b7bbe26e2c740572a7fb58 | 1 | 1488436203 |
| live | 298631 | 2016-05-30T09:22:03.211Z | j298631L4yBt | gsp-aeu001 | 298631-snapshot-live | 298631 | 574c062c5b917005f46b9149 | 1 | 1464600123 |
| live | 298643 | 2016-06-04T18:46:46.856Z | d298643A6s83 | gsp-auw001 | 298643-snapshot-live | 298643 | 57532202d0c95d0607e2f4f2 | 1 | 1465066007 |
| live | 298655 | 2016-10-15T15:41:03.302Z | S298655oEG0D | gsp-aeu004 | 298655-snapshot-live | 298655 | 58024dd85e5528a4ff1e4735 | 1 | 1476546063 |

Load Balancer

Runtime Services - Datastores

Websockets

The Good

- Reduced handshake overhead

- Minimal headers

- Asynchronous messaging

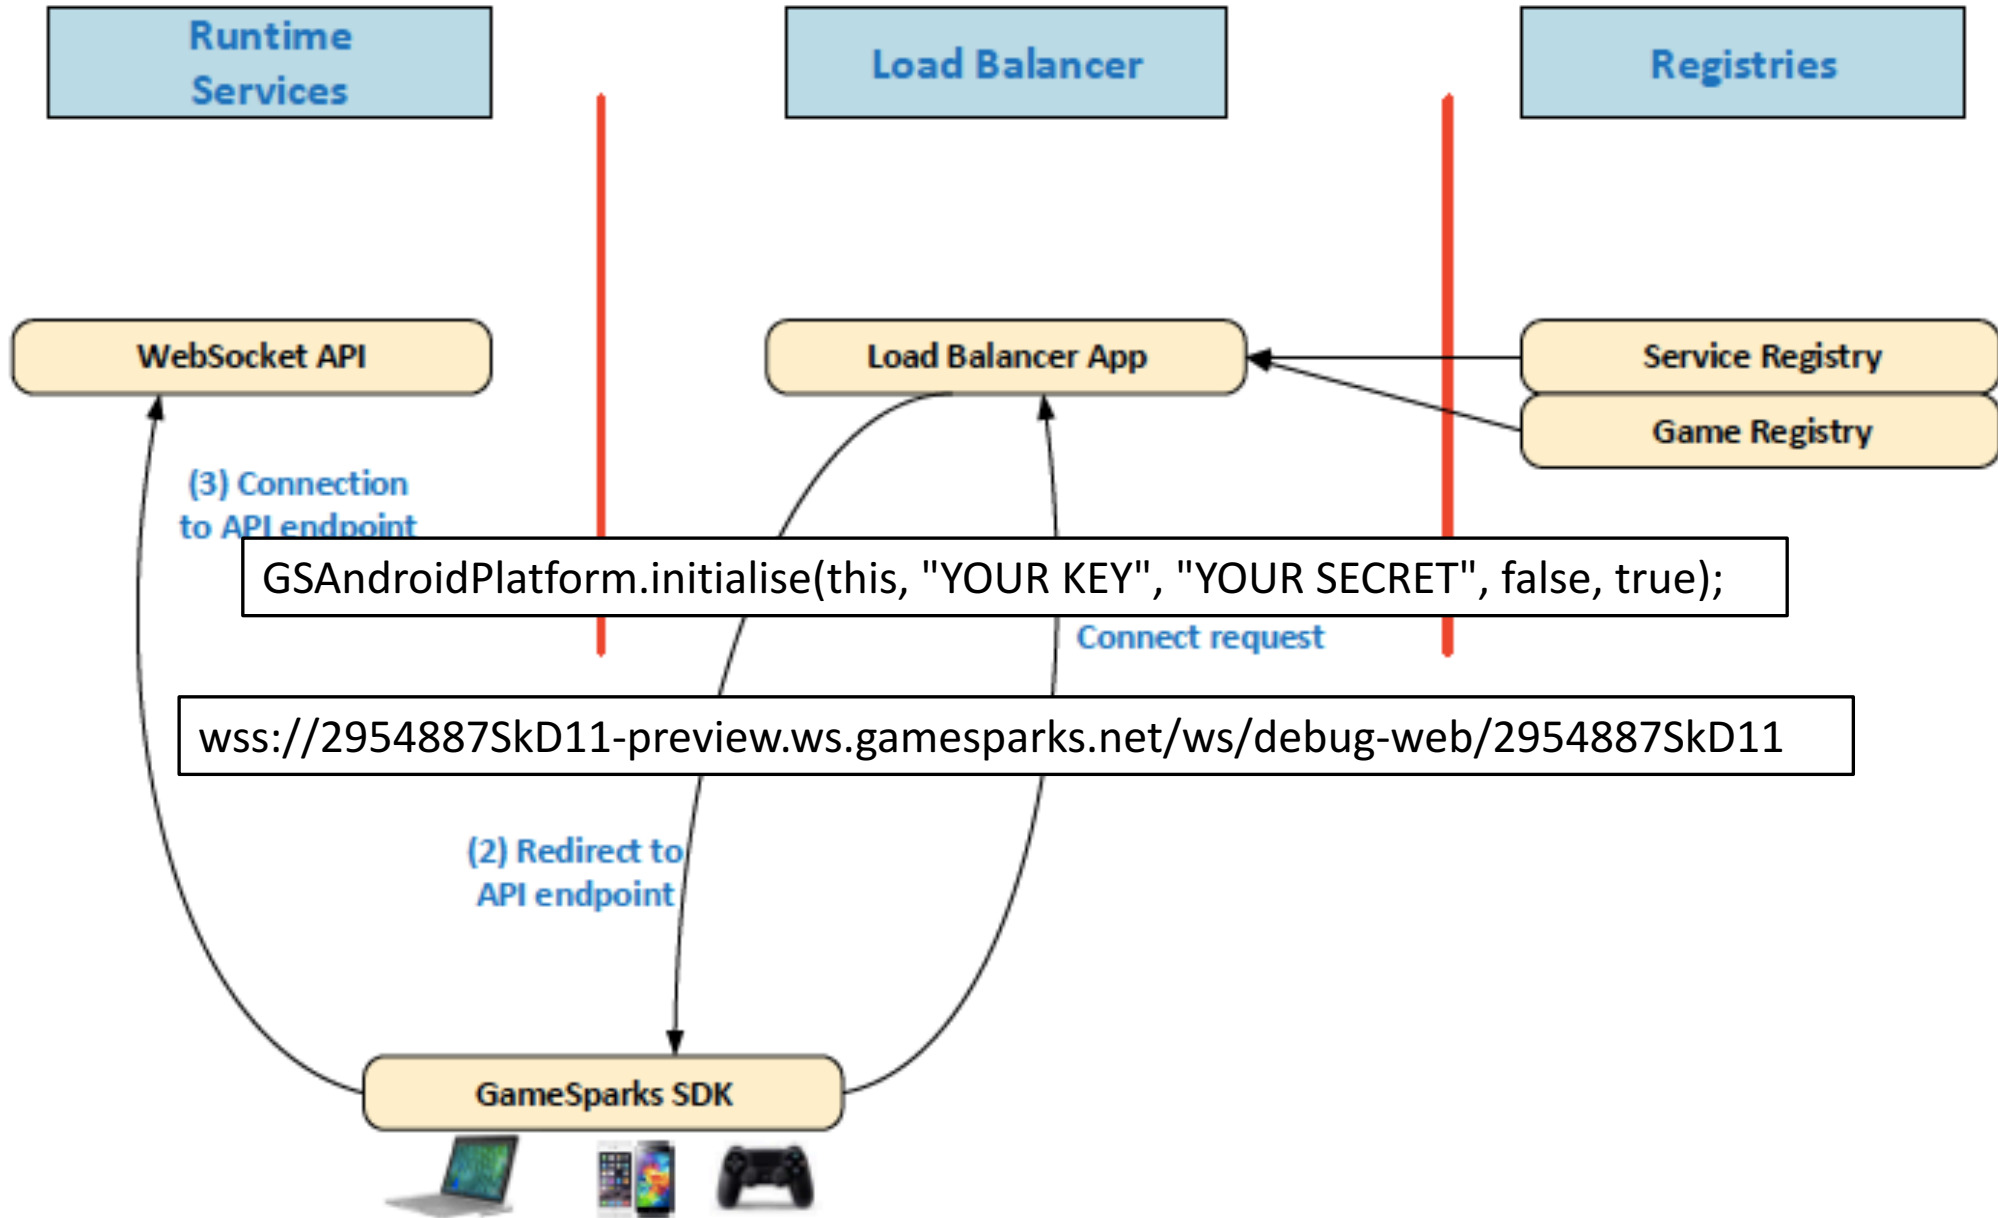
- No polling

The Bad

- Load balancing!

The Ugly

- The Internet!



Workload segregation

| PartitionKey | RowKey ^ | Timestamp | capabilities | domainName | excludedGames | includedGames | localhostname | sessionCount | status | ts | version |
|-------------------|------------------------|---------------------------------|----------------------|-----------------------|---------------|---------------|------------------------|--------------|------------------|-------------------|----------------------|
| gsp-aeu005 | gsp-aeu005-se01 | 2017-02-24T15:54:57.303Z | server,rest | gamesparks.net | | | gsp-aeu005-se01 | 5697 | available | 1487951696 | 3b2feeea85a9+ |
| gsp-aeu005 | gsp-aeu005-se02 | 2017-02-24T15:54:58.404Z | server,rest | gamesparks.net | | | gsp-aeu005-se02 | 5732 | available | 1487951697 | 3b2feeea85a9+ |
| gsp-aeu005 | gsp-aeu005-se03 | 2017-02-24T15:54:59.182Z | server,rest | gamesparks.net | | | gsp-aeu005-se03 | 5722 | available | 1487951698 | 3b2feeea85a9+ |
| gsp-aeu005 | gsp-aeu005-se04 | 2017-02-24T15:54:58.365Z | server,rest | gamesparks.net | | | gsp-aeu005-se04 | 5731 | available | 1487951697 | 3b2feeea85a9+ |
| gsp-aeu005 | gsp-aeu005-se05 | 2017-02-24T15:54:56.730Z | server,rest | gamesparks.net | | | gsp-aeu005-se05 | 5739 | available | 1487951696 | 3b2feeea85a9+ |
| gsp-aeu005 | gsp-aeu005-se06 | 2017-02-24T15:54:56.853Z | server,rest | gamesparks.net | | | gsp-aeu005-se06 | 5699 | available | 1487951696 | 3b2feeea85a9+ |
| gsp-aeu005 | gsp-aeu005-se07 | 2017-02-24T15:55:00.671Z | server,rest | gamesparks.net | | | gsp-aeu005-se07 | 5761 | available | 1487951700 | 3b2feeea85a9+ |
| gsp-aeu005 | gsp-aeu005-se08 | 2017-02-24T15:54:57.163Z | server,rest | gamesparks.net | | | gsp-aeu005-se08 | 5711 | available | 1487951696 | 3b2feeea85a9+ |
| gsp-aeu005 | gsp-aeu005-se09 | 2017-02-24T15:54:57.039Z | server,rest | gamesparks.net | | | gsp-aeu005-se09 | 5684 | available | 1487951696 | 3b2feeea85a9+ |
| gsp-aeu005 | gsp-aeu005-se10 | 2017-02-24T15:54:59.290Z | server,rest | gamesparks.net | | | gsp-aeu005-se10 | 5711 | available | 1487951698 | 3b2feeea85a9+ |
| gsp-aeu005 | gsp-aeu005-se38 | 2017-02-24T15:54:57.977Z | server,rest | gamesparks.net | | | gsp-aeu005-se38 | 5664 | available | 1487951697 | 3b2feeea85a9+ |
| gsp-aeu005 | gsp-aeu005-se39 | 2017-02-24T15:54:56.911Z | server,rest | gamesparks.net | | | gsp-aeu005-se39 | 5696 | available | 1487951696 | 3b2feeea85a9+ |
| gsp-aeu005 | gsp-aeu005-se41 | 2017-02-24T15:54:56.868Z | debug,scheduler,bulk | gamesparks.net | | | gsp-aeu005-se41 | 3 | available | 1487951696 | 3b2feeea85a9+ |
| gsp-aeu005 | gsp-aeu005-se42 | 2017-02-24T15:54:58.884Z | debug,scheduler,bulk | gamesparks.net | | | gsp-aeu005-se42 | 1 | available | 1487951698 | 3b2feeea85a9+ |

Auto Scaling and Healing

We wrote our own auto-scaler – eek!

Metric driven

- CPU

- Heap usage

- Garbage Collection

- Current Connections

- Arrival Rate

- Throughput

Prediction via scikit-learn Python module

Durable requests

Some requests don't matter, but some really do

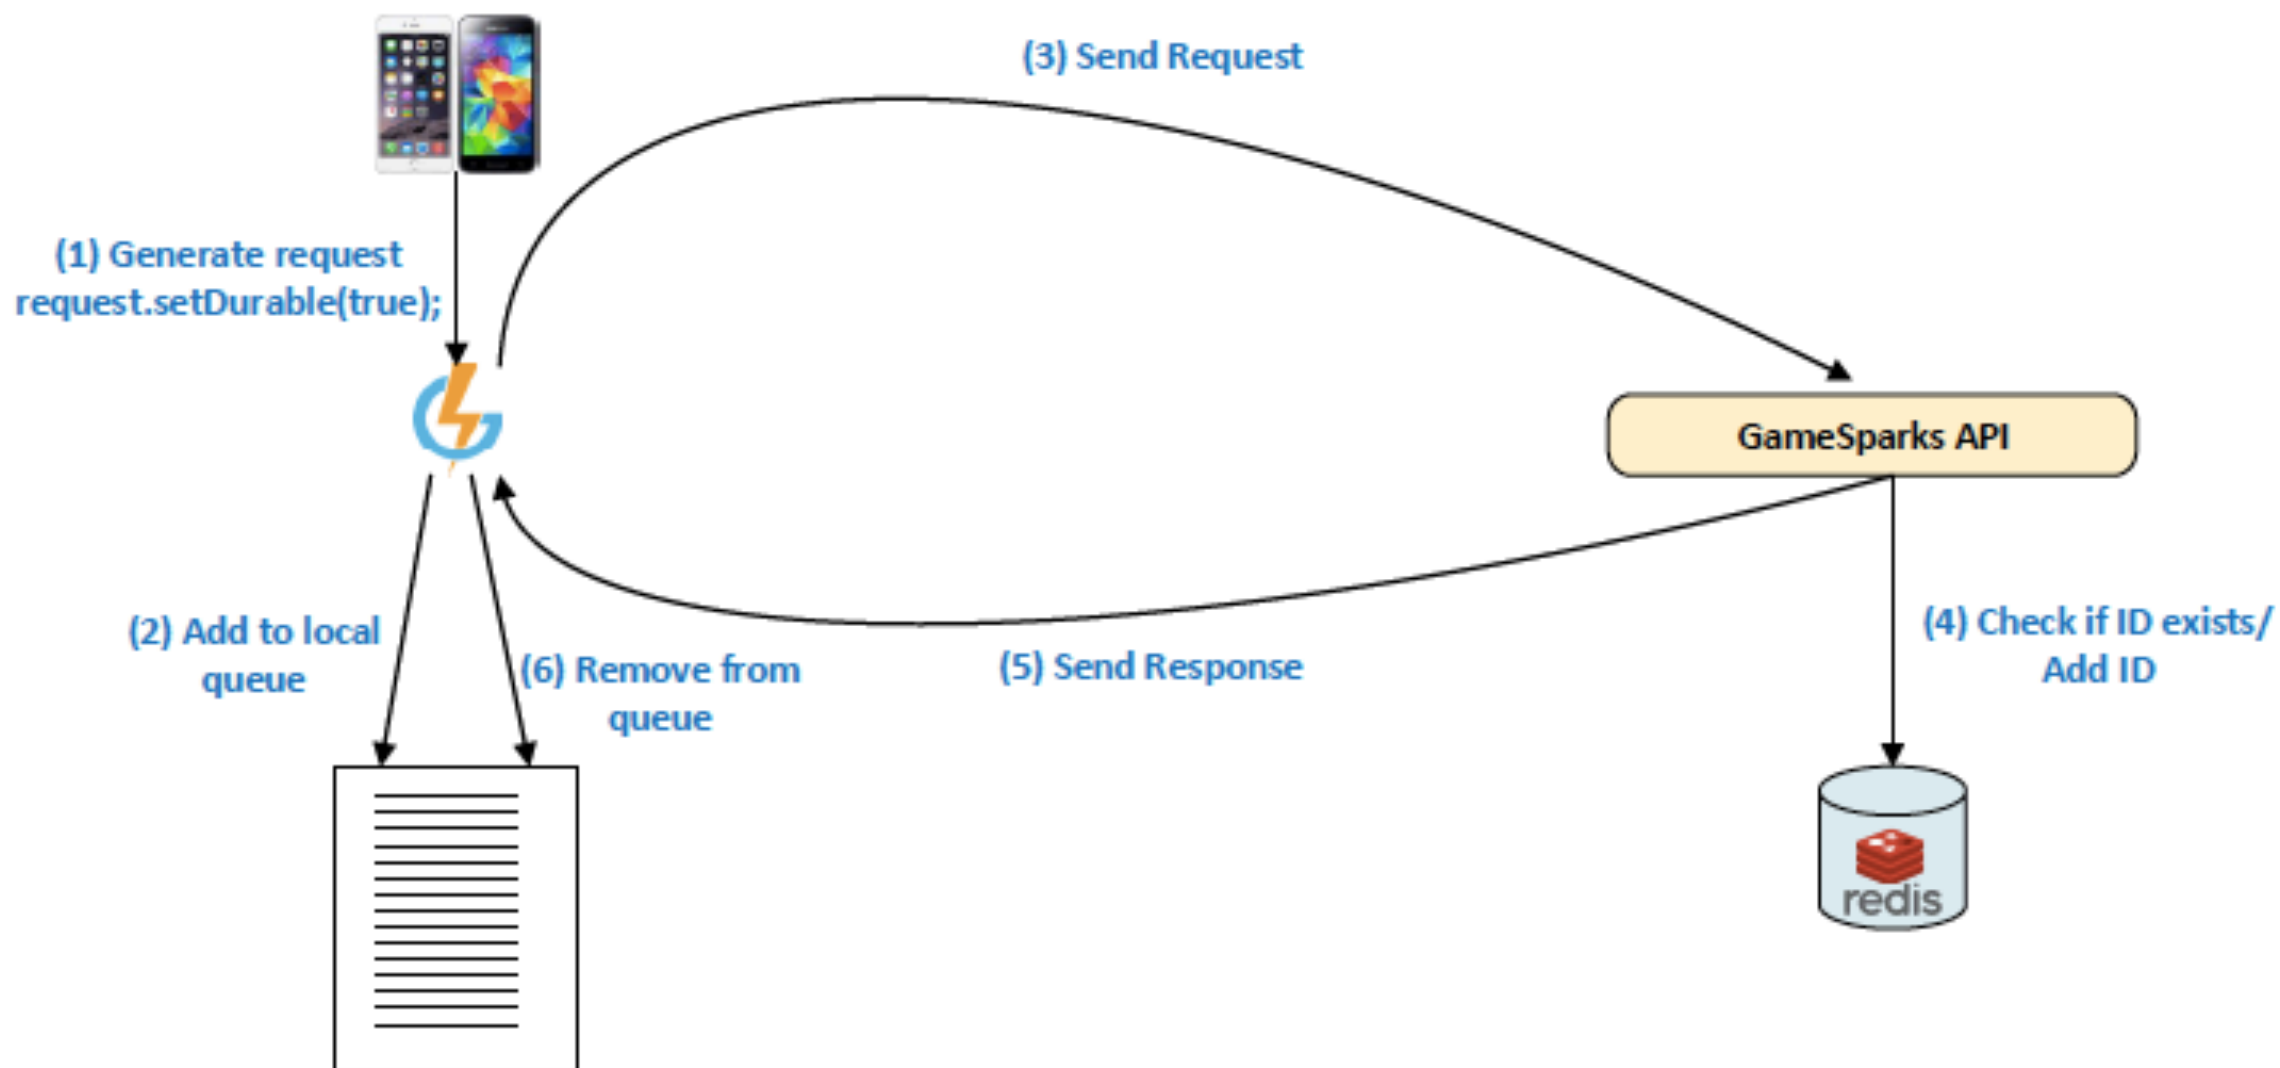
Request failure – why does it happen?

- Error processing the request

- Network failure between client and server

- Network failure between server and client

```
request.setDurable(true);
```



Resource Management – code

```
for (;;) {}
```

Instrumentation

- Execution time

- Statement count

- Bytecode instructions

```
var ms = getRemainingMilliseconds()
```

com.sun.management.ThreadMXBean

Method Summary

Methods

| Modifier and Type | Method and Description |
|-------------------|--|
| long | getThreadAllocatedBytes (long id) Returns an approximation of the total amount of memory, in bytes, allocated in heap memory for the thread of the specified ID. |
| long[] | getThreadAllocatedBytes (long[] ids) Returns an approximation of the total amount of memory, in bytes, allocated in heap memory for each thread whose ID is in the input array <i>ids</i> . |
| long[] | getThreadCpuTime (long[] ids) Returns the total CPU time for each thread whose ID is in the input array <i>ids</i> in nanoseconds. |
| long[] | getThreadUserTime (long[] ids) Returns the CPU time that each thread whose ID is in the input array <i>ids</i> has executed in user mode in nanoseconds. |
| boolean | isThreadAllocatedMemoryEnabled () Tests if thread memory allocation measurement is enabled. |
| boolean | isThreadAllocatedMemorySupported () Tests if the Java virtual machine implementation supports thread memory allocation measurement. |
| void | setThreadAllocatedMemoryEnabled (boolean enable) Enables or disables thread memory allocation measurement. |

Resource Management – data

Data persistence + flexibility = danger!

Issues we see with data persisted in MongoDB:

- Unindexed data

- Low cardinality data

- Poor data models


- Inefficient access

- Full updates

- Query Repetition

MongoDB Auto-indexing

```
try { Spark.runtimeCollection("map").dropIndex({"userId": 1, "Building.Id": 1}); } catch (e) { }
try { Spark.runtimeCollection("map").dropIndex({"X": 1, "Y": 1}); } catch (e) { }
try { Spark.runtimeCollection("map").dropIndex({"userId": 1, "Building.UniqId": 1}); } catch (e) { }
try { Spark.runtimeCollection("map").dropIndex({"userId": 1}); } catch (e) { }
try { Spark.runtimeCollection("map").dropIndex({"Path": 1}); } catch (e) { }
try { Spark.runtimeCollection("map").dropIndex({"X": 1, "Y": 1, "Path": 1}); } catch (e) { }
try { Spark.runtimeCollection("map").dropIndex({"X": 1, "Y": 1, "Path": 1, "Rubble": 1}); } catch (e) { }
try { Spark.runtimeCollection("map").dropIndex({"Rubble": 1}); } catch (e) { }
try { Spark.runtimeCollection("map").dropIndex({"Pit": 1}); } catch (e) { }
try { Spark.runtimeCollection("map").dropIndex({"userId": 1, "X": 1, "Y": 1}); } catch (e) { }
Spark.runtimeCollection("map").ensureIndex({"userId": 1, "X" : 1, "Y" : 1, "Building.Id": 1, "Building.EndConstructionTime" : 1});
Spark.runtimeCollection("map").ensureIndex({"userId": 1, "X" : 1, "Y" : 1, "Building.EndConstructionTime" : 1});
Spark.runtimeCollection("map").ensureIndex({"userId": 1, "X" : 1, "Y" : 1, "Building.Expedition.EndExpeditionTime": 1});
Spark.runtimeCollection("map").ensureIndex({"userId": 1, "Building.Id": 1, "Building.Level": 1});
Spark.runtimeCollection("map").ensureIndex({"userId": 1, "Building.UniqId": 1});
Spark.runtimeCollection("map").ensureIndex({"userId": 1, "Pit.StartCollectingTime" : 1, "Pit.EndCollectingTime" : 1});
Spark.runtimeCollection("map").ensureIndex({"userId": 1, "X" : 1, "Y" : 1, "Path": 1, "Building": 1, "Rubble": 1, "Pit": 1});
```




```
{
  "_id" : ObjectId("58a6cf1effdbd06e93fb71bd"),
  "collection" : "script.jsTestRuntime",
  "query" : {
    "fieldA" : "?",
    "fieldB" : "?",
    "numericValue" : "?"
  },
  "lastOccurrence" : ISODate("2017-02-22T17:09:21.041Z"),
  "lastExample" : {
    "query" : {
      "fieldA" : "fieldA_1",
      "fieldB" : "fieldB_1",
      "numericValue" : 1
    }
  },
  "occurrences" : {
    "2017-02-17" : {
      "update" : {
        "count" : 28,
        "time" : NumberLong(147)
      },
      "findOne" : {
        "count" : 7,
        "time" : NumberLong(34)
      },
      "count" : {
        "count" : 7,
        "time" : NumberLong(7)
      }
    }
  }
}
```

The collection being queried

The query itself (plus projections and sorts)

Example variables

Types of query and counts



```
{"fieldA": "fieldA_1", "fieldB": "fieldB_1", "numericValue": 1}
```

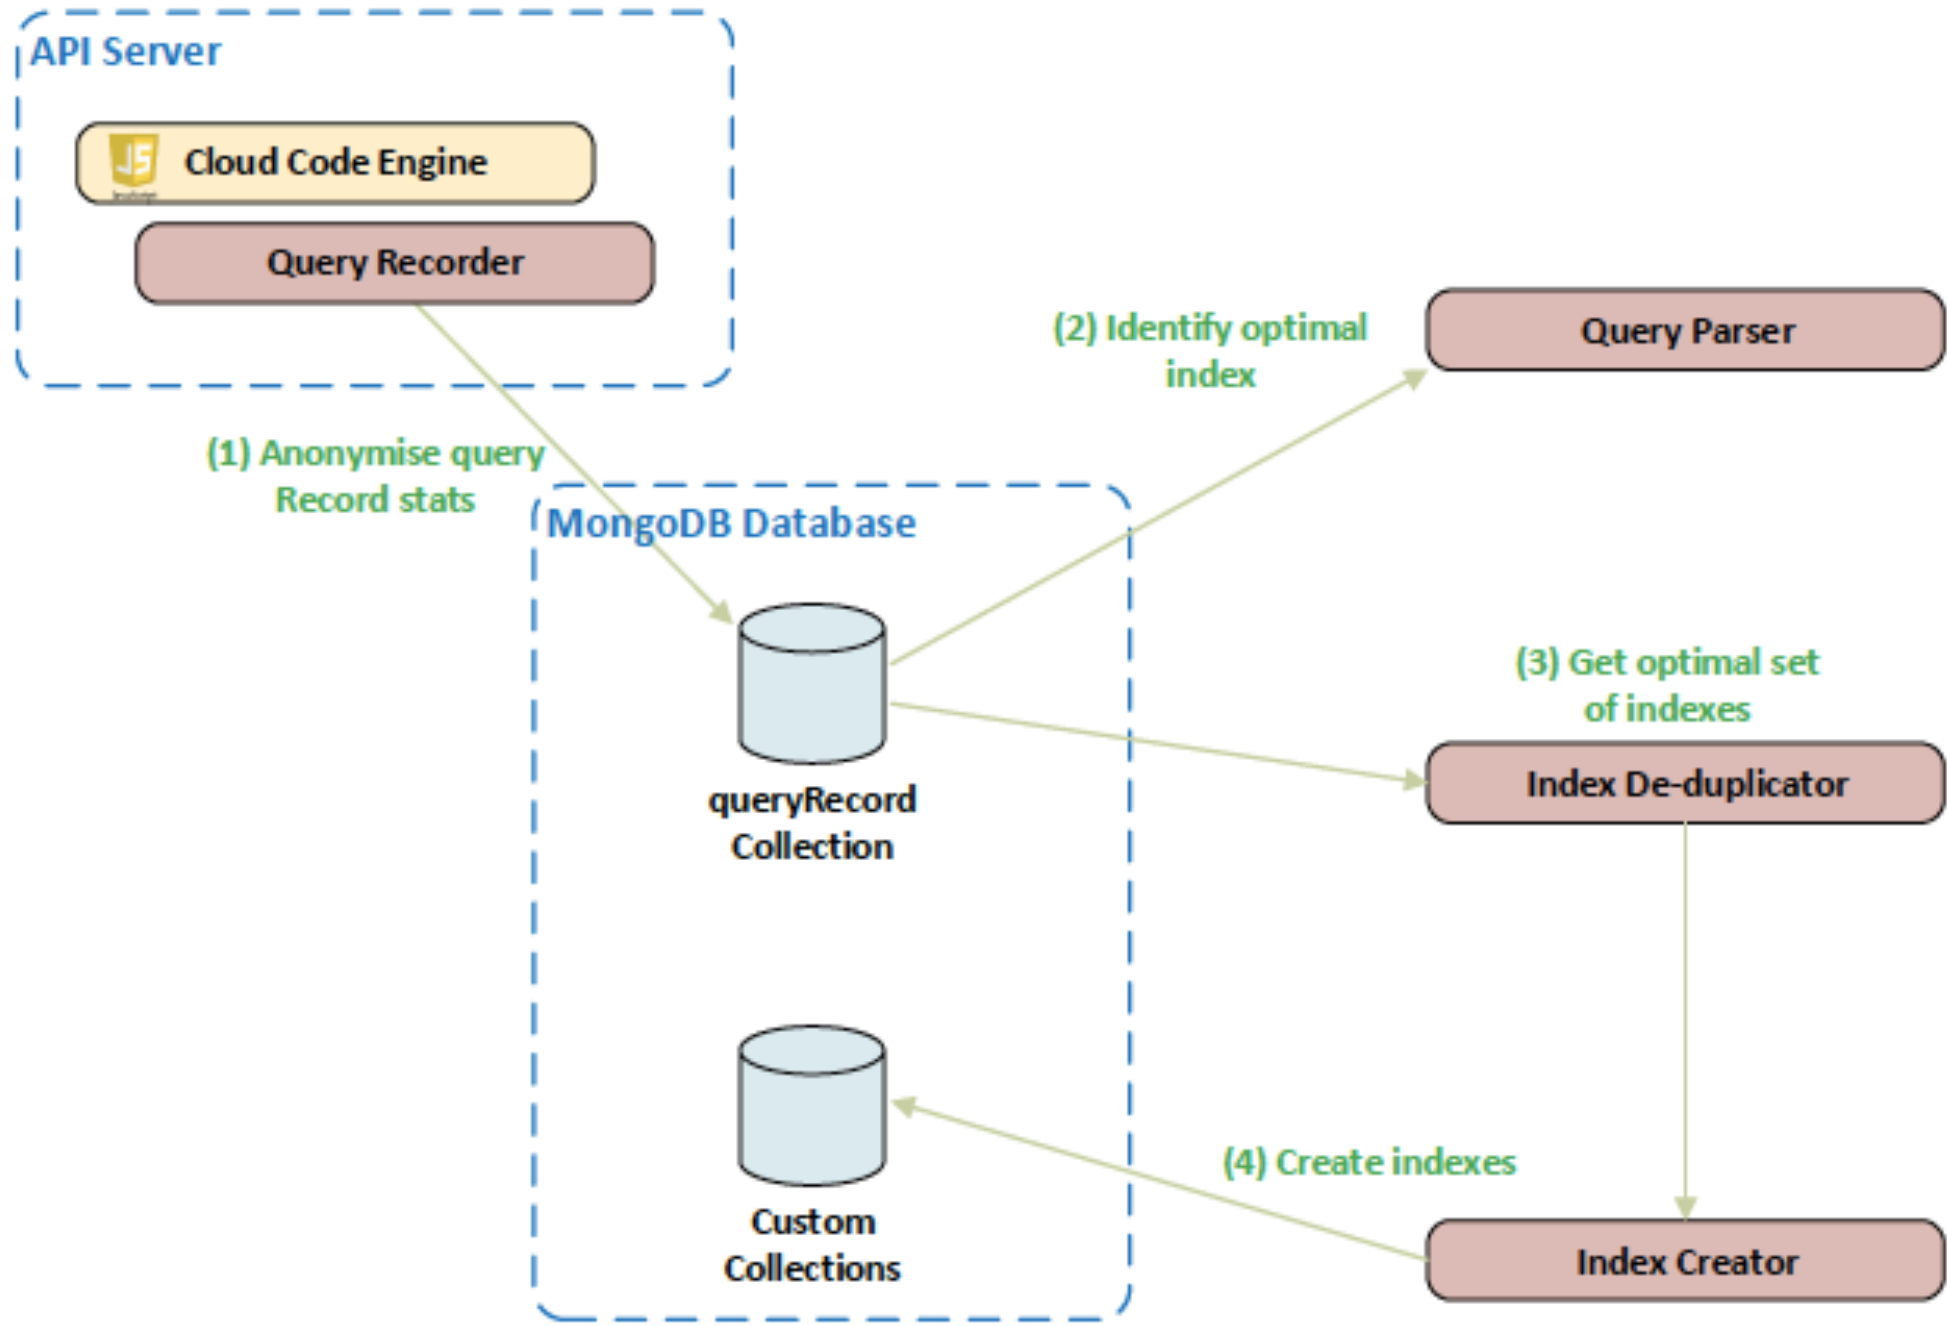
```
Index: {"fieldA": 1, "fieldB": 1, "numericValue": 1}
```

```
{"fieldA": "fieldA_1", "fieldB": "fieldB_1"}
```

```
Index: {"fieldA": 1, "fieldB": 1}
```

```
{"fieldA": "fieldA_1"}
```

```
Index: {"fieldA": 1}
```



Partial updates

```
var myRuntimeCollection = Spark.runtimeCollection('runtimetest');
```

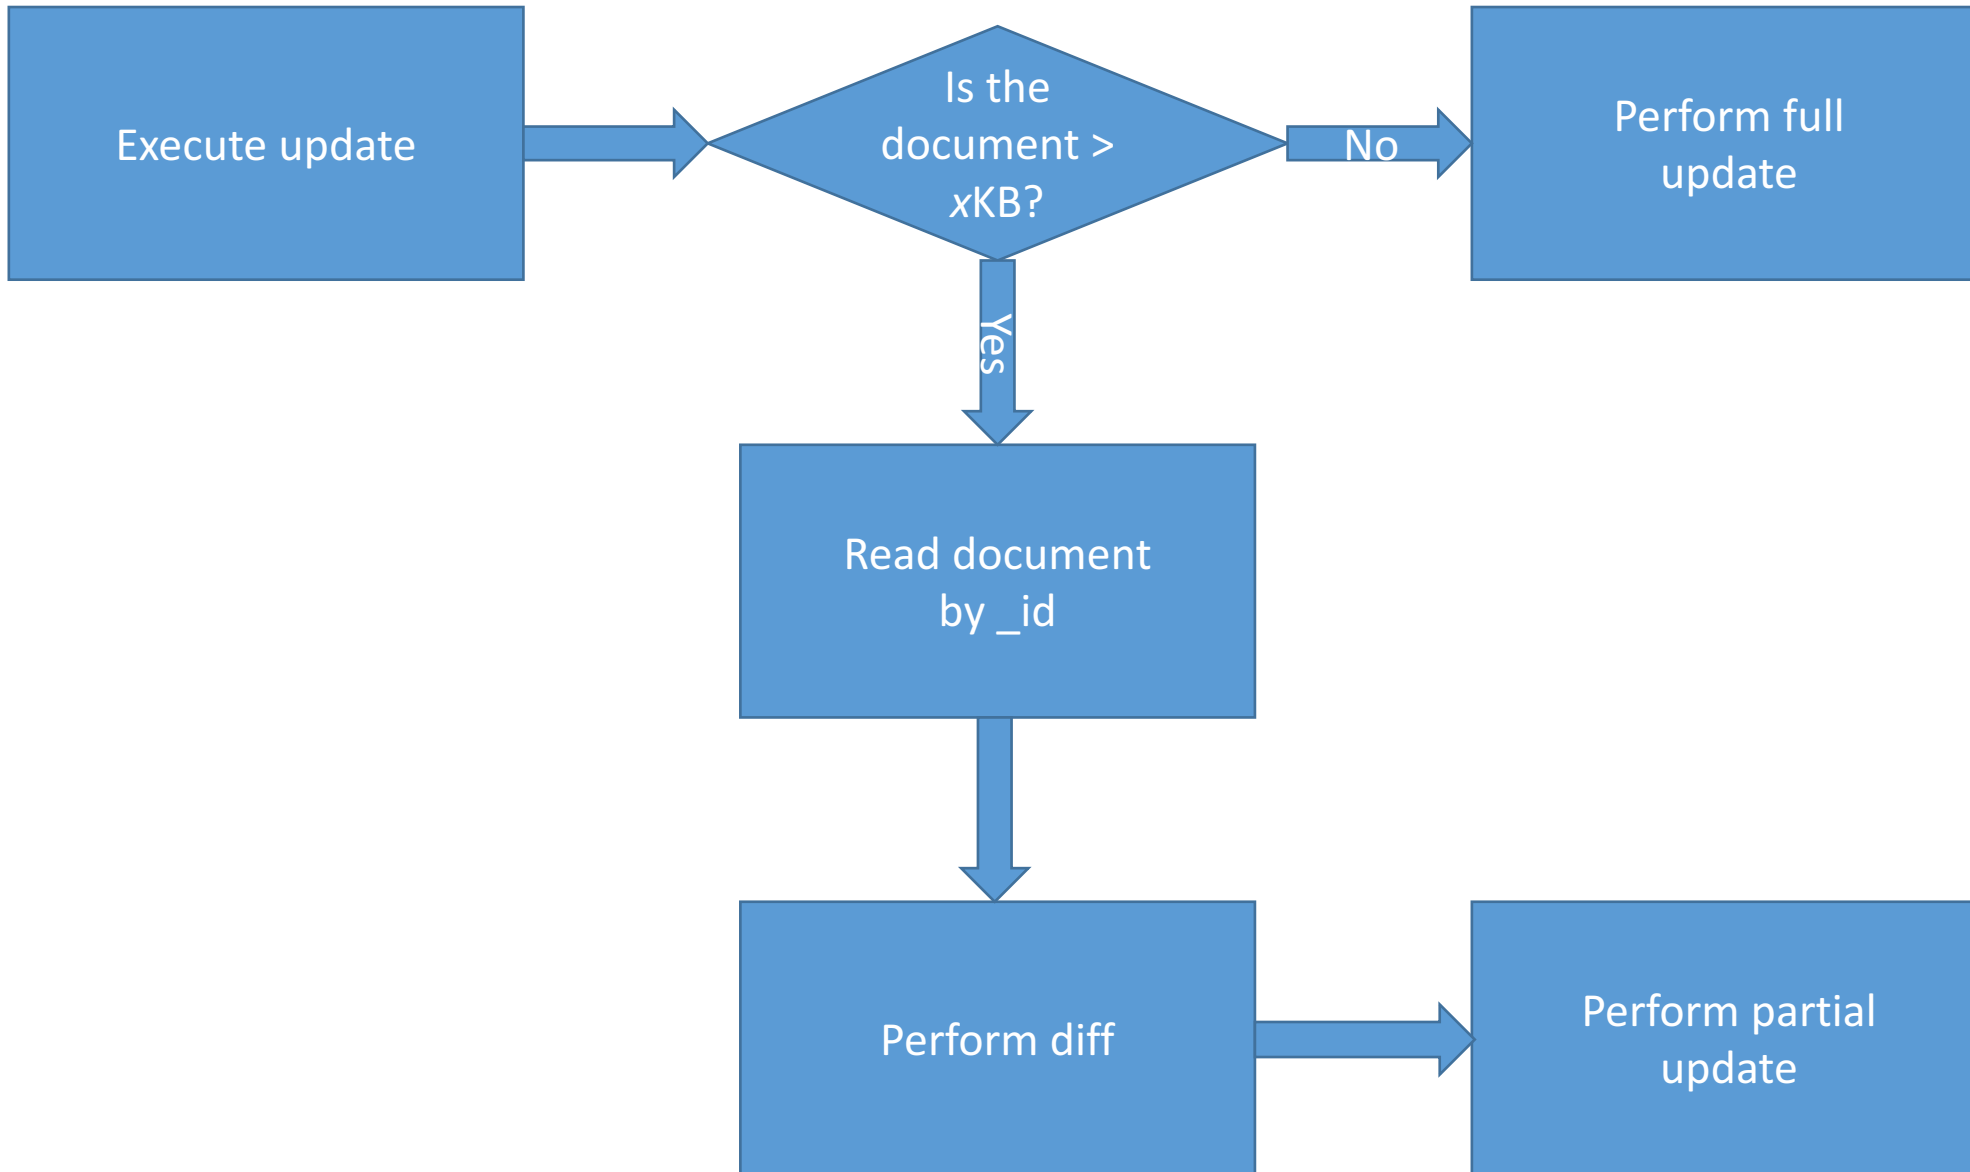
```
var results = myRuntimeCollection.findOne({"_id": "abc123"});
```

```
<<do something>>
```

```
var success = myRuntimeCollection.update({"_id" : "abc123"}, results);
```

```
<<do something>>
```

```
var success = myRuntimeCollection.update({"_id" : "abc123"}, results);
```

Resource tracking

Track the resource usage of every request

Identify hotspots and high consumers

Highlight anomalies

Track performance trends

```

"metrics": {
  "redisTimePlatformTotal": 0,
  "redisCountPlatformTotal": 0,
  "redisTimeScriptTotal": 2,
  "redisCountScriptTotal": 8,
  "mongoTimePlatform": {},
  "mongoCountPlatform": {},
  "mongoTimePlatformTotal": 0,
  "mongoCountPlatformTotal": 0,
  "mongoTimeScript": {
    "find": {
      "script.Matches": 0,
      "script.FieldPlayers": 0,
      "script.ScheduleActions": 0
    },
    "findOne": {
      "script.MatchSnapShot": 1,
      "script.ObjectCache": 0,
      "script.Sponsoring": 0,
      "script.Clubs": 4,
      "script.AchievementTracker": 0,
      "script.Leagues": 0,
      "player": 0
    },
    "save": {
      "script.ObjectCache": 0,
      "script.AchievementTracker": 1,
      "script.ScheduleActions": 2
    },
    "count": {
      "script.Matches": 1
    },
    "update": {
      "script.ClubLeagueStatistics": 2,
      "script.Leagues": 0,
      "script.SquadDynamic": 2
    },
    "remove": {
      "script.ObjectCache": 0,
      "script.ScheduleActions": 1
    }
  },

```

```

    "findAndModify": {
      "script.Matches": 1,
      "script.AchievementTracker": 0,
      "script.ScheduleActions": 1
    }
  },
  "mongoCountScript": {
    "find": {
      "script.Matches": 1,
      "script.FieldPlayers": 1,
      "script.ScheduleActions": 1
    },
    "findOne": {
      "script.MatchSnapShot": 1,
      "script.ObjectCache": 1,
      "script.Sponsoring": 3,
      "script.Clubs": 14,
      "script.AchievementTracker": 3,
      "script.Leagues": 1,
      "player": 1
    },
    "save": {
      "script.ObjectCache": 1,
      "script.AchievementTracker": 1,
      "script.ScheduleActions": 7
    },
    "count": {
      "script.Matches": 2
    },
    "update": {
      "script.ClubLeagueStatistics": 5,
      "script.Leagues": 1,
      "script.SquadDynamic": 2
    },
    "remove": {
      "script.ObjectCache": 2,
      "script.ScheduleActions": 1
    },
    "findAndModify": {
      "script.Matches": 1,
      "script.AchievementTracker": 1,
      "script.ScheduleActions": 1
    }
  },
  "mongoTimeTotalScript": 16,
  "mongoCountTotalScript": 52
}

```

- DIMENSIONS**
- Time
 - Api
 - Api.sub Type
 - Api Key
 - Error
 - Platform.cluster
 - Platform.server
 - Stage
 - Type


- MEASURES**
- Count
 - Developer.mongo.operations.sum
 - Developer.mongo.time.sum
 - Developer.redis.operations.sum
 - Developer.redis.time.sum
 - Developer.script.duration.sum
 - Developer.script.statements.sum
 - Request.duration.sum
 - Request.size.sum
 - Response.size.sum

FILTER Latest Day

SPLIT Api Key Time (Hour)



Developer.mongo.operations.sum: 327.7 m



LEGEND Developer...

| API KEY | Count |
|--------------|---------|
| B298955Dhf4I | 183.8 m |
| 294937i173W5 | 124.9 m |
| 10009214mb0U | 15.7 m |
| 294090X0PqDS | 3.2 m |
| S299170zPE4N | 0.0 m |

Show selection

PINBOARD Count

Click or drag dimensions to pin them

Learnings

Minimise the Failure Domain

Give the benefit of the doubt

Think of the worst case scenario

Measure as much as you can

Questions?

greg.murphy@gamesparks.com