



Real-Time

Schema Discovery

Austin, TX March, 2023





Daniel Selans

20+ years of software and systems experience

Past 10 years spent writing backend Go

All about event driven, event sourcing and all things async

Reside in Portland, OR (originally from Riga, Latvia =)

PREVIOUSLY AT



Community

DigitalOcean





Greetings from Coachella

Why real-time schema discovery?

Why real-time schema discovery?

Need to start with some context

CONFIDENTIAL

What is a schema?

What is a schema? Pre-defined and agreed-upon

structure and layout for data

Data used by many teams; everyone needs data

Data used by many teams; everyone needs data

Data Engineers \rightarrow AI, machine learning, analytics

CONFIDENTIAL

Data used by many teams; everyone needs data

Data Engineers \rightarrow Al, machine learning, analytics

CONFIDENTIAL

DATA COUNCIL 2023

Developers \rightarrow Build features

Data used by many teams; everyone needs data

Data Engineers \rightarrow Al, machine learning, analytics

Platform & DevOps \rightarrow Troubleshooting/debug

CONFIDENTIAL

DATA COUNCIL 2023

Developers \rightarrow Build features

CONFIDENTIAL

Incidents & outages

Unexpected upstream schema changes impact downstream consumers

Incidents & outages

Unexpected upstream schema changes impact downstream consumers

Duplicate effort

Everyone normalizes data according to their requirements

Incidents & outages

Unexpected upstream schema changes impact downstream consumers

Duplicate effort

Everyone normalizes data according to their requirements DATA COUNCIL 2023

Wasted time Tool and process updates

Schema examples) cloudevents

"data":{ EVENT_DATA / }, "datacontenttype": "application/json", "id": "MESSAGE_ID / ", "source": "//storage.googleapis.com/projects/_/buckets/BUCKET_NAME / ", "specversion": "1.0", "type": "google.cloud.storage.object.v1.finalized", "time": "EVENT_GENERATION_TIME / ", "subject": "objects/my-file.txt",



Schema examples



**	
* An example protocol in Avro IDL	
*/	
mamespace("org.apache.avro.test")	
protocol Simple (
<pre>@aliases(["org.foo.KindOf"])</pre>	
enum Kind {	
F00,	
BAR, // the bar enum value	
BAZ	2000 200 - 00
<pre>} = F00; // For schema evolution purposes, unmatched values do not throw an</pre>	error, but are resolved
/** Holds MD5 hash; good enough to avoid most collisions, and smaller than	(for example) SHA256. */
fixed MD5(16);	
record TestRecord {	
string name;	
int age;	
float height;	
boolean isAlive;	
// Custom type	
Kind kind;	
// Custom type	
MD5 hash;	
// Array/Slice/List	
<pre>array<long> arrayOfLongs;</long></pre>	
// Union shorthand	
<pre>MD5? @aliases(["hash"]) nullableHash = null;</pre>	
1	

DATA COUNCIL 2023

solved to FOO.

Schema examples

43 lines (37 sloc) 857 Bytes

```
1 syntax = "proto3";
 2
3 package fakes;
 4
 5 import "billing.proto";
 6 import "billingv2.proto";
    import "coin.proto";
 8 import "post.proto";
    import "product.proto";
 9
10 import "search.proto";
11 import "user.proto";
    import "weather.proto";
12
13
14
     option go_package = "github.com/batchcorp/schemas/build/go/events/fakes";
15
16 enum EventType {
      EVENT_TYPE_UNSET = 0;
17
      EVENT_TYPE_BILLING = 1;
18
19
      EVENT_TYPE_SEARCH = 4;
20
      EVENT TYPE PRODUCTS = 10;
      EVENT_TYPE_USERS = 11;
21
22
      EVENT_TYPE_POSTS = 12;
      EVENT_TYPE_COINS = 13;
23
      EVENT_TYPE_WEATHER = 14;
24
25 }
26
27 message Event {
28
      EventType type = 1;
29
      int64 timestamp_nano = 2;
30
      string request_id = 3;
31
      string source = 4;
32
33
      oneof event {
34
        Billing billing = 100;
35
        Search search = 101;
        Product product = 102;
36
37
        User user = 103;
38
        Post post = 104;
39
        Coin coin = 105;
40
        Weather weather = 106;
41
        BillingV2 billing_v2 = 107;
42
     }
```

protobuf Protocol Buffers

```
3
 4
10
11
12
13
14
15
16
17
18
19
20
```

43 }

21 lines (17 sloc) 409 Bytes

```
1 syntax = "proto3";
2
     package fakes;
    option go_package = "github.com/batchcorp/schemas/build/go/events/fakes";
 5
 6
     // Used for fake event generation
7
 8
    message Billing {
       string order_id = 1;
9
      string order state = 2;
      string customer_id = 3;
       int64 order_date = 4;
       repeated BillingProduct products = 5;
    - 3
    message BillingProduct {
      string id = 1;
       string name = 2;
       int32 quantity = 3;
       float price = 4;
21 }
```

Conclusion

- Having a schema is a foundational component of your platform
- Lack of a schema will result in:
 - Waste of engineering resources
 - Slower feature & work delivery
 - Incidents & Outages
 - Difficult to debug
 - Slower mean time to recovery



STEP 1 Define schema

STEP1

Define schema

...but who's the source of truth?

CONFIDENTIAL

STEP1

Define schema

...but who's the source of truth?

STEP 2 Establish schema update process

STEP 1

Define schema

...but who's the source of truth?

STEP 2 Establish schema update process

...but what if someone doesn't follow it?

STEP1

Define schema

...but who's the source of truth?

STEP 2 Establish schema update process

...but what if someone doesn't follow it?

DATA COUNCIL 2023

STEP 3

Real-time Schema Discovery

Real-Time Schema Discovery In 17 easy steps

• •

CONFIDENTIAL







Determine and nominate one or more sources of truth for data (ie. Kafka) \checkmark

- Determine and nominate one or more sources of truth for data (ie. Kafka) \checkmark
- Write a service that consumes from Kafka

- Determine and nominate one or more sources of truth for data (ie. Kafka) \checkmark
- Write a service that consumes from Kafka \checkmark
- Write a library that parses JSON and generates a typed \checkmark schema (ie. parquet)

- Determine and nominate one or more sources of truth for data (ie. Kafka) \checkmark
- Write a service that consumes from Kafka
- Write a library that parses JSON and generates a typed \checkmark schema (ie. parquet)
- Research and implement schema conflict resolution \checkmark

- Determine and nominate one or more sources of truth for data (ie. Kafka) \checkmark
- Write a service that consumes from Kafka
- Write a library that parses JSON and generates a typed \checkmark schema (ie. parquet)
- Research and implement schema conflict resolution \checkmark
- Implement distributed schema election support in service (learn raft) \checkmark

- Determine and nominate one or more sources of truth for data (ie. Kafka) \checkmark
- Write a service that consumes from Kafka
- Write a library that parses JSON and generates a typed \checkmark schema (ie. parquet)
- Research and implement schema conflict resolution \checkmark
- Implement distributed schema election support in service (learn raft)
- Write a distributed schema-election service

- Determine and nominate one or more sources of truth for data (ie. Kafka)
- Write a service that consumes from Kafka
- Write a library that parses JSON and generates a typed \checkmark schema (ie. parquet)
- Research and implement schema conflict resolution
- Implement distributed schema election support in service (learn raft)
- Write a distributed schema-election service
- Figure out what to do during unresolvable schema conflicts (dead-letter?)



- Determine and nominate one or more sources of truth for data (ie. Kafka)
- Write a service that consumes from Kafka
- Write a library that parses JSON and generates a typed schema (ie. parquet)
- Research and implement schema conflict resolution
- Implement distributed schema election support in service (learn raft)
- Write a distributed schema-election service
- Figure out what to do during unresolvable schema conflicts (dead-letter?)
- And while you're at it, probably patent something you discovered while building this 😘



We actually did all of this

CONFIDENTIAL

We actually did all of this and so did other folks...

CONFIDENTIAL
Companies with schema inference support aWS snowflake Sdatabricks

ClickHouse 😡 BigQuery 😪 streamdal



And several others ...

CONFIDENTIAL

6 months to build solution

CONFIDENTIAL

- 6 months to build solution
- 4 months to get it "ready for prod"

- 6 months to build solution
- 4 months to get it "ready for prod"
- ✓ 3 months to deal with edge cases

- 6 months to build solution \checkmark
- 4 months to get it "ready for prod" \checkmark
- 3 months to deal with edge cases \checkmark
- 1 week every 2 months for troubleshooting/debug of "strange" behavior \checkmark

- 6 months to build solution \checkmark
- 4 months to get it "ready for prod" \checkmark
- 3 months to deal with edge cases \checkmark
- 1 week every 2 months for troubleshooting/debug of "strange" behavior \checkmark

2 principal-level engineers for 1.2 years (\$200K+/year) \checkmark

- 6 months to build solution \checkmark
- 4 months to get it "ready for prod" \checkmark
- 3 months to deal with edge cases \checkmark
- 1 week every 2 months for troubleshooting/debug of "strange" behavior \checkmark

2 principal-level engineers for 1.2 years (\$200K+/year) \checkmark

2 senior-level engineers for 8 hrs per month – ongoing support (~\$150K/year) \checkmark

- 6 months to build solution \checkmark
- 4 months to get it "ready for prod" \checkmark
- 3 months to deal with edge cases \checkmark
- 1 week every 2 months for troubleshooting/debug of "strange" behavior \checkmark

- 2 principal-level engineers for 1.2 years (\$200K+/year) \checkmark
- 2 senior-level engineers for 8 hrs per month ongoing support (~\$150K/year) \checkmark
- ~\$500K first year spend; ~\$40K/year ongoing spend \checkmark

Or in other words...

Oon't do it.

(probably)

CONFIDENTIAL



Instead, establish a **Foundational Schema**



CONFIDENTIAL

Protocol buffers are Google's language-neutral, platform-neutral, extensible mechanism for serializing structured data - think XML, but smaller, faster, and simpler.

Some other serialization formats you may have heard of: Thrift, Flatbuffer, Avro.

DATA COUNCIL 2023

PROTO DEFINITION

syntax = "proto3";

package flags;

message SubCommand 4 string action = 1; repeated string args = 2; int64 count = 3;

```
message Base {
  string action = 1;
  SubCommand sub = 2;
  bool debug = 3;
  string version = 4;
```

Step 2: Intro to protobuf

COMPILE SCHEMAS

You compile the .protos with `protoc` tool (part of protobuf toolchain)

 `protoc` generates "compiled" *.pb.go files which contain generated code for accessing the attributes and structures you defined in .protos

docker run --rm -w \$(PWD) -v \$(PWD):\$(PWD) -w\${PWD} jaegertracing/protobuf:0.2.0 \
--proto_path=./protos/args \
--go_out=plugins=grpc:\$(GO_PROTOS_DIR)/args \
--go_opt=paths=source_relative \
-o ./\$(GO_DESCRIPTOR_SET_DIR)/args.fds \
--include_imports \
--include_source_info \
protos/args/*.proto

SCREENSHOT OF COMPILING PROTOS IN DOCKER

Step 3: Intro to protobuf

What does the generated protobuf look like?



COMPILED

type SubCommand struct {
 Action string
 Args []string
 Count int64

```
type Base struct {
Action string
Sub SubCommand
Debug bool
Version string
```

Schema Management



CONFIDENTIAL





Protobuf schema lives in a Github repo

Repois governed, protected by architects & stakeholders



Schema changes are proposed via PR

STEP 4 Once merged, Github kicks off notification about schema update

CONFIDENTIAL

Stakeholders see schema update notification and update their components (with updated schema)

Backwards incompatible changes are NOT permitted (During PR process)



Consumers are deployed with new schema

CONFIDENTIAL



Producers are deployed with new schema

What does it look like in practice?

CONFIDENTIAL





All services use common schema

CONFIDENTIAL



Services encode ingres data as protobuf

Services emit encoded data to data bus (kafka, pulsar, nats, etc)

CONFIDENTIAL



Consumers read data from data bus and know how to decode data due to schema



Consumers do XYZ with decoded data

CONFIDENTIAL

1 Single source of truth

Everyone knows where to find out what data is supposed to look like.

Single source of truth

Everyone knows where to find out what data is supposed to look like.

2 No guesswork surrounding data structure

Team A does not need to find Person X on Team Y who knows about possible field values in a certain data structure.

Single source of truth

Everyone knows where to find out what data is supposed to look like.

2 No guesswork surrounding data structure

Team A does not need to find Person X on Team Y who knows about possible field values in a certain data structure. 3

No need for normalization and/or

Sanitization step Everyone knows where to find out

Everyone knows where to find out what data is supposed to look like.

Single source of truth

Everyone knows where to find out what data is supposed to look like.

2 No guesswork surrounding data structure

Team A does not need to find Person X on Team Y who knows about possible field values in a certain data structure.

4

Cost reduction

Less data to transport (ie. via Kafka) – 75% size reduction VS JSON Smaller memory footprint Lower network costs 3

No need for normalization and/or

Sanitization step Everyone knows where to find out

Everyone knows where to find out what data is supposed to look like.

Single source of truth

Everyone knows where to find out what data is supposed to look like.

2 No guesswork surrounding data structure

Team A does not need to find Person X on Team Y who knows about possible field values in a certain data structure.



Cost reduction

Less data to transport (ie. via Kafka) – 75% size reduction VS JSON Smaller memory footprint Lower network costs



Less Bugs

Reduce number of data structure related mistakes

3

No need for normalization and/or

Sanitization step Everyone knows where to find out

Everyone knows where to find out what data is supposed to look like.

BENEFITS

1

Single source of truth

Everyone knows where to find out what data is supposed to look like.

2 No guesswork surrounding data

structure

Team A does not need to find Person X on Team Y who knows about possible field values in a certain data structure.



Cost reduction

Less data to transport (ie. via Kafka) – 75% size reduction VS JSON Smaller memory footprint Lower network costs



Less Bugs

Reduce number of data structure related mistakes



Establish a solid foundation

for the future

Data structure affects all of engineering. Adopting a solid schema process will continue to pay huge dividends in every part of your org.

3 No need for normalization and/or sanitization step

Everyone knows where to find out what data is supposed to look like.

MORE IMPORTANTLY

Conclusion (1)

DIY schema discovery in distributed
 systems is **really** hard




DIY schema discovery in distributed systems is really hard

a. Hardest part == conflict resolution



DIY schema discovery in distributed
systems is **really** hard
a. Hardest part == conflict resolution

b. Huge time sink



DIY schema discovery in distributed systems is really hard
a. Hardest part == conflict resolution
b. Huge time sink
c. Only necessary if you use a

semi-structured data format (JSON)



- DIY schema discovery in distributed systems is really hard
 - a. Hardest part == conflict resolution
 - b. Huge time sink
 - c. Only necessary if you use a
 - semi-structured data format (JSON)
- Time better spent adopting a strong schema format (Protobuf)



CONFIDENTIAL



- DIY schema discovery in distributed systems is really hard
 - a. Hardest part == conflict resolution
 - b. Huge time sink
 - c. Only necessary if you use a
 - semi-structured data format (JSON)
- Time better spent adopting a strong schema format (Protobuf)
 - a. But... will require organization-wide adoption



 You should spend time on DIY schema discovery IF:



- You should spend time on DIY schema discovery IF:
 - a. You have **NO** control over ingres data structure



- You should spend time on DIY schema discovery IF:
 - a. You have **NO** control over ingres
 - data structure
 - b. ... And even then, the ingres data should be wrapped in protobuf



- You should spend time on DIY schema \checkmark discovery IF:
 - a. You have **NO** control over ingres
 - data structure
 - b. ... And even then, the ingres data should be wrapped in protobuf



Adopting protobuf is easier at small orgs



- You should spend time on DIY schema \checkmark discovery IF:
 - a. You have **NO** control over ingres
 - data structure
 - b. ... And even then, the ingres data should be wrapped in protobuf



- Adopting protobuf is easier at small orgs
 - a. Large orgs need to do piecemeal adoption (which is harder)



Thank you.



Daniel Selans Co-Founder & CTO / Streamdal

STREAMDAL.COM







daniel@streamdal.com

in

 \sim

linkedin.com/in/dselans/

github.com/dselans



github.com/batchcorp/plumber

