



Towards new data solutions

Views and opinions expressed **are those of the speakers** and not necessarily Novo Nordisk

Steffen Falgreen // Anders Bilgrau
2023-09-28

Current setup - Introduction to our CDR & SCE

What's the problem(s)?


What's a solution?

Results on data formats & database systems


Final notes

Setup

What's the purpose?





What's the delivery?




EDC: Electronic Data Capture; SDTM: Study Data Tabulation Model; ADaM: Analysis Data Model; TFL: Tables, Figures, Lists

SCE & CDR







Each bucket contains flat files

Our setup



- ae.sas7bdat
- ag.sas7bdat
- apae.sas7bdat
- apce.sas7bdat
- apcm.sas7bdat
- apdd.sas7bdat
- apdm.sas7bdat
- apface.sas7bdat
- apho.sas7bdat
- apmh.sas7bdat
- apqs.sas7bdat
- aprebsub.sas7bdat



- adadj.sas7bdat
- adae.sas7bdat
- adcgmen.sas7bdat
- adcgms.sas7bdat
- adcm.sas7bdat
- adcen.sas7bdat
- addec.sas7bdat
- addeg.sas7bdat
- adhypo.en.sas7bdat
- adhypo.sas7bdat
- advs.sas7bdat
- axe.sas7bdat




- acgmfas.txt
- acgmfas.xml
- ahyposfas.txt
- ahyposfas.xml
- faesas.png
- faesas.xml
- flbbiosas.png
- flbbiosas.xml
- tvsbptsas.txt
- tvsbpsas.xml



- demographic-data.docx
- efficacy-data.docx
- adverse-events.docx
- deaths-saes.docx
- lab-displays.docx


Data organization - current



What's the problem(s)?

THINGS
ARE SLOW

Typically 50-100 times faster reading speeds



rds vs. sas7bdat (via haven)

Secondary issues

Hard to automate


Data not easily
accessible from all
languages

Solution does not
scale well


Very big data
does not work
(larger-than-
memory)

What should a solution look like?


Data organization - current



Data organization – desired alternative





Execution & use of data within trial



Directed Acyclic Graphs only!

Execution & use of data within trial





Workloads

OLTP or **OLAP**
(Online Transactional Processing) (Online Analytical Processing)

Tabular data

Logical model

A1	B1	C1
A2	B2	C2
A3	B3	C3
A4	B4	C4

Physical model

Row-major

A1	B1	C1	A2	B2	B2	A3	B3	C3	A4	B4	C4
----	----	----	----	----	----	----	----	----	----	----	----

Column-major

A1	A2	A3	A4	B1	B2	B3	B4	C1	C2	C3	C4
----	----	----	----	----	----	----	----	----	----	----	----

Tabular data

Logical model

A1	B1	C1
A2	B2	C2
A3	B3	C3
A4	B4	C4

Physical model

Row-major

A1	B1	C1	A2	B2	B2	A3	B3	C3	A4	B4	C4
----	----	----	----	----	----	----	----	----	----	----	----

Column-major

A1	A2	A3	A4	B1	B2	B3	B4	C1	C2	C3	C4
----	----	----	----	----	----	----	----	----	----	----	----

Tabular data

Logical model

A1	B1	C1
A2	B2	C2
A3	B3	C3
A4	B4	C4

Physical model

Row-major

A1	B1	C1	A2	B2	B2	A3	B3	C3	A4	B4	C4
----	----	----	----	----	----	----	----	----	----	----	----

Column-major

A1	A2	A3	A4	B1	B2	B3	B4	C1	C2	C3	C4
----	----	----	----	----	----	----	----	----	----	----	----


OLTP

(Online Transactional Processing)

or

OLAP


(Online Analytical Processing)



read

(read intensive system)


or



write

(write intensive system)

Properties



Formats for storing data

Compare different file formats stored in the cloud

Focus on ADaM data



- adadj.sas7bdat
- adae.sas7bdat
- adcgmen.sas7bdat
- adcgmm.sas7bdat
- adcm.sas7bdat
- adecen.sas7bdat
- adec.sas7bdat
- adeg.sas7bdat
- adhypo.en.sas7bdat
- adhypo.sas7bdat
- advs.sas7bdat
- axe.sas7bdat

adcgm.sas7bdat


USUBJID	PARAMCD	AVISITN	ADTM	STUDYID	SUBJID	SITEID	STRATA	...	PARAM	PARAMN	AVAL
Unique Subject Identifier	Parameter Code	Analysis Visit (N)	Analysis Datetime	Study Identifier	Subject Identifier for the Study	Study Site Identifier	Randomized Strata	...	Parameter	Parameter (N)	Analysis Value
NN1234-4321/101002	C105585M	400	2019-04-04 09:28:00	NN1234-4321	101002	101	DPP-4 INHIBITOR	...	CGM - Meter value (mmol/L)	600	7.2
NN1234-4321/101002	C105585M	400	2019-04-04 09:43:00	NN1234-4321	101002	101	DPP-4 INHIBITOR	...	CGM - Meter value (mmol/L)	600	7.2
NN1234-4321/101002	C105585M	400	2019-04-04 09:58:00	NN1234-4321	101002	101	DPP-4 INHIBITOR	...	CGM - Meter value (mmol/L)	600	7.2
NN1234-4321/101002	C105585M	400	2019-04-04 10:13:00	NN1234-4321	101002	101	DPP-4 INHIBITOR	...	CGM - Meter value (mmol/L)	600	7.6
NN1234-4321/101002	C105585M	400	2019-04-04 10:28:00	NN1234-4321	101002	101	DPP-4 INHIBITOR	...	CGM - Meter value (mmol/L)	600	8.6
NN1234-4321/101002	C105585M	400	2019-04-04 10:43:00	NN1234-4321	101002	101	DPP-4 INHIBITOR	...	CGM - Meter value (mmol/L)	600	9.7
NN1234-4321/101002	C105585M	400	2019-04-04 10:58:00	NN1234-4321	101002	101	DPP-4 INHIBITOR	...	CGM - Meter value (mmol/L)	600	8.3
NN1234-4321/101002	C105585M	400	2019-04-04 11:13:00	NN1234-4321	101002	101	DPP-4 INHIBITOR	...	CGM - Meter value (mmol/L)	600	8.7

Rows: 2,804,396; columns: 62; Sas7bdat: 6.3 GB

Formats chosen for comparison



C, S, V







Next generation transport format

Dataset-JSON

Dataset-JSON



Actual Metadata from data set



```
{
  "creationDateTime": "2023-03-22T11:53:27",
  "datasetJSONVersion": "1.0.0",
  "fileOID": "www.sponsor.org.project123.final",
  "asOfDateTime": "2023-02-15T10:23:15",
  "originator": "Sponsor XYZ",
  "sourceSystem": "Software ABC",
  "sourceSystemVersion": "1.2.3",
  "clinicalData": {
    "studyOID": "xxx",
    "metaDataVersionOID": "xxx",
    "metaDataRef": "https://metadata.location.org/api.link",
    "itemGroupData": {
      "IG.DM": {
        "records": 600,
        "name": "DM",
        "label": "Demographics",
        "items": [
          {"OID": "ITEMGROUPDATASEQ", "name": "ITEMGROUPDATASEQ", "label": "Record identifier", "type": "integer"},
          {"OID": "IT.STUDYID", "name": "STUDYID", "label": "Study identifier", "type": "string"},
          {"OID": "IT.USUBJID", "name": "USUBJID", "label": "Unique Subject Identifier", "type": "string"},
          {"OID": "IT.DOMAIN", "name": "DOMAIN", "label": "Domain Identifier", "type": "string"},
          {"OID": "IT.AGE", "name": "AGE", "label": "Subject Age", "type": "integer"}
        ]
      },
      "itemData": [
        [1, "MyStudy", "001", "DM", 56],
        [2, "MyStudy", "002", "DM", 26],
        ...
        [600, "MyStudy", "900", "DM", 45],
      ]
    }
  }
}
```

Dataset-JSON

Logical model

TRIALID	PARAM CD	TIME	AVAL
A1	B1	C1	D1
A2	B2	C2	D2
A3	B3	C3	D3
A4	B4	C4	D4

Physical model

Row-major


A1	B1	C1	D1	A2	B2	C2	D2	A3	B3	C3	D3	A4	B4	C4	D4
----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----



Formats chosen for comparison



C, S, V





- Free and open-source file format
- **Good for storing big data of any kind**
 - **Structured data tables**, images, videos, documents

Characteristics:

- Utilizes schema, Contains metadata on data, Can contain user specified metadata, Supports different levels of compression, On-disk

Tabular data

Logical model

TRIALID	PARAM CD	TIME	AVAL
A1	B1	C1	D1
A2	B2	C2	D2
A3	B3	C3	D3
A4	B4	C4	D4

Physical model

Column-major

A1	A2	A3	A4	B1	B2	B3	B4	C1	C2	C3	C4	D1	D2	D3	D4
----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----

RLE

RLE

Delta encoding

Plain

Encoding: NN1234-4321 * 4

C105585M * 4

TIME diff diff diff 7 5 7 8




Learnings & tests



File size by format



File size by format




Where do we store the files



FSx for lustre



- Pros
 - Cost-effective storage
 - Entirely managed service
 - Supports more than 16 groups
 - Works in combination with Simple Storage Service (S3)
 - The most used file system for super computers
- Cons
 - Specify size upfront





Elastic File Storage

- Pros
 - Cost-effective storage
 - Entirely managed service
 - Set-and-forget
 - Grows with your data
 - Only pay for what you use
- Cons
 - No support for more than 16 groups

Read and write by type of drive




Read and write by type of drive



Partial conclusion file storage systems



FSx for lustre



Elastic File Storage



Speed-up is format dependent



Always faster reads

(although speed-up often is not substantial)

Approx 2× write


Performance of formats



A single format may have many readers / writers




A single format may have many readers / writers




Read and write csv by different implementations



Read



Write



Performance of formats



Read and write of adcg

Read



0.27



28



46



71



221

adcgm (sec):

Read and write of adcg

Read



0.27



28



46



71



221

Write



8



27



45



110





148



adcgm (sec):

adcgm (sec):

Read and write on FSX




Read and write on FSX



Partial data reads

Queryability (yes that it as word, no I will not say it)

- Only read what is necessary
- Example with **CGM** stored in parquet:
 - Read data
 - Filter on subjects < 56 years
 - Calculate summary statistics for each individual



Comparison to DBMS



PostgreSQL

Read and write for postgresSQL

Driver comparison




R specific postgresSQL driver



RPostgres



Open Database Connectivity





ODBC

Read and write for postgresSQL

Driver comparison



Read and write complete data set



Possible use as transport format

Wishes for a transport format

Support
metadata

Support
for rules

Human
readable


Small file
size

Wishes for a transport format

Small file
size




File size by format (adcgm)



Each file must be 5 GB or less




SAS V5 transport: 50 GB (conventional/could be double)




Authorities

Each file must be 5 GB or less




File size by format



Each file must be 5 GB or less




Dataset-JSON: 30 GB




Authorities

Each file must be 5 GB or less




Each file must be 5 GB or less



R supports characters of length $2^{31} - 1$


File size by format



Each file must be 5 GB or less





Parquet: 1 GB



Authorities

Parquet supports automatic split to subfiles by row-value or maximum file size

Read and write on FSX



Wishes for a transport format

	Support metadata	Support for rules	Human readable	Small file size
	✓	✓	✗	✓
	✓	✓	✗	✗
	✓	✓	✓	✗

Final notes




	CSV	sas7bdat	JSON	rds	Parquet
Fast	✓	✗	✗	✓	✓
Storage-efficient	✗	✗ (✓)	✗	✓	✓
Open & interoperable	✗	✗	✓	✗	✓
Metadata-support	✗	✓	(✓)	✓	✓
Rule-support	✗	✗	(✓)	✗	✓
Queryable	✗	(✓)	✗	✗	✓
Stable checksums	✓	✗	✓	✓	✓

Deliverable format $\neq$ internal format

- Delivery and transport formats does **not** need to be identical to internal production
 - Internal and external needs may be very different
 - Formats are optimized for different things
 - Internal format may be optimized for speed, automation, ease for use
 - External delivery may be optimized for cross-industry harmonization, transparency,
 - **If** they differ, import/export capabilities are obviously needed

Trade-offs

- Compression
 - Saves space but costs time
 - Compression/decompression time should not offset reduced transfer time
- Implementation matters (features/support/speed)
- Read vs write
- Languages
 - **R** is designed around column-major ordering
 - **SAS** (data steps) works is row-major ordering
 - **Python**
 - **numpy**: row-major
 - **Pandas**: column-major



Conclusions

- **Parquet** seems to be superior by almost any metric **if** you want to stick with files
 - Enables efficient queries on data unlike other file formats
 - *Disclaimer: SAS is claimed to work with parquet but **we** have not tested yet*
 - **Future-proof:** Compatible with most high-performance big-data processing and storage frameworks (Hadoop, Spark, **Delta Lake**, ...)
- Relational databases like Postgres would be nice, but has limitations in terms of organization (only supports two levels of “name-spacing”)
 - Primary *functional* appeal from database management systems (DBMS) are queries (now diminished).
 - Consider change management! People *really* like their files.
 - DBMS' may be too hard to learn



