

# CS460: Intro to Database Systems

## Class 26: NoSQL Systems

Instructor: Manos Athanassoulis

<https://bu-disc.github.io/CS460/>

# What is NoSQL?

from "Geek and Poke"

## HOW TO WRITE A CV



Leverage the NoSQL boom

# What is NoSQL?

An emerging “movement” around non-relational software for Big Data

Roots are in the Google, Amazon, Facebook homegrown software stacks



A NoSQL system provides a **mechanism for storage and retrieval of** data that uses looser consistency models than traditional relational databases in order to achieve horizontal scaling and higher availability.

NoSQL comes from “Not SQL” or “Not only SQL” to emphasize that some NoSQL systems allow SQL-like queries.

# NoSQL Stores

offer an easy to program storage model

**simplification of relational**

two attributes (a key and a value)

value has variable size

# NoSQL features

## Scalability is crucial!

- load increased rapidly for many applications

## Large servers are expensive

## Solution: use clusters of small commodity machines

- need to **shard** the data (maybe use replication)
- cheap (usually open source!)
- cloud-based storage

```
create table people (
  user varchar(24),
  first varchar(24),
  last varchar(24),
  shard key (first)
);
```

people

people_0			people_1			people_2			people_3		
user	first	last	user	first	last	user	first	last	user	first	last
areece	alex	reece	jdoh	john	doe	ndrew	nancy	drew	amonson	amanda	monson
anick	alex	nick	jsmith	john	smith						
amace	alex	mace									
tholmes	tom	holmes									

image from: [singlestore.com](https://singlestore.com)

# NoSQL features

Sometimes not a well-defined schema

Allow for semi-structured data

- still need to provide ways to query efficiently (use of index methods)
- need to express specific types of queries easily

# Scalability

Often cited as the main reason for moving from DB technology to NoSQL

DB Position: there is no reason a parallel DBMS cannot scale to 1000's of nodes

NoSQL Position: a) Prove it; b) it will cost too much anyway

# Flavors of NoSQL

Four main types:

*key-value* stores

*document* databases

*column-family* (aka big-table) stores

*graph* databases

Here we will talk more about “Document” databases (MongoDB)



# Key-Value Stores

There are many systems like that:



redis



**amazon**  
DynamoDB

Simple data model: key/value pairs

the DBMS *does not attempt to interpret* the value

Queries are limited to *query by key*

- get/put/update/delete a key/value pair
- iterate over key/value pairs

# Document Databases

Examples include:



Special type of key/value that ***value is a document***.

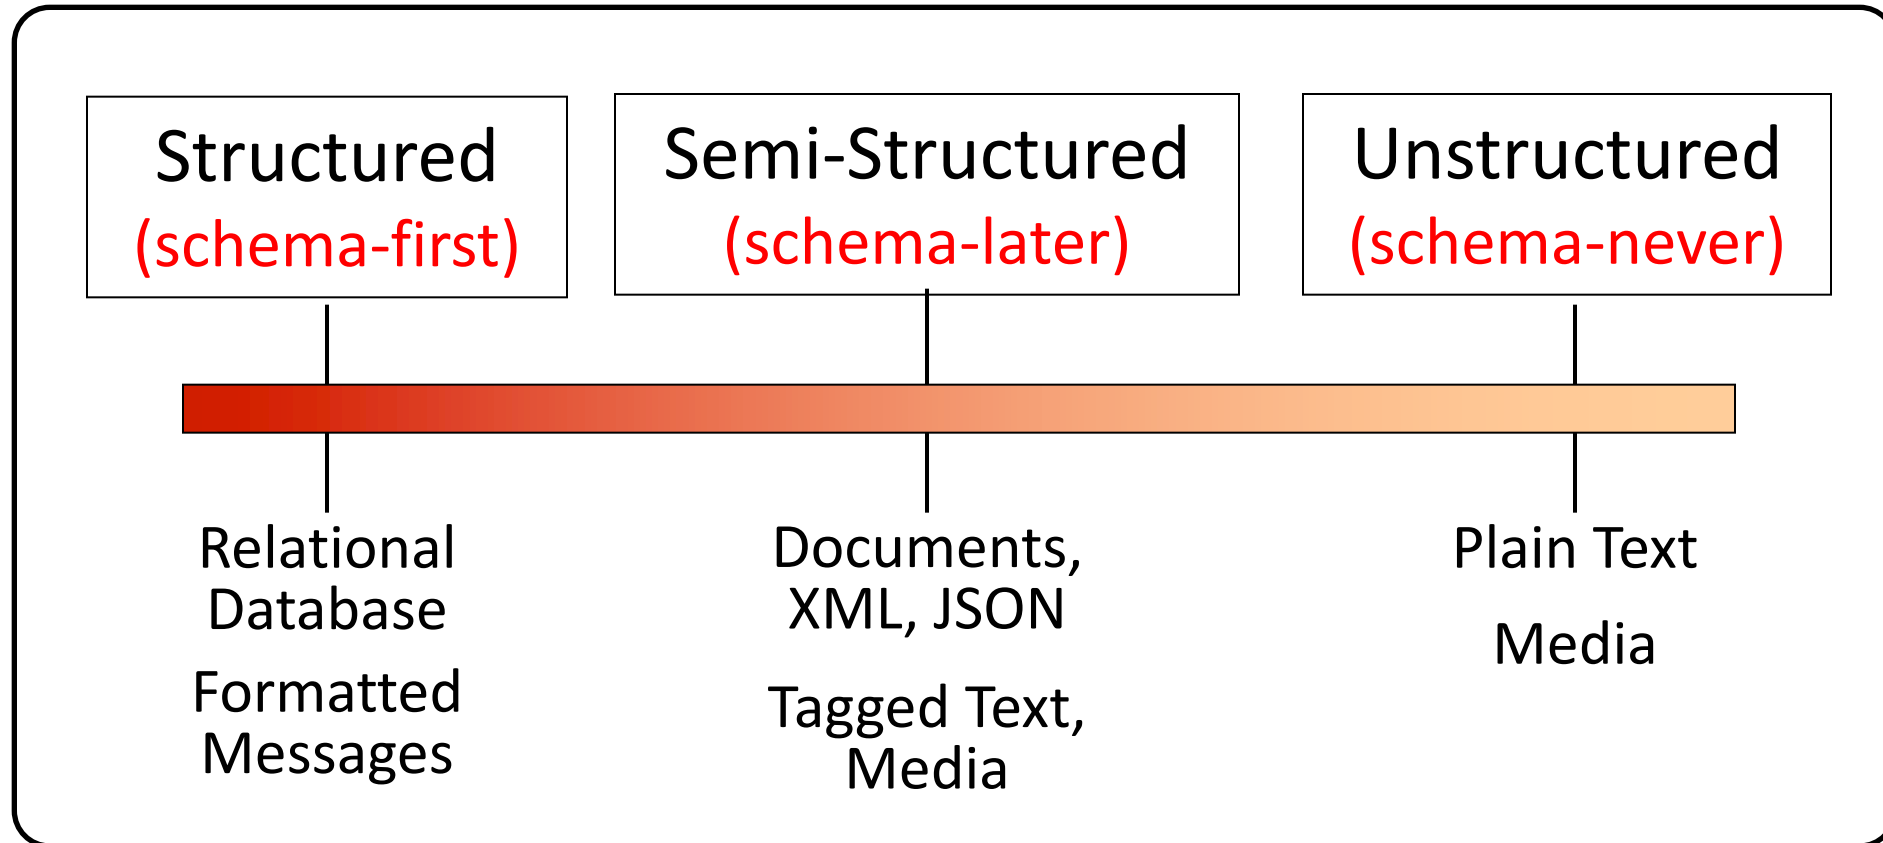
- use some sort of ***semi-structured data model***: XML/JSON
- the ***value can be examined*** and used by the system (unlike in key/data stores)

Queries based on key (as in key/value stores), but ***also on the document*** (value).

Here again, there is support for ***sharding*** and ***replication***.

- the sharding can be based on values within the document

# The Structure Spectrum



# Recap: NoSQL

**simplification of relational: {key, value}**

unique key

variable-size value

Key/Value stores:

- get/put/update/delete a key/value pair
- iterate over key/value pairs

Document stores (value is a ***semi-structured*** document ):

- use some sort of semi-structured data model: XML/JSON
- the value can be examined and used by the system (unlike in key/data stores)
- queries based on key (as in key/value stores), but also on the document (value).

# MongoDB (An example of a Document Database)

Data are organized in *collections*. A collection stores a *set of documents*.

Collection (*like table*) and document (*like record*)

- **BUT** each document can have different attributes even in the same collection
- Semi-structured schema!

Only requirement: every document should have an “\_id” field

- humongous => Mongo

# Example MongoDB

```
{ "_id": ObjectId("4efa8d2b7d284dad101e4bc9"),  
  "Last Name": " Cousteau",  
  "First Name": " Jacques-Yves",  
  "Date of Birth": "06-1-1910" },  
  
{ "_id": ObjectId("4efa8d2b7d284dad101e4bc7"),  
  "Last Name": "PELLERIN",  
  "First Name": "Franck",  
  "Date of Birth": "09-19-1983",  
  "Address": "1 chemin des Loges",  
  "City": "VERSAILLES" }
```

# Example Document Database: MongoDB

Key features include:

JSON-style documents

– actually, uses BSON (JSON's binary format)

***replication*** for high availability

***auto-sharding*** for scalability

***key & document-based*** queries

can create an index on any attribute for faster reads

under the hood, a simple key-value store called WiredTiger!  
design based on LSM-trees

# MongoDB Terminology

relational term  $\Leftrightarrow$  MongoDB equivalent

---

database  $\Leftrightarrow$  database

table  $\Leftrightarrow$  collection

row  $\Leftrightarrow$  document

attributes  $\Leftrightarrow$  fields (field-name:value pairs)

primary key  $\Leftrightarrow$  the `_id` field, which is the key associated with the document



# JSON

JSON is an alternative data model for semi-structured data

- JavaScript Object Notation

Built on two key structures:

- an **object**, which is a sequence of name/value pairs  
`{ "_id": "1000", "name": "Sanders Theatre", "capacity": 1000 }`
- an **array of values** [ "123", "222", "333" ]

A **value** can be:

- an atomic value: string, number, true, false, null
- an object
- an array

# The `_id` Field

Every MongoDB document must have an `_id` field.

its value must be unique within the collection

acts as the primary key of the collection

it is the key in the key/value pair

If you create a document without an `_id` field:

MongoDB adds the field for you

assigns it a unique BSON (binary JSON) ObjectID

example from the MongoDB shell:

```
> db.test.save({ rating: "PG-13" })  
> db.test.find() { "_id" : ObjectId("528bf38ce6d3df97b49a0569"), "rating" : "PG-13" }
```

Note: quoting field names is optional (see `rating` above)

# Capturing Relationships in MongoDB

Two options:

1. store references to other documents using their `_id` values
2. embed documents within other documents

# Example relationships

```
{
  "_id":ObjectId("52ffc33cd85242f436000001"),
  "name": "Tom Benzamin ",
  "contact": "987654321",
  "dob": "01-01-1991"
}
{
  "_id":ObjectId("52ffc4a5d85242602e000000"),
  "building": "22 A, Indiana Apt",
  "pincode": 123456,
  "city": "Los Angeles",
  "state": "California"
}
```

## Here an example of reference-based relationship

```
{
  "_id":ObjectId("52ffc33cd85242f436000001"),
  "contact": "987654321",
  "dob": "01-01-1991",
  "name": "Tom Benzamin",
  "address_ids": [
    ObjectId("52ffc4a5d85242602e000000")
  ]
}
```

## And, here is an example of embedded relationship:

```
{
  "_id":ObjectId("52ffc33cd85242f436000001"),
  "contact": "987654321",
  "dob": "01-01-1991",
  "name": "Tom Benzamin",
  "address": [
    {
      "building": "22 A, Indiana Apt",
      "pincode": 123456,
      "city": "Los Angeles",
      "state": "California"
    },
    {
      "building": "170 A, Acropolis Apt",
      "pincode": 456789,
      "city": "Chicago",
      "state": "Illinois"
    }
  ]
}
```

# Queries in MongoDB

Each query can only access a single collection of documents.

Use a method called

```
> db.collection.find(<selection>, <projection>)
```

**Example:** find the names of all R-rated movies:

```
> db.movies.find({ rating: 'R' }, { name: 1 })
```

# Projection

Specify the name of the fields that you want in the output with 1 ( 0 hides the value)

Example:

```
> db.movies.find({}, {"title":1, _id:0})
```

(will report the title but not the id)

# Selection

You can specify the condition on the corresponding attributes using the find:

```
> db.movies.find({ rating: "R", year: 2000 }, { name: 1, runtime: 1 })
```

Operators for other types of comparisons:

MongoDB	SQL equivalent
\$gt, \$gte	>, >=
\$lt, \$lte	<, <=
\$ne	!=

**Example:** find the names of movies with an earnings <= 200000

```
> db.movies.find({ earnings: { $lte: 200000 } })
```

For logical operators \$and, \$or, \$nor

use an array of conditions and apply the logical operator among the array conditions:

```
> db.movies.find({ $or: [ { rating: "R" }, { rating: "PG-13" } ] })
```

# Aggregation

Recall the aggregate operators in SQL: AVG(), SUM(), etc.

More generally, aggregation involves computing a result from a collection of data.

MongoDB supports several approaches to aggregation:

- single-purpose aggregation methods
- an aggregation pipeline
- map-reduce

Aggregation pipelines are more flexible and useful (see next):

<https://docs.mongodb.com/manual/core/aggregation-pipeline/>



# Simple Aggregations

## **db.collection.count(<selection>)**

returns the number of documents in the collection  
that satisfy the specified selection document

**Example:** how many R-rated movies are shorter than 90 minutes?

```
> db.movies.count({ rating: "R", runtime: { $lt: 90 } })
```

## **db.collection.distinct(<field>, <selection>)**

returns an array with the distinct values of the specified field  
in documents that satisfy the specified selection document  
if omit the query, get all distinct values of that field

**Example:** which actors have been in one or more of the top 10 grossing movies?

```
> db.movies.distinct("actors.name", { earnings_rank: { $lte: 10 } })
```

# Aggregation Pipeline

A very powerful approach to write queries in MongoDB is to use pipelines.

We execute the query in *stages*.

Every stage gets as *input some documents*, applies filters/aggregations/projections and *outputs some new documents*.

These documents are the input to the next stage (next operator) and so on

**Similar to a traditional query plan. But always with one child (no joins!)**

# Aggregation Pipeline example

Example for the zipcodes database:

```
> db.zipcodes.aggregate( [
  { $group: { _id: "$state", totalPop: { $sum: "$pop" } } },
  { $match: { totalPop: { $gte: 10000000 } } }
] )
```

```
{
  "_id": "10280",
  "city": "NEW YORK",
  "state": "NY",
  "pop": 5574,
  "loc": [
    -74.016323,
    40.710537
  ]
}
```

Here we use `group_by` to group documents per state, compute sum of population and output documents with `_id`, `totalPop` (`_id` has the name of the state). The next stage finds a match for all states the have more than 10M population and outputs the state and total population.

More here: <https://docs.mongodb.com/manual/tutorial/aggregation-zip-code-data-set/>

continued:

In SQL:

Output example:

```
{
  "_id" : "NY",
  "totalPop" : 19750000
}
```

```
SELECT state, SUM(pop) AS totalPop
FROM zipcodes
GROUP BY state
HAVING totalPop >= (10000000)
```

```
db.zipcodes.aggregate( [
  { $group: { _id: "$state", totalPop: { $sum: "$pop" } } },
  { $match: { totalPop: { $gte: 10000000 } } }
] )
```

# more examples:

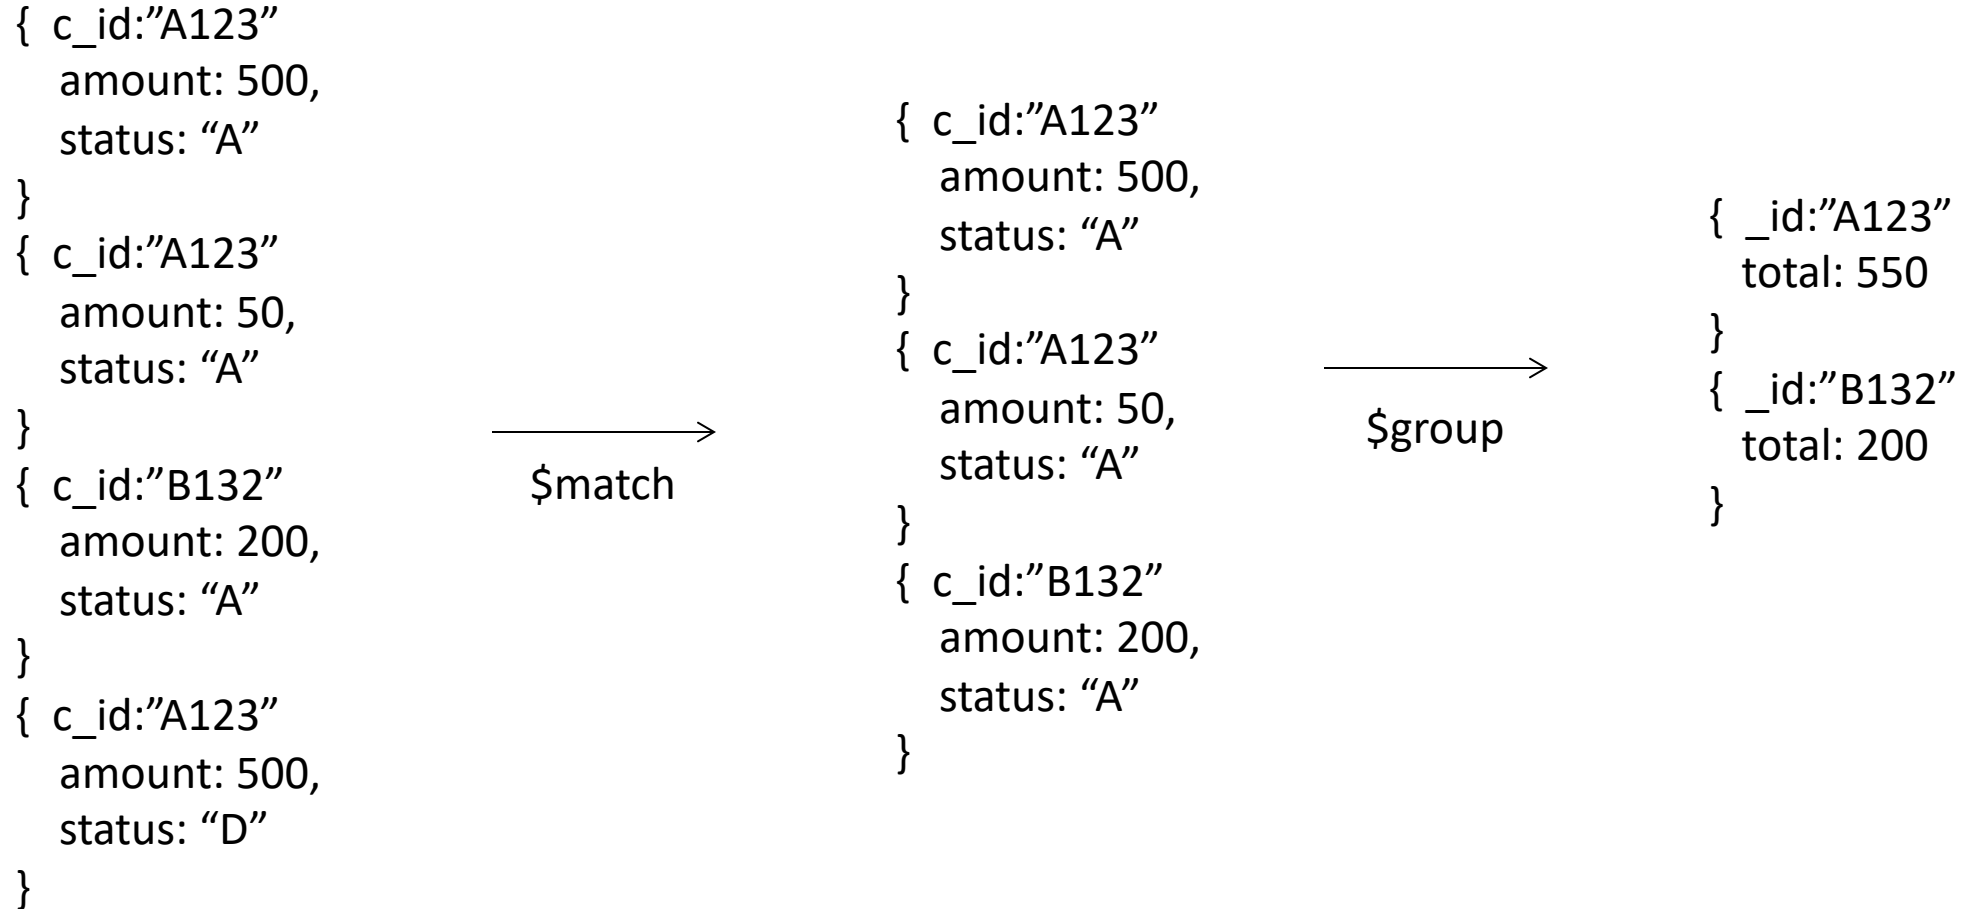
```
db.zipcodes.aggregate( [
  { $group: { _id: { state: "$state", city: "$city" }, pop: { $sum: "$pop" } } },
  { $group: { _id: "$_id.state", avgCityPop: { $avg: "$pop" } } }
] )
```

What we compute here?

First we get groups by city and state and for each group we compute the population.  
Then we get groups by state and compute the average city population

```
{
  "_id" : {
    "state" : "MN",
    "city" : "EDGEWATER"
  },
  "pop" : 13154
}
      →
{
  "_id" : "MN",
  "avgCityPop" : 5335
}
```

# Aggregation Pipeline example



```

db.orders.aggregate([ { $match: {status: "A"}}
                      { $group: {_id:"c_id", total: {$sum: $amount}}
                      ])
  
```

# Demo

<https://mongoplayground.net/>

# Other Structure Issues

## NoSQL

- a) Tables are unnatural
- b) “joins” are evil
- c) need to be able to “grep” my data

## DB

- a) Tables are a natural/neutral structure
- b) data independence lets you precompute joins under the covers
- c) this is a price of all the DBMS goodness you get

This is an Old Debate – Object-oriented databases, XML DBs, Hierarchical, ...



# Fault Tolerance

DBs: coarse-grained FT – if trouble, restart transaction

- Fewer, Better nodes, so failures are rare
- Transactions allow you to kill a job and easily restart it

NoSQL: Massive amounts of cheap HW, **failures are the norm** and massive data means **long running jobs**

- So must be able to do mini-recoveries
- This causes some overhead (file writes)

## CS460: Intro to Database Systems

# Database Systems and Beyond

Instructor: Manos Athanassoulis

<https://midas.bu.edu/classes/CS460/>

# Database Systems


we spent a whole semester on Database Systems  
what is next?

*what can we do with data?*

data-driven science

data-driven discovery

data-driven governance



*“Experimental, theoretical, and computational science are all being affected by the data deluge, and a fourth, ‘**data-intensive**’ science paradigm is emerging.*”

*The goal is to have a world in which all of the science literature is online, all of the science data is online, and they interoperate with each other. **Lots of new tools are needed to make this happen.**”*

Faster Innovation through Data-Intensive Approaches

Need for Innovation in Data Management!



# DATA & AI LANDSCAPE 2019

## INFRASTRUCTURE

**HADOOP ON-PREMISE**  
 cloudera Hortonworks  
 MAPR Pivotal  
 IBM InfoSphere  
 jethro

**HADOOP IN THE CLOUD**  
 aws Microsoft Azure  
 Google Cloud  
 SAP Cloud Platform  
 IBM InfoSphere BigInsights  
 arm  
 Dubble CAZENA

**STREAMING / IN-MEMORY**  
 Amazon Kinesis  
 databricks  
 SAP Cloud Platform  
 ORACLE  
 confluent  
 strimzi  
 hazelcast  
 GridGain  
 GIGASPACEs  
 Wallaroo.ai  
 FASTDATA  
 KX

**NoSQL DATABASES**  
 Google Cloud AWS  
 ORACLE  
 mongoDB MarkLogic  
 Couchbase DISTRAX  
 redislabs  
 KEROPIKE  
 ArangoDB SCYLLA

**NewSQL DATABASES**  
 SAP Clustrix  
 Pivotal  
 MEMSQL infuxdata  
 Cockroach LABS  
 VOLTDB splice  
 paradigms  
 IBM Amazon Neptune  
 ORACLE  
 OrientDB  
 InfiniteGraph  
 Objectivity

**GRAPH DBs**  
 neo4j  
 Amazon Neptune  
 ORACLE  
 OrientDB  
 InfiniteGraph  
 Objectivity

**MPP DBs**  
 TERADATA  
 VERTICA  
 IBM Data Warehouse Systems  
 Kognitio  
 Exasol  
 dremio  
 Yellowbrick

**CLOUD EDW**  
 aws  
 Google Cloud  
 Microsoft Azure  
 Pivotal  
 snowflake  
 Infoworks

**SERVERLESS**  
 Amazon Lambda  
 Google Cloud  
 Microsoft Azure  
 PULSAR  
 nuclio  
 PaaS Function Service

**DATA TRANSFORMATION**  
 talend pentaho  
 alteryx TRIFACTA  
 tamr Paxata  
 StreamSets UNIFI

**DATA INTEGRATION**  
 SAP Data Services Informatica  
 MuleSoft TEALUM  
 snaplogic enigma  
 Segment ATTUNITY  
 xplenty ZALONI import.io  
 Infosys Fivetran  
 SNOWFLOW MATILLION

**DATA GOVERNANCE**  
 Informatica  
 SailPoint  
 IBM  
 collibra  
 Alation  
 HMDTA  
 OKERA  
 MANTA  
 dataworld

**MGMT / MONITORING**  
 aws New Relic actifio  
 rubrik  
 dynatrace  
 SignalFx  
 splunk  
 Moogsoft pagerduty  
 Unavai Numerify  
 zentao OpsRamp  
 APPDYNAMICS  
 WAVEFRONT  
 druuu  
 pageduty  
 VEEAM  
 MAGNITUDE

**STORAGE**  
 aws  
 Google Cloud  
 Microsoft Azure  
 PURE STORAGE  
 ALLUXIO wasabi  
 nimbustorage  
 Qumulo panache  
 COHERITY

**CLUSTER SVCS**  
 Amazon ELB  
 IBM  
 Microsoft Azure  
 MESSORANGE  
 packet  
 nimbustorage  
 Qumulo panache  
 COHERITY

**DATA GENERATION & LABELLING**  
 amazon mechanicalturk  
 upwork  
 appen  
 HIVE  
 Labelbox  
 Mighty AI  
 ALBERT  
 LIONBRIDGE

**AI OPS**  
 ALGORITHMIA  
 SPILL comet  
 Verta.ai datmo  
 distatlon  
 Whiggle AI  
 Determined AI  
 fiddler

**GPU DBs & CLOUD**  
 kinetica  
 SQUEAM  
 bryllit  
 BLAZINGDB  
 PG-Ström  
 LOYDHUB

**HARDWARE**  
 Google TPU arm  
 intel AI NVIDIA  
 GRAPHCORE  
 MYTHIC  
 Movidius habana  
 WAVE  
 CERNAMI  
 PALADIN  
 INTEL  
 DEFINIX

## CROSS-INFRASTRUCTURE/ANALYTICS

aws Google Cloud Microsoft IBM SAP Hewlett Packard Enterprise SAS IOIGDATA vmware TIBCO TERADATA ORACLE NetApp syncsort MAPR cloudera

## ANALYTICS & MACHINE INTELLIGENCE

**DATA ANALYST PLATFORMS**  
 Microsoft pentaho alteryx  
 Digital Reasoning  
 GUAVUS AYASDI  
 ATTIVO Datameer incorta  
 interana MODE ENDOR  
 sisu switchboard Starburst

**DATA SCIENCE PLATFORMS**  
 IBM databricks dataiku  
 DOMINO rapidminer TIBCO  
 ANACONDA SAS  
 KNIME MathWorks

**BI PLATFORMS**  
 looker  
 DOME ARCADIA DATA ThoughtSpot  
 ATSCALE Qlik  
 GoodData Information Builders birst  
 MicroStrategy Keon IO

**VISUALIZATION**  
 tableau Power BI  
 SAP  
 Google Cloud  
 celonis  
 Periscope Data  
 zepl  
 CHARTIO TOUCAN TCOO

**MACHINE LEARNING**  
 Azure Machine Learning  
 Amazon SageMaker  
 Google Cloud AI Platform Vision H2O  
 DataRobot gamalon  
 VISENZE ELEMENT  
 deepsense.ai

**COMPUTER VISION**  
 Microsoft Azure  
 Amazon Rekognition  
 clarifai  
 EVERAI deepomatic  
 neuro  
 UBUIQUITY ABBE  
 YITU trax  
 synthesis DataClz

**HORIZONTAL AI**  
 IBM Watson Cortana Face++  
 sentient Voyager  
 Affictive  
 Numenta  
 narologics  
 BLUE VISION  
 II  
 Fortress

**SPEECH & NLP**  
 Google Cloud twilio  
 amazon alexa Amazon Transcribe  
 narrative science semantic machines  
 Mobval  
 SoundHound Inc PRIMER  
 Mindfield  
 cogito snips  
 SMARTLING Unbabel PolyAI

**SEARCH**  
 elasticsearch ORACLE ENDECA  
 algolia covéo  
 Lucidworks ATTIVO  
 swifttype EXALEAD  
 alphasense MAANA  
 omnius SINEQUA

**LOG ANALYTICS**  
 splunk  
 sumologic  
 solarwinds  
 TIMBER  
 kibana  
 logz.io

**SOCIAL ANALYTICS**  
 Hootsuite sprinklr  
 NETBASE  
 synthosio track  
 simple reach  
 bitly SimilarWeb

**WEB / MOBILE / COMMERCE ANALYTICS**  
 Google Analytics  
 mixpanel AMPITUDE  
 Airtable RESCI  
 SIGOPT granify  
 custora

## OPEN SOURCE

**FRAMEWORKS**  
 Spark  
 Flink YARN TEZ  
 MESOS  
 docker CDOP  
 Red Hat HELIX

**QUERY / DATA FLOW**  
 Spark SQL  
 presto  
 SLAMDATA  
 ARAPACHE DRILL  
 GraphQL  
 Flink

**DATA ACCESS & DATABASES**  
 cassandra mongoDB redis  
 CouchDB  
 Cockroach LABS druid  
 riak HBASE  
 CLOUD TRICKS SciDB  
 Cloud Spanner ACCUTULO

**ORCHESTRATION & MGMT**  
 talend  
 Apache Ambari  
 Apache Airflow  
 MESOS  
 etcd Kong

**STREAMING & MESSAGING**  
 Spark  
 nifi  
 Flink  
 beam  
 kafka STORM  
 Apache RocketMQ

**STAT TOOLS & LANGUAGES**  
 python  
 Scala  
 Studio  
 SciPy julia

**AI OPS & INFRA**  
 miflow  
 Kubeflow  
 mlops  
 SELDON  
 PyLemon

**AI / MACHINE LEARNING / DEEP LEARNING**  
 TensorFlow Keras  
 OpenAI DM TK  
 inxnet VELES  
 theano  
 Apache SINGA  
 DIMSUM FeatureFu  
 neom DSSTNE ml4b  
 DL4J MAHOUT Aerosolve FastAI mir

**SEARCH**  
 elasticsearch  
 Solr

**LOGGING & MONITORING**  
 kibana  
 logstash Prometheus  
 fluentbit fluentd Grafana  
 VECTOR

**VISUALIZATION**  
 matplotlib  
 TensorBoard  
 seaborn  
 Bokeh

**COLLABORATION**  
 BeakerX  
 jupyter  
 ANACONDA

**SECURITY**  
 Apache Ranger  
 KNOX  
 Sentry  
 ACCUTULO

## DATA SOURCES & APIs

HEALTH IOT FINANCIAL & ECONOMIC DATA AIR / SPACE / SEA PEOPLE / ENTITIES LOCATION INTELLIGENCE OTHER DATA SERVICES INCUBATORS & SCHOOLS RESEARCH

see full version at: [http://mattturck.com/wp-content/uploads/2019/07/2019\\_Matt\\_Turck\\_Big\\_Data\\_Landscape\\_Final\\_Fullsize.png](http://mattturck.com/wp-content/uploads/2019/07/2019_Matt_Turck_Big_Data_Landscape_Final_Fullsize.png)

## APPLICATIONS - ENTERPRISE

**SALES**  
 CHORUS  
 INSIDESALES.COM peopleai  
 conversica  
 clari aviso tact.ai  
 fuse machines Clearbit

**MARKETING - B2B**  
 RADIUS App Annie  
 EVERSTRING  
 MINTIGO  
 sense  
 tubular  
 ENGAGIO  
 KNOTCH mpe

**MARKETING - B2C**  
 Zeta  
 bloomreach SendGrid  
 braze ACTIONIQ BLUECORE  
 CONTENT SQUARE TEALUM  
 Ampero amperity  
 Simon Lytix PERSADO  
 remesh

**CUSTOMER EXPERIENCE / SERVICE**  
 qualtrics MEDALLIA SurveyMonkey UserTesting  
 CLARABRIDGE zendesk Customer freshdesk  
 INTERCOM Drift LIVEPERSON Gainsight pendo  
 HEAP Amplitude Watson Assistant  
 DigitalGenius A.S.A.P.P ada NUTOMAT ahniti  
 CaSDesk

**ENTERPRISE PRODUCTIVITY**  
 slack  
 ORACLE  
 GURU lumia  
 DIFFBOT clara  
 talla Kasisto

**HUMAN CAPITAL**  
 pyometrics  
 hiQ  
 Allyo textio  
 Wade & Wendy Stella  
 entelo  
 uncommo  
 eat beemey

**LEGAL**  
 RAVEL  
 Everlaw  
 DISCO kiro  
 JUDICATA  
 BREVIATA  
 IRONCLAD  
 PREPOSITION  
 ROSS  
 Casetext

**REGTECH & COMPLIANCE**  
 BigID  
 text IQ  
 Comply Advantage

**FINANCE**  
 fnanplan  
 ZUORO  
 SP5AHANA  
 TRADESHIFT  
 SCALE FACTOR  
 baakeeper  
 pilot

**BACK OFFICE AUTOMATION & RPA**  
 UiPath  
 blueprints  
 VIDADO  
 WorkFusion workato  
 Reworks  
 SCALE FACTOR  
 ANTWORKS  
 ALKYRI

**SECURITY**  
 TANIUM  
 CYLANCE  
 zscaler  
 StackPath  
 illumio  
 CODE42  
 CIPHERCLoud  
 DARKTRACE ANOMALI  
 threatMetrix  
 VECTRA  
 Guardian  
 DATAVISOR  
 sift science  
 pindrop  
 exabeam  
 SICNIFYD  
 SentinelOne  
 SecurityScorecard  
 SCURE  
 CodeSecure  
 bitglass  
 Recorded Future  
 feedzai  
 Cybex  
 BITSIGHT  
 sparkcognition  
 CyberArk  
 FORTER  
 riskrecon  
 J.A.S.K  
 ARCA SECURITY  
 BLUEHEXAGON  
 Semble  
 SHIELD AI  
 Armorblox

## APPLICATIONS - INDUSTRY

**ADVERTISING**  
 AppNexus  
 critico  
 ORACLE  
 MOAT  
 theTradeDesk  
 dstillery  
 LiveIntent  
 TAPAD dataxum gumgum  
 Cppier

**EDUCATION**  
 Lullishou  
 猿辅导  
 KNEWTON  
 Clever  
 Dicleara  
 kidaptive  
 PANORAMA  
 knowre  
 gradescope

**REAL ESTATE**  
 REDFIN  
 Opndoor  
 VTS  
 CREDIFI  
 GEOPHY  
 reonomy  
 COMPSTAK  
 SPACEMAKER  
 SKYLINE

**GOV'T**  
 OPENGOV  
 mark43  
 LiveStories  
 Passport  
 SmartProcure  
 STREETLIGHTDATA  
 OpenDataSoft

**INTELLIGENCE**  
 Palantir  
 Dataminr  
 Quid  
 PRIMER  
 FORGE

**FINANCE - INVESTING**  
 KENSHC  
 Quantopian  
 ADAPPER  
 NUMERA  
 ISENTIUM  
 ALGORITZ  
 FlavourPack  
 PAGAYA

**FINANCE - LENDING**  
 ondeck  
 affirm  
 Kreditech  
 AVANT  
 TALA  
 CLEARBANC  
 Upgrade  
 100Credit  
 WeLab  
 Weccash  
 TrueAccord  
 MoneyLion  
 Active AI  
 aire  
 cignifi

**INSURANCE**  
 Metromile  
 Lemonade  
 CYENCE  
 Hippo  
 Shift Technology  
 ROOT  
 zesty.ai  
 TRAVELAI  
 CAPE

**HEALTHCARE**  
 flatiron Clover KYRUS  
 HealthTap  
 METABIOTA  
 Gingerio Glow  
 babylion  
 3DMed zebra  
 PathAI  
 OVIA  
 TEMPUX  
 patientslikeme  
 AiCuro insitro  
 LIAODOO  
 citizen  
 notable  
 humandotai  
 RECURSION  
 prognos  
 Qventus  
 ARTERYS  
 IMA GEN  
 innovacore  
 PAIGE  
 DATAVANT

**LIFE SCIENCES**  
 Zebra  
 color  
 BenevolentAI  
 verily  
 WuXiNextCODE  
 Clear Labs  
 fraxnoma  
 MANOPRO  
 DNAnexus  
 Phosphorix  
 CITRINE  
 twoAR  
 Qventus  
 OWKIN

**TRANSPORTATION**  
 UBER TESLA  
 CLEARPATH  
 CRUISE  
 NUFU  
 drive ai  
 CAMBRIDGE  
 Aurora  
 nauto  
 AMOTIVE  
 G7  
 PILOT AI  
 NIO  
 OPTIMUS  
 moovit  
 Ilic  
 nexar  
 Kodiak  
 comma.ai  
 netradynae  
 German Autotech  
 Civil Maps  
 cognate  
 thinc  
 INRIX

**AGRICULTURE**  
 FARMERS  
 Granular  
 JOHN DEERE  
 BLUE RIVER  
 FarmersEdge  
 AgroStar  
 FarmLogs  
 TARANIS  
 GAMAYA  
 terraviva  
 prospera

**COMMERCE**  
 instacart  
 STITCH FIX  
 Stitch Fix  
 lowGood  
 heurix

**INDUSTRIAL**  
 AVEVA SIEMENS  
 PREDIX UPTAKE  
 SCORTEX  
 TACHYUS  
 OTHER  
 eharmony stem  
 Amper  
 ByteDance  
 Huggins  
 elect  
 SOJERN  
 BBOXEVER  
 VERDIGIS  
 duetto  
 Mendeck  
 Electric  
 ZINER  
 Spoke  
 Moveworks

Apple VALIDIC  
 practicefusion

IOT  
 GE Digital  
 IOTAKE

FINANCIAL & ECONOMIC DATA

AIR / SPACE / SEA

PEOPLE / ENTITIES

LOCATION INTELLIGENCE

OTHER

DATA SERVICES

INCUBATORS & SCHOOLS

RESEARCH  
 facebook research  
 OpenAI  
 MIRI



# DATA & AI LANDSCAPE 2019

## INFRASTRUCTURE

**HADOOP ON-PREMISE**  
cloudera Hortonworks  
MAPR Pivotal  
IBM InfoSphere  
jethro

**HADOOP IN THE CLOUD**  
aws Microsoft Azure  
Google Cloud  
SAP Cloud Platform  
IBM InfoSphere BigInsights  
arm  
Luible CAZENA

**STREAMING / IN-MEMORY**  
Amazon Kinesis  
databricks  
SAP Cloud Platform  
ORACLE  
confluent  
strim hazelcast  
GridGain  
GIGASPACE  
FASTDATA  
kx

**NO SQL DATABASES**  
aws  
ORACLE  
mongoDB MarkLogic  
Couchbase DISTRAX  
redislabs  
KEROPIKE  
MongoDB SCYLLA

**NEWSQL DATABASES**  
SAP Clustrix  
Pivotal  
MEMSQL  
infuxdata  
cockroach LABS  
VOLTDB splice  
paradigm

**GRAPH DBs**  
neo4j  
Amazon Neptune  
IBM  
ORACLE  
InfinteGraph

**MPP DBs**  
TERADATA  
IBM Data Warehouse Systems  
Qcon  
Kognitio  
Exasol  
dremio  
Yellowbrick

**CLOUD EDW**  
aws  
Microsoft Azure  
Pivotal  
snowflake  
Infoworks

**SERVERLESS**  
aws  
Microsoft Azure  
Pivotal  
snowflake  
Infoworks

**DATA TRANSFORMATION**  
talend pentaho  
alteryx TRIFACTA  
tamr Paxata  
StreamSets UNIFI

**DATA INTEGRATION**  
Microsoft Informatica  
snoplogic enigma  
Qlik Data Catalyst  
Segment ATTUNITY  
ZALONI import.io  
Infovorks Fivetran  
SNOWFLOW MATILLION

**DATA GOVERNANCE**  
Microsoft Informatica  
IBM  
Maize Skyhigh Security Cloud  
colibra dremio  
Alation  
HMDTA  
OKERA  
MANTA data.world

**MGMT / MONITORING**  
New Relic actifio  
rubrik  
APPDYNAMICS  
dynamtrace  
WAVEFRONT  
SignalFx  
drivu  
splunk  
Moogsoft pagerduty  
unavai Numerify  
SEALY VEBAM  
zabbix OpsRamp  
MAGNITUDE

**STORAGE**  
aws  
Microsoft Azure  
PURE STORAGE  
ALLUXIO wasabi  
Qumulo panache  
COHERITY

**CLUSTER SVCS**  
Amazon EMR  
IBM  
Microsoft Azure  
MESSENGER  
packet  
nimbustorage  
Quanto  
big5 Computing  
CYCLOCLOUD

**DATA GENERATION & LABELLING**  
amazon mechanicalturk  
upwork  
open  
hive  
Mighty AI  
ALEXEE  
LIONBRIDGE

**AI OPS**  
ALGORITHMIA  
comet  
verta.ai datmo  
statstion  
WhisperAI  
Determined AI  
tidder

**GPU DBs & CLOUD**  
kinetica  
SQUEAM  
bryllyt  
BLAZINGDB  
PG-Strm  
LOYDHUB

**HARDWARE**  
Google TPU arm  
intel AI NVIDIA  
GRAPHICORE MYTHIC  
Intel Power  
GRAPHICORE MYTHIC  
Movidius habana  
WAVE  
CERENAM  
PARADIGM  
DEFINIX

**CROSS-INFRASTRUCTURE/ANALYTICS**  
aws Google Cloud Microsoft IBM SAP Hewlett Packard Enterprise SAS IOIOWATA vmware TIBCO TERADATA ORACLE NetApp syncsort MAPR cloudera

## ANALYTICS & MACHINE INTELLIGENCE

**DATA ANALYST PLATFORMS**  
Microsoft pentaho alteryx  
Digital Reasoning  
GUAVUS AYASDI  
ATTIVO Datameer incorta  
interana MODE ENDOR  
sisu switchboard Starburst

**DATA SCIENCE PLATFORMS**  
IBM databricks dataiku  
DOMINO rapidminer TIBCO  
ANACONDA SAS  
KNIME MathWorks

**BI PLATFORMS**  
looker  
amazon analytics  
aws  
DOMO  
ARCADIA DATA ThoughtSpot  
ATSCALE  
Qlik  
GoodData Information Builders birst  
MicroStrategy Keon IO

**VISUALIZATION**  
tableau  
Power BI  
SAP  
Google Cloud  
celonis  
Periscope Data  
zepl  
VIZIENZE ELEMENT  
CHARTIO TOUCAN TCOO

**MACHINE LEARNING**  
Amazon SageMaker  
Google Cloud  
H2O  
DataRobot gamalon  
VISENZE ELEMENT  
deepense.ai

**COMPUTER VISION**  
Microsoft Azure  
Amazon Rekognition  
clarifai  
EVERAI  
deepomatic  
neura  
UBIQUITY ABBE  
synthesia DataClz

**HORIZONTAL AI**  
IBM Watson Cortana Face++  
sentient  
Voyager  
Affectiva  
Numenta  
PETRUM  
narologics  
CURIOUS AI  
OSARO  
BLISS VISION  
Fortress

**SPEECH & NLP**  
Google Cloud  
amazon alexa  
amazon translate  
narrative science  
semantic machines  
Mobval  
PRIMER  
SoundHound Inc  
PRIMER  
Mindfield  
cogito snips  
SMARTLING Unbabel PolyAI

**SEARCH**  
elasticsearch  
algolia covéo  
Lucidworks ATTIVO  
swiftype  
alphasense MAANA  
omni:us SINEQUA

**LOG ANALYTICS**  
splunk  
sumologic  
solarwinds  
TIMBER  
kibana  
logz.io

**SOCIAL ANALYTICS**  
Hootsuite  
sprinklr  
NETBASE  
synthesio track  
simple reach  
bitly SimilarWeb

**WEB / MOBILE / COMMERCE ANALYTICS**  
Google Analytics  
mixpanel AMPITUDE  
Airtable RESCI  
SIGOPT granify  
custora

## APPLICATIONS - ENTERPRISE

**SALES**  
CHORUS  
INSIDESALES.COM people.ai  
conversica  
clari aviso tact.ai  
fuse machines Clearbit

**MARKETING - B2B**  
RADIUS App Annie  
EVERSTRING  
Lattice  
MINTIGO  
sense  
tubular  
KNOTCH mrpe

**MARKETING - B2C**  
Zeta  
bloomreach  
SendGrid  
braze ACTIONIQ BLUECORE  
CONTENT SQUARE TEALIUM  
mparticle  
Amplero amperity  
QUANTIFIND  
Simon Lytice PERSADO  
remesh

**CUSTOMER EXPERIENCE / SERVICE**  
qualtrics MEDALLIA SurveyMonkey UserTesting  
CLARABRIDGE zendesk Customer freshdesk  
INTERCOM Drift LIVEPERSON Gainsight pendo  
HEAP Amplitude Watson Assistant  
DigitalGenius A.S.A.P.P ada AUTOMAT afiniti  
CaRDesk  
talla Kasisto

**ENTERPRISE PRODUCTIVITY**  
slack  
ORACLE  
GURU lumina  
DIFFBOT clara

**HUMAN CAPITAL**  
Hiive  
pyometrics  
hiQ  
GOSTER  
mya  
Aillyo textio  
Wade&Wendy Stella  
entelo  
uncommon  
eat beemey

**LEGAL**  
RAVEL  
Everlaw  
DISCO kiro  
JUDICATA  
BREVIA  
IRONCLAD  
PREPOSITION  
ROSS  
Casetext

**REGTECH & COMPLIANCE**  
BigID  
Tessian  
text IQ  
Comply Advantage

**FINANCE**  
Anaplan  
ZUORA  
SAP S/4HANA  
TRADESHIFT  
mineral tree  
SCALE FACTOR  
SCALE FACTOR  
baulkeeper  
pilot

**BACK OFFICE AUTOMATION & RPA**  
UiPath  
Blue Prism  
Vidado  
AppZen  
WorkFusion workato  
Reworks  
Catalytic  
ANTWORKS  
KRYON

**SECURITY**  
Tanium  
Cylance  
zscaler  
StackPath  
illumio  
CODE42  
CyberCloud  
DARKTRACE ANOMALI  
ThreatMetrix  
VECTRA  
Guardian  
DATAVISOR  
sift science  
pindrop  
exabeam  
SICNIFYD  
SentinelOne  
SecurityScorecard  
SCURE  
Vade Secure bitglass  
BlueTalon Recorded Future  
freedat Cybe  
BIT SIGHT  
sparkcognition  
CyberArk  
FIRTER  
riskrecon  
J.A.S.K  
ARCA SECURITY  
BLUEHEXAGON  
Semble  
OSISSIM  
XANONIS  
SHIELD AI  
ArmedBlock

## APPLICATIONS - INDUSTRY

**ADVERTISING**  
AppNexus  
critico  
ORACLE  
MOAT  
theTradeDesk  
dstillery  
LiveIntent  
TAPAD dataxum gumgum  
Cupier  
yieldmo  
mobilemax

**EDUCATION**  
Lullishou  
Knewton  
KNEWTON  
Clever  
Cleclara  
kidaptive  
PANORAMA  
knowre  
gradescope

**REAL ESTATE**  
REDFIN  
Opendoor  
VTS  
CREDIFI  
GEOPHY  
reonomy  
COMPSTAK  
SPACE MAKER  
SKYLINE

**GOV'T**  
OPENGOV  
mark43  
FXI  
LiveStories  
Passport  
SmartProcure  
STREETLIGHT DATA  
OpenDataSoft

**INTELLIGENCE**  
Palantir  
Dataminr  
Quid  
PRIMER  
FORGE

**FINANCE - INVESTING**  
KENSHC  
Quantopian  
ADAPPER  
NUMERA  
SENTIUM  
ALGORIZ  
FlavourPack  
PAGAYA

**FINANCE - LENDING**  
ondeck  
affirm  
拍拍贷  
JIANPU.AI  
Kreditech AVANT  
TALA  
aurea finance  
Upstart  
100Credit  
WeLab  
WeCASH  
aire  
cignif

**INSURANCE**  
Metromile  
Lemonade  
CYENCE  
Hippo  
Shift Technology  
ROOT  
zesty.ai  
CAPE

**HEALTHCARE**  
flatiron Clover KYRUS  
HealthTap  
METABIOTA  
Gingerio Glow  
babyon  
3D Med zebra  
PathAI  
OVIA  
TEMPUS  
patientslikeme  
AiCure insitro  
LIMCOO  
citizen  
notable  
sumodivert  
RECURSION  
prognos  
enlitic  
img  
BlackThorn  
BAYLABS  
Qventus  
ARTERYS  
IMAGEN  
innovacore  
PAIGE  
DATAVANT

**LIFE SCIENCES**  
Zelus  
color  
BenevolentAI  
verily  
WuXiNextCODE  
SIBIRIN  
Clear Labs  
fronostme  
MANOPHORE  
DNAnexus  
Phosphorix  
CITRINE  
twoAR  
Atomize  
QWIKIN

**TRANSPORTATION**  
UBER TESLA  
CLEARPATH  
CRUISE  
NUFO  
drive ai  
CAMBRIDGE  
Aurora  
nauto  
AMOTIVE  
G7  
PILOT.AI  
NIO  
OPTIMUS  
moovit  
Ilic  
nexar  
Kodiak  
comma.ai  
netradynne  
Civil Maps  
cogitate  
thinc  
INRIX

**AGRICULTURE**  
FARMERS  
Granular  
JOHN DEERE  
BLUERIVER  
FarmersEdge  
AgroStar  
FarmLogs  
TARANIS  
GAMAYA  
Terraviva  
prospera

**COMMERCE**  
instacart  
FAIRFIRE  
STITCH FIX  
Doo & Co  
lowGood  
heuristic  
OTHER  
eharmony stem  
Amper  
ByteDance  
happn  
elect  
SOJERN  
BIXEVER  
VERDIGIR  
duetto  
Miledeck  
Electric  
ZINER  
Moveworks

**INDUSTRIAL**  
AVEVA SIEMENS  
PREDIX UPTAKE  
SCORTEX  
KONIG  
TACHYUS  
Apex

## OPEN SOURCE

**FRAMEWORKS**  
Spark  
Flink  
YARN  
TEZ  
MESOS  
docker  
CDAP  
RedHat  
HELDX

**QUERY / DATA FLOW**  
Spark SQL  
presto  
SLM

**DATA ACCESS & DATABASES**  
cassandra  
mongoDB  
redis  
cockroachdb

**ORCHESTRATION & MGMT**  
talend  
Apache Zookeeper

**STREAMING & MESSAGING**  
Spark  
nifi  
kafka

**STAT TOOLS & LANGUAGES**  
python  
Scala

**AI OPS & INFRA**  
miflow

**AI / MACHINE LEARNING / DEEP LEARNING**  
TensorFlow  
Keras  
PyTorch  
theano  
Caffe  
Microsoft Cognitive Toolkit  
DM  
MICHELANGLO  
ONNX  
WEXA  
LUDWIG  
PyTorch  
mlflow  
DL4J  
MAHOUT  
Aerosolve  
FastAI  
mir

**SEARCH**  
elasticsearch  
Solr

**LOGGING & MONITORING**  
elasticsearch  
kibana  
sentry  
logstash  
Prometheus  
fluentbit  
fluentd  
Grafana  
VECTOR

**VISUALIZATION**  
matplotlib  
TensorBoard  
seaborn  
Bokeh

**COLLABORATION**  
BeakerX  
jupyter  
Anaconda

**SECURITY**  
Apache Ranger  
KNOX  
Sentry  
ANACONDA  
accruiolo

The Backbone is Database Systems, Storage Engines, & Frameworks for Parallelization

increase throughput by parallelization

“scale-up”

use more powerful machines (>#CPUs, >RAM)

“scale-out”


use more machines

# Scale Up Execution

how to use more cores (threads)?

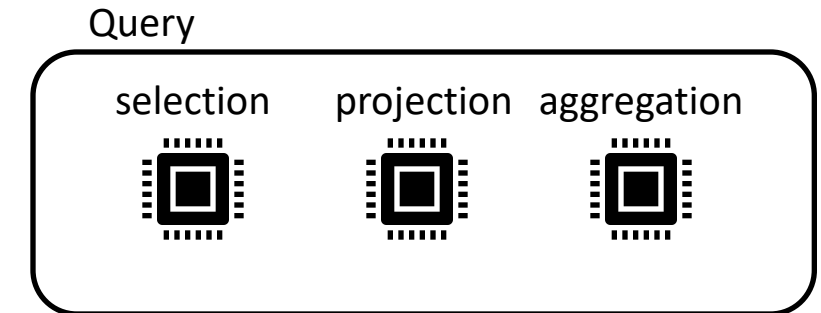
## **inter-query parallelism**

each query runs on one processor




## **inter-operator parallelism**

each query runs on multiple processors  
an operator runs on one processor



## **intra-operator parallelism**

an operator runs on multiple processors





# Scale Up Storage

needs more disks!

how to distribute data?

**block partition**

**hash partition**

**range partition**

how to distribute data accesses?

# Scale Out

similar questions across machines

new bottlenecks?

move data across machines: network!



diving into the internals of modern data systems

*cutting-edge* designs / *research* projects / *engineering* projects

# CS 561: Data Systems Architectures

Spring 2022

# Open Discussion

Questions?

for NoSQL

for DBMS

for next semester (CS561!)

for life after college (Academia vs. Industry vs. ?)

Next: Review and questions for final