

A Tale of Two Stores

DIMES+STORE

The official unofficial Nickelback Fan-Store



NICKELBACK

FEED THE MACHINE TOUR



The Beginning of DimeStore

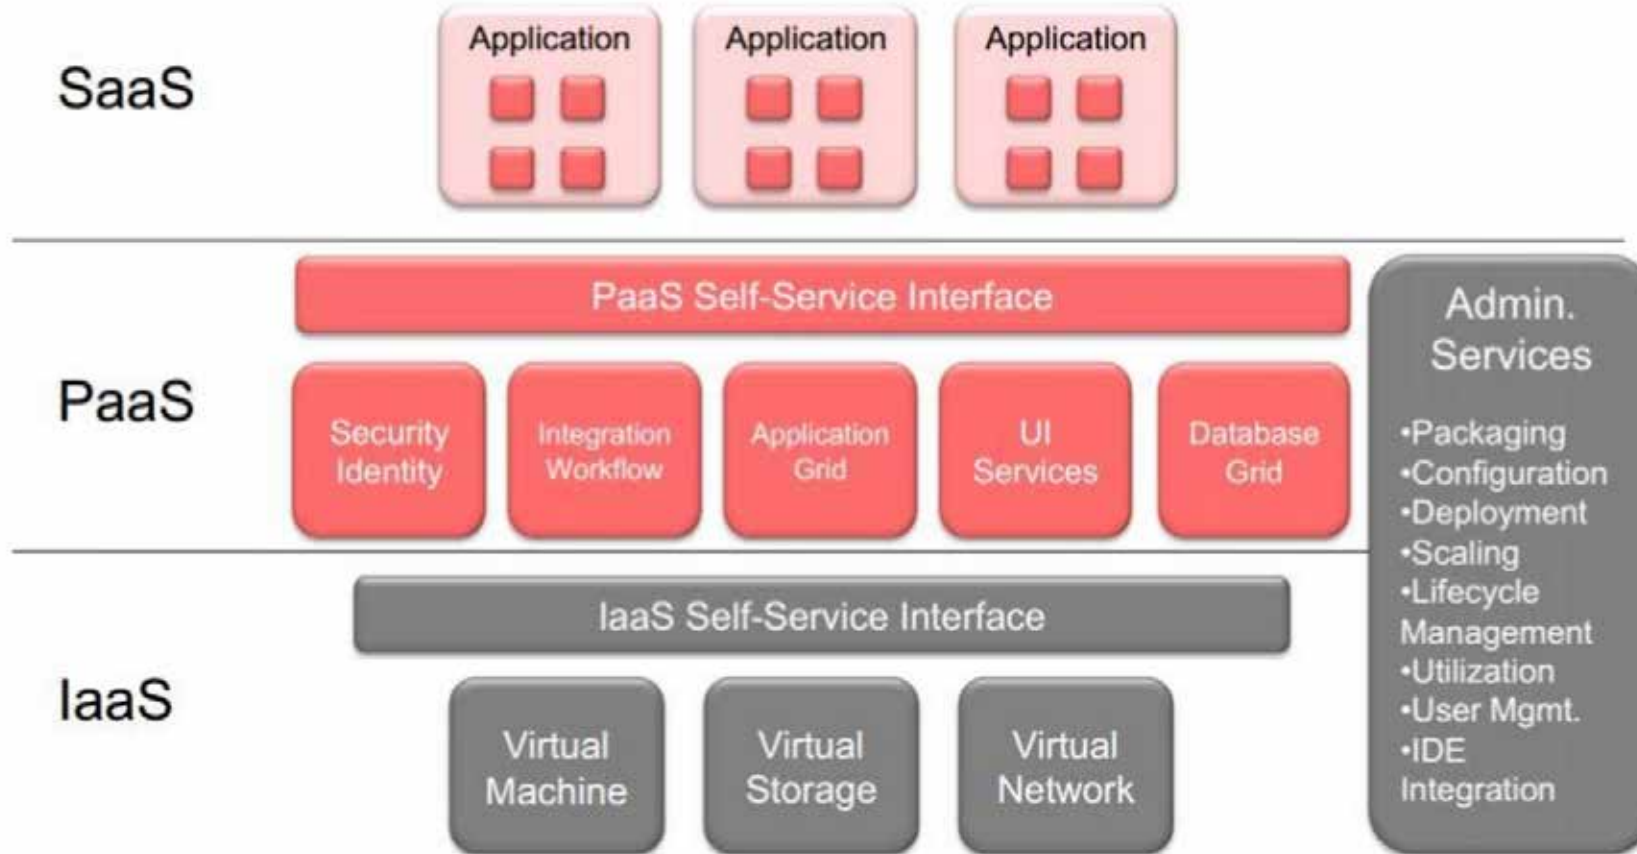




Chad Kroeger



Cloud Computing



Static Web-page with e-mail order form



1. Sign in to the AWS Management Console and open the **Amazon S3** console at <https://console.aws.amazon.com/s3/>.
2. Create two buckets that match your domain name and subdomain (*dimestore.com* and *www.dimestore.com*).

Use the Amazon S3 console to configure the bucket for website hosting

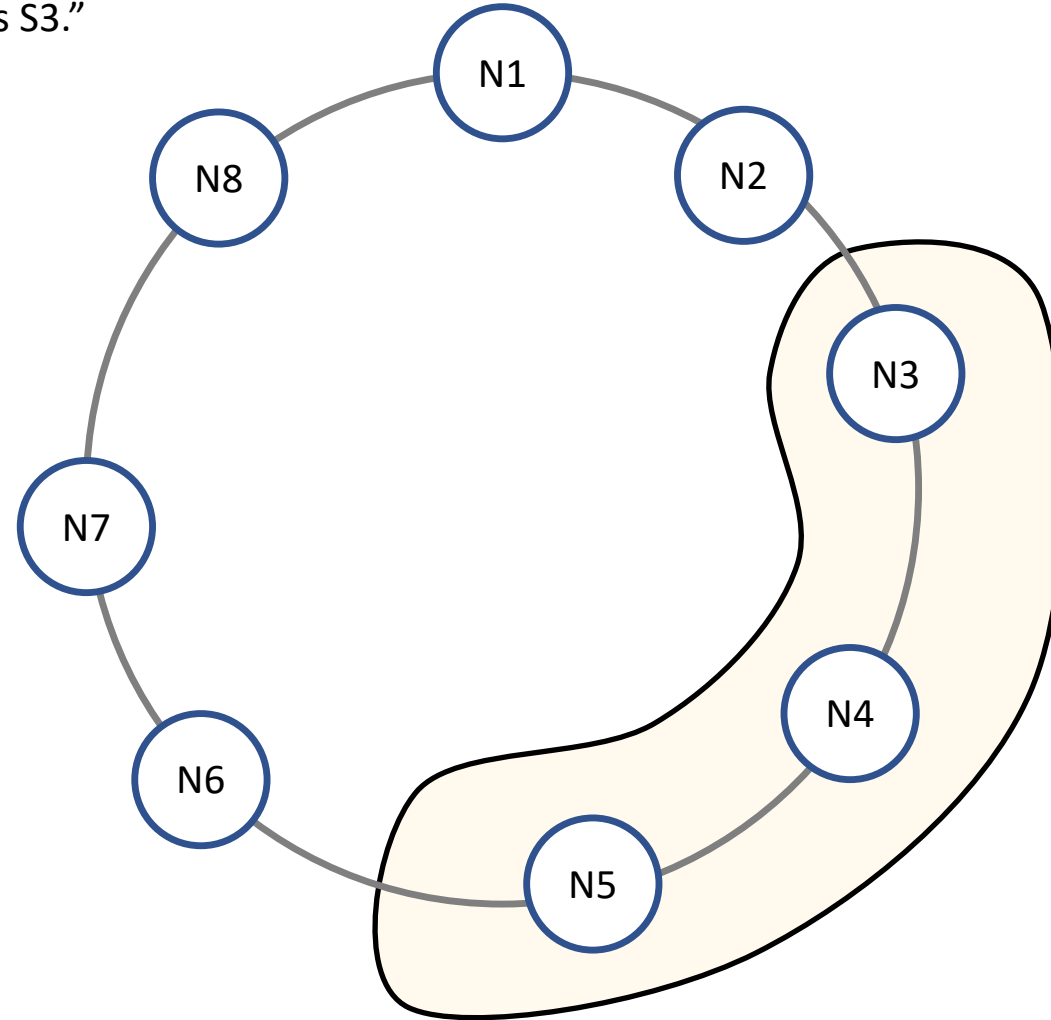
1. In the S3 buckets list, choose the bucket with the same name as your domain. Then Choose Properties -> Static website hosting.-> Use this bucket to host a website & Index Document box, enter the name of your index page (index.html)
2. Redirect requests from *www.dimestore.com* to *dimestore.com*: Choose Properties -> Choose Static website hosting -> Choose Redirect requests. In the Target bucket or domain box, enter your domain (for example, *example.com*).
3. Upload Index and Website Content
4. Edit Block Public Access Settings: choose Edit public access settings. Clear Block all public access, and choose Save.
5. Get Your Endpoints and Test Your Domain Endpoint
6. Add Alias Records for *dimestore.com* and *www.dimestore.com*

Amazon S3

Werner Vogel: “Let me emphasize the internal technology part before it gets misunderstood: Dynamo is not **directly** exposed externally as a web service; however, Dynamo and similar Amazon technologies **are** used to power parts of our Amazon Web Services, such as S3.”



Eventual Consistency



TakeYourNickelBack.com



Meanwhile...

1. I hate the establishment
2. I only listen to local bands founded by my former roommates and despise Nickelback and their irritating fans.



JAKALS



- After one too many arguments about how all the music that everyone else likes is terrible and I'm the only one with impeccable musical taste. I decide to start my own Anti-Nickelback website for making fun of these people.
- I distrust all large businesses (Except hardware manufacturers for some reason) and decide to host the website on a web server running in my basement



NGINX

DIMES+STORE

The official unofficial Nickelback Fan-Store



NICKELBACK

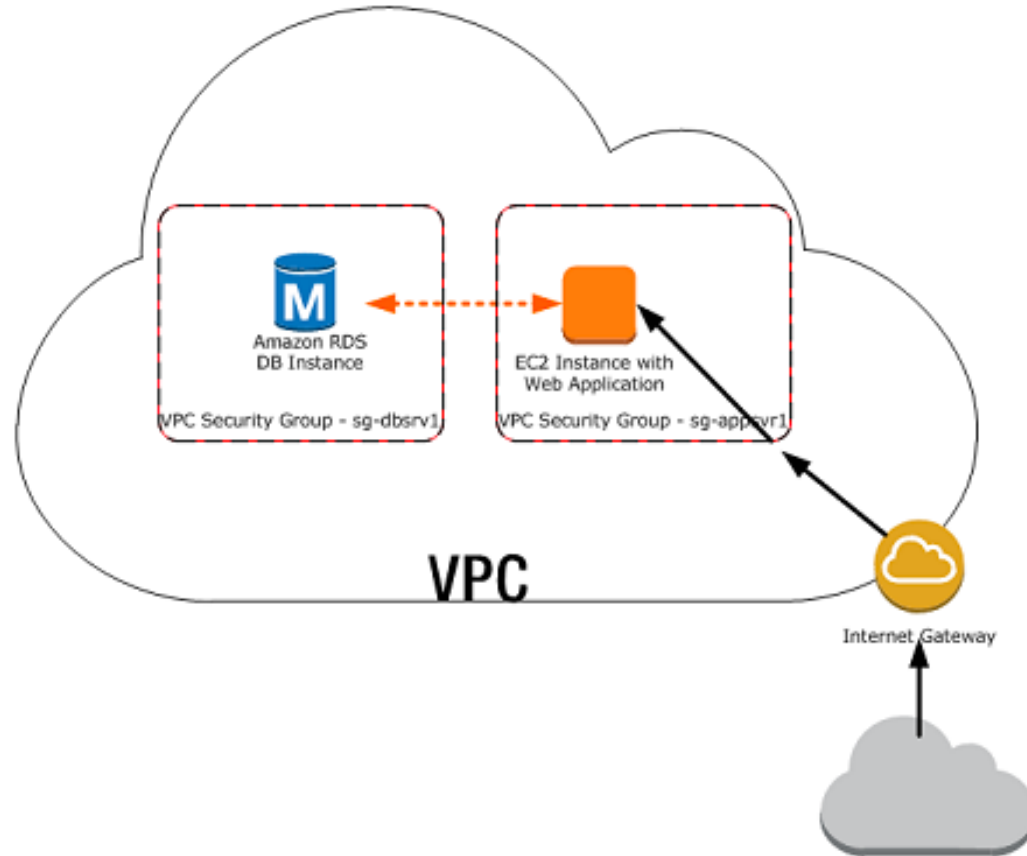
FEED THE MACHINE TOUR



DimeStore is getting traction



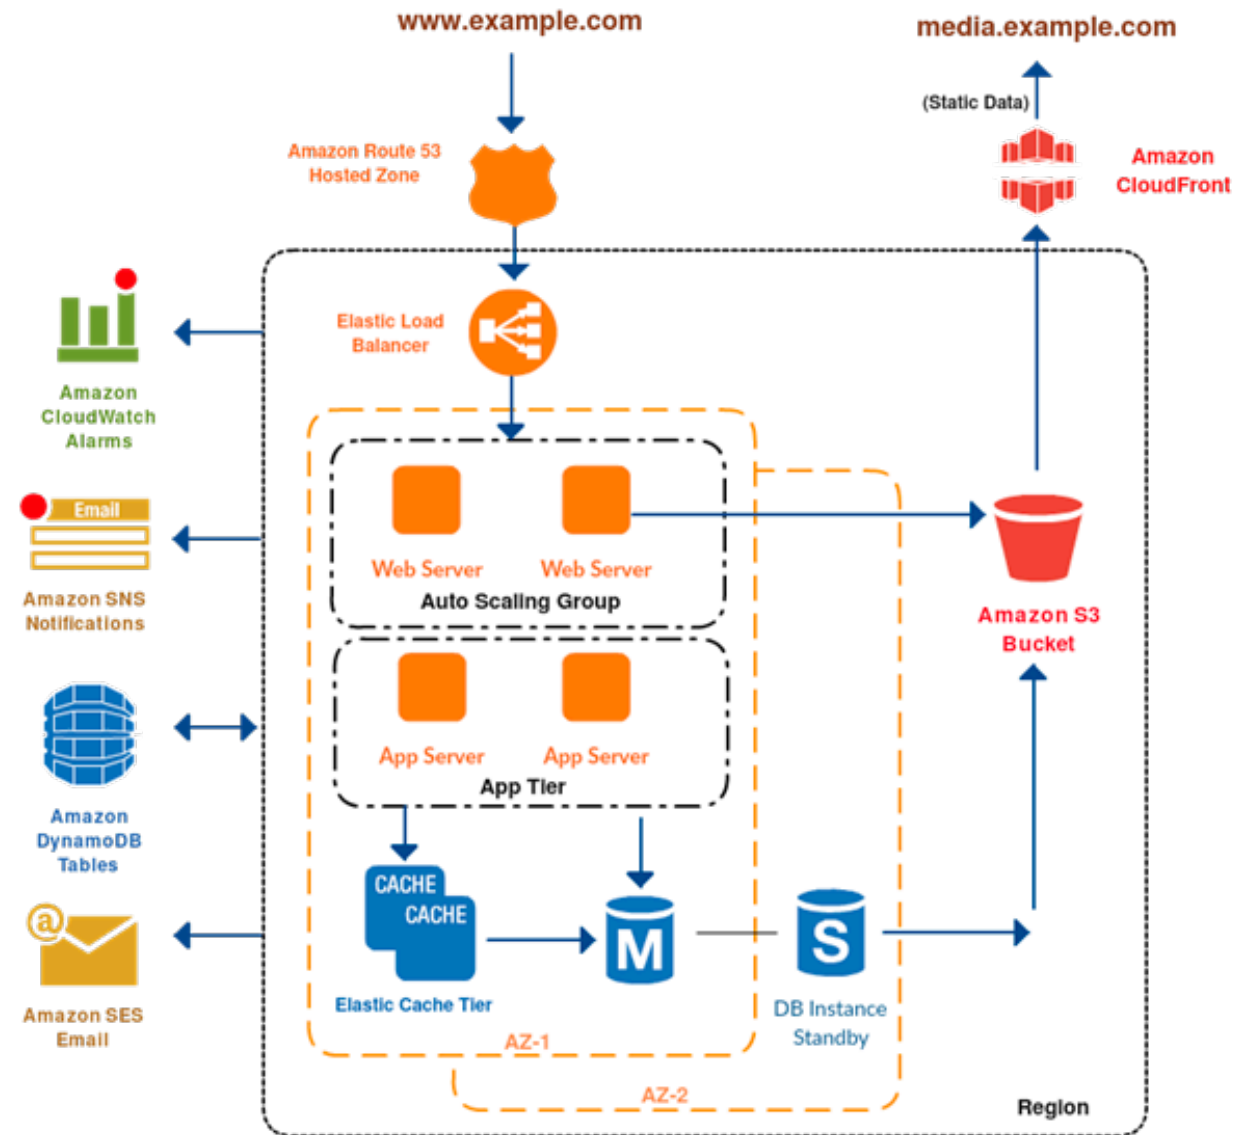
DimeStores first real web-store



Amazon RDS database engines



Typical scalable architecture on Amazon



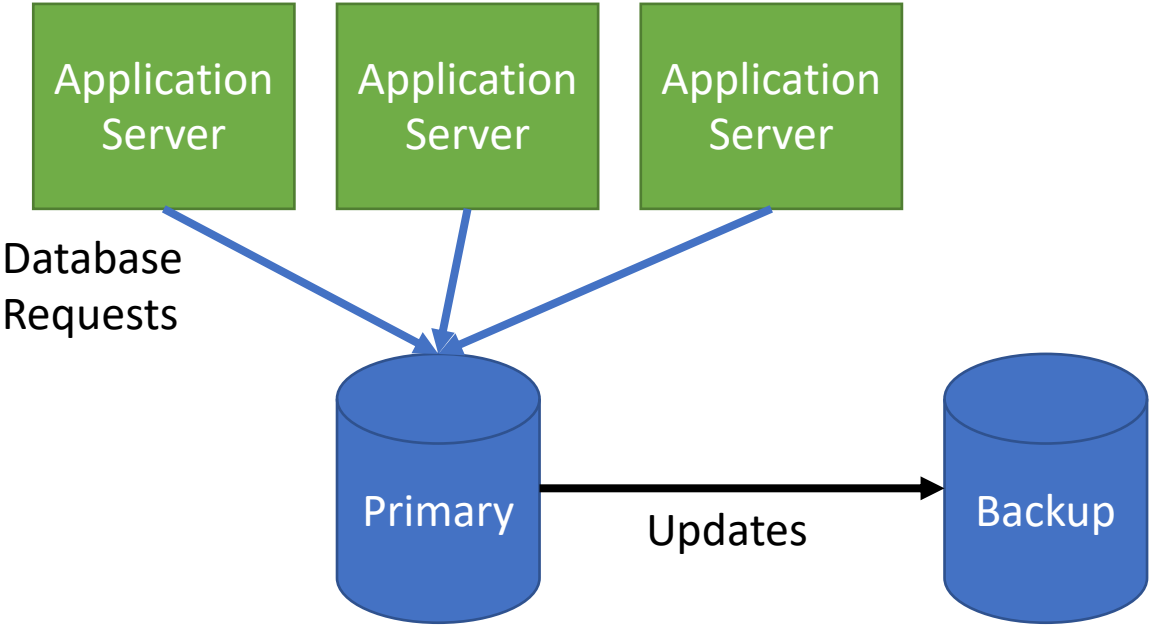
Why do we need a database and transactions?



What do you do when your database fails?



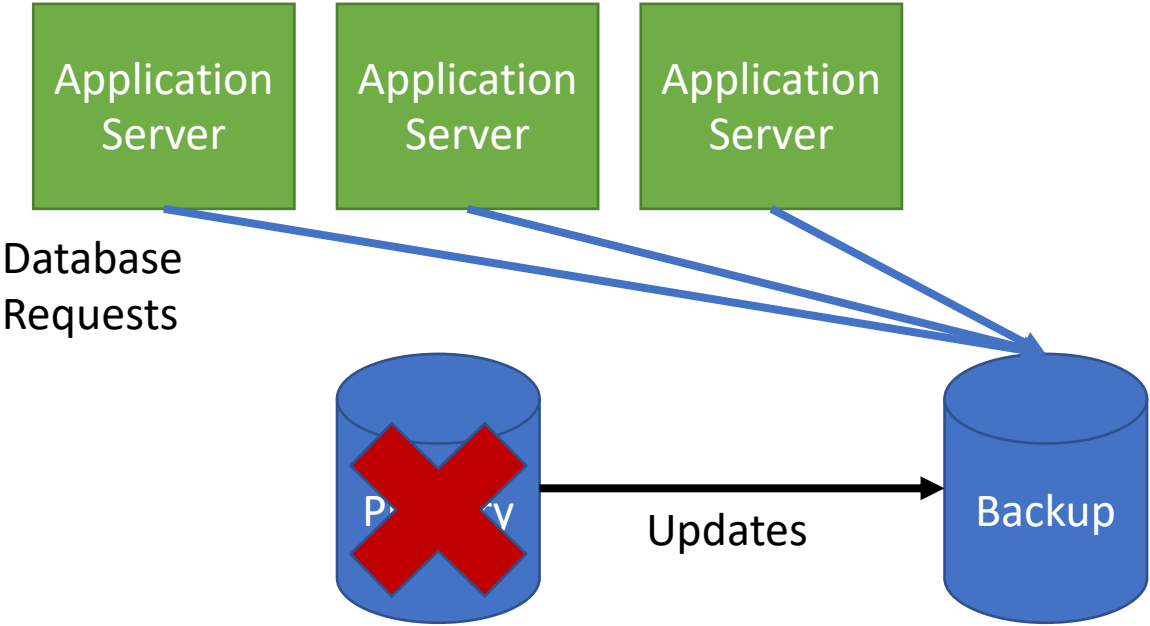
Primary Backup



What do you do when your database fails?



Primary Backup

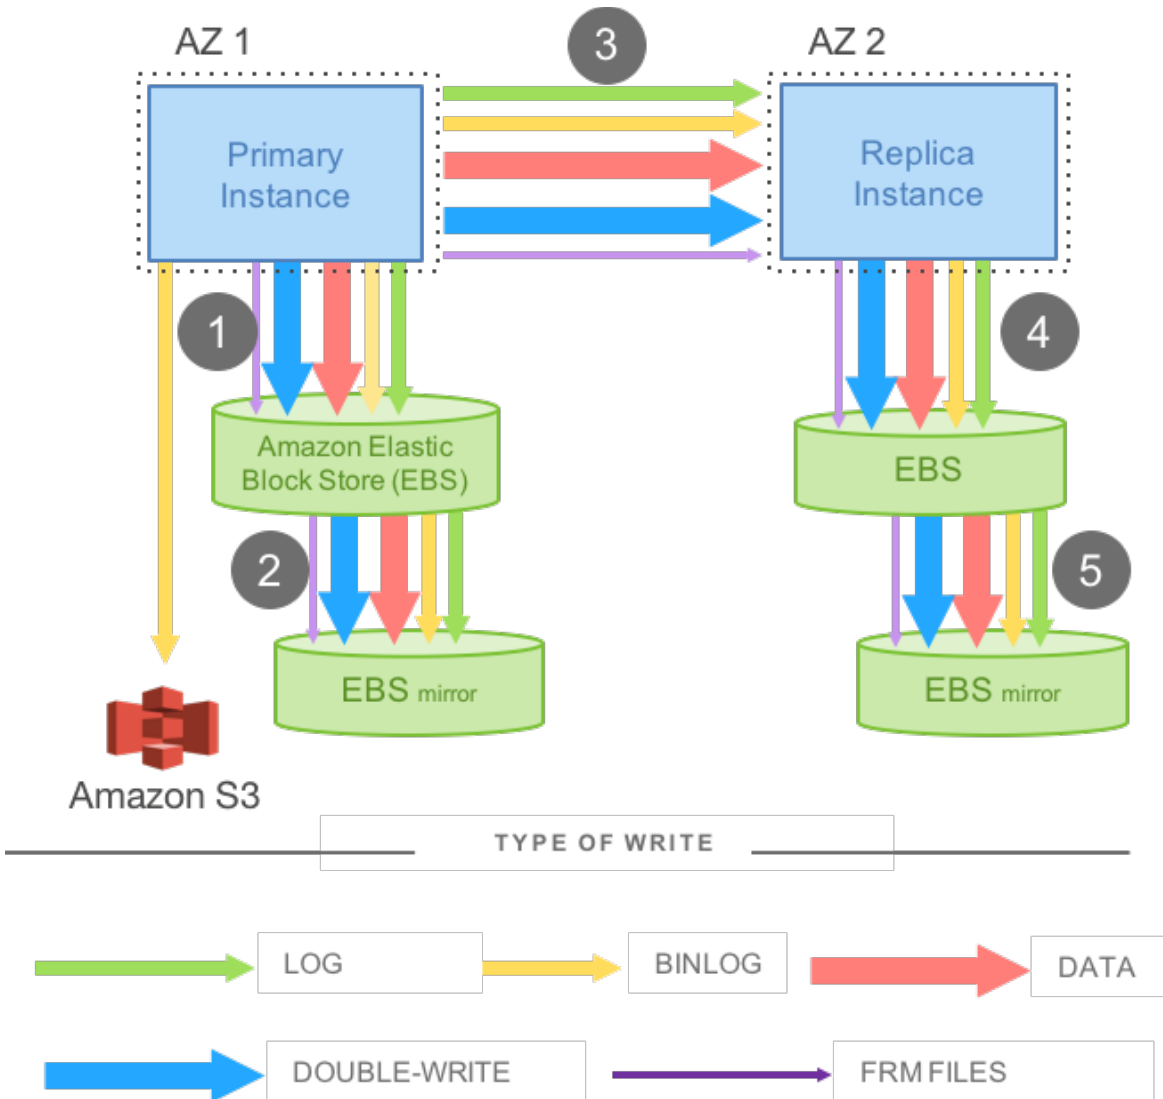




Amazon
RDS



Amazon RDS - Traditional Replication



Steps 1, 3, and 5 are sequential and synchronous.

Vertical Scaling

Modify DB Instance: Im1ks7xpixmry6w

Instance Specifications

DB Engine Version	MySQL 5.6.27 (default)
DB Instance Class	db.t2.small — 1 vCPU, 2 GiB RAM
Multi-AZ Deployment	db.t2.micro — 1 vCPU, 1 GiB RAM db.t2.small — 1 vCPU, 2 GiB RAM db.t2.medium — 2 vCPU, 4 GiB RAM
Storage Type	db.t2.large — 2 vCPU, 8 GiB RAM db.m4.large — 2 vCPU, 8 GiB RAM
Allocated Storage*	db.m4.xlarge — 4 vCPU, 16 GiB RAM db.m4.2xlarge — 8 vCPU, 32 GiB RAM db.m4.4xlarge — 16 vCPU, 64 GiB RAM db.m4.10xlarge — 40 vCPU, 160 GiB RAM db.m3.medium — 1 vCPU, 3.75 GiB RAM db.m3.large — 2 vCPU, 7.5 GiB RAM db.m3.xlarge — 4 vCPU, 15 GiB RAM db.m3.2xlarge — 8 vCPU, 30 GiB RAM
DB Instance Identifier	db.r3.large — 2 vCPU, 15 GiB RAM db.r3.xlarge — 4 vCPU, 30.5 GiB RAM db.r3.2xlarge — 8 vCPU, 61 GiB RAM db.r3.4xlarge — 16 vCPU, 122 GiB RAM db.r3.8xlarge — 32 vCPU, 244 GiB RAM
New Master Password	db.m2.xlarge — 2 vCPU, 17.1 GiB RAM db.m2.2xlarge — 4 vCPU, 34 GiB RAM

Settings

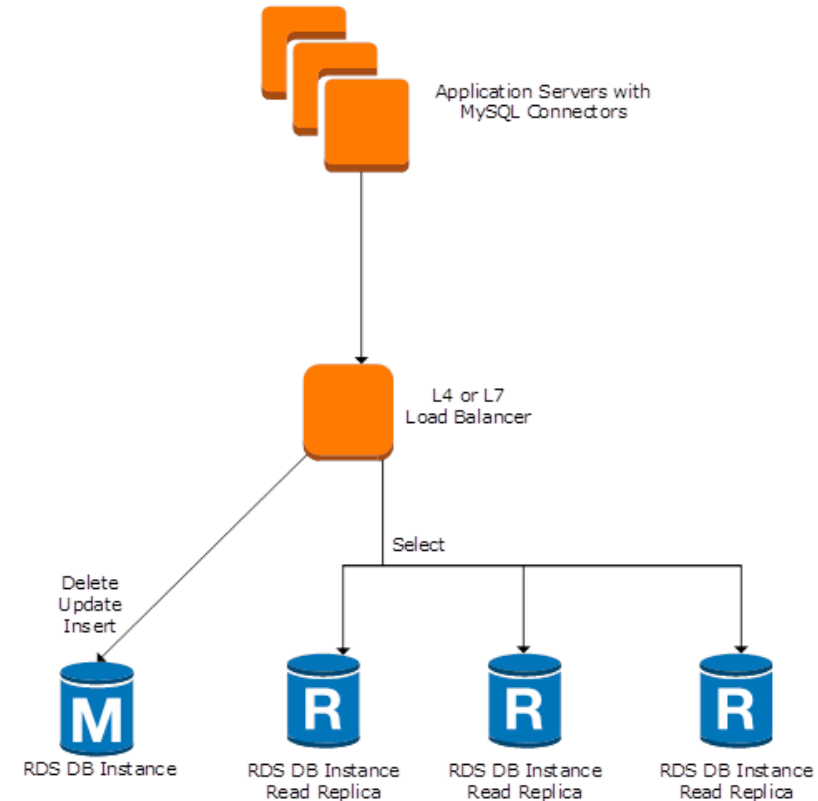
DB Instance Identifier
New Master Password

Network & Security

Security Group

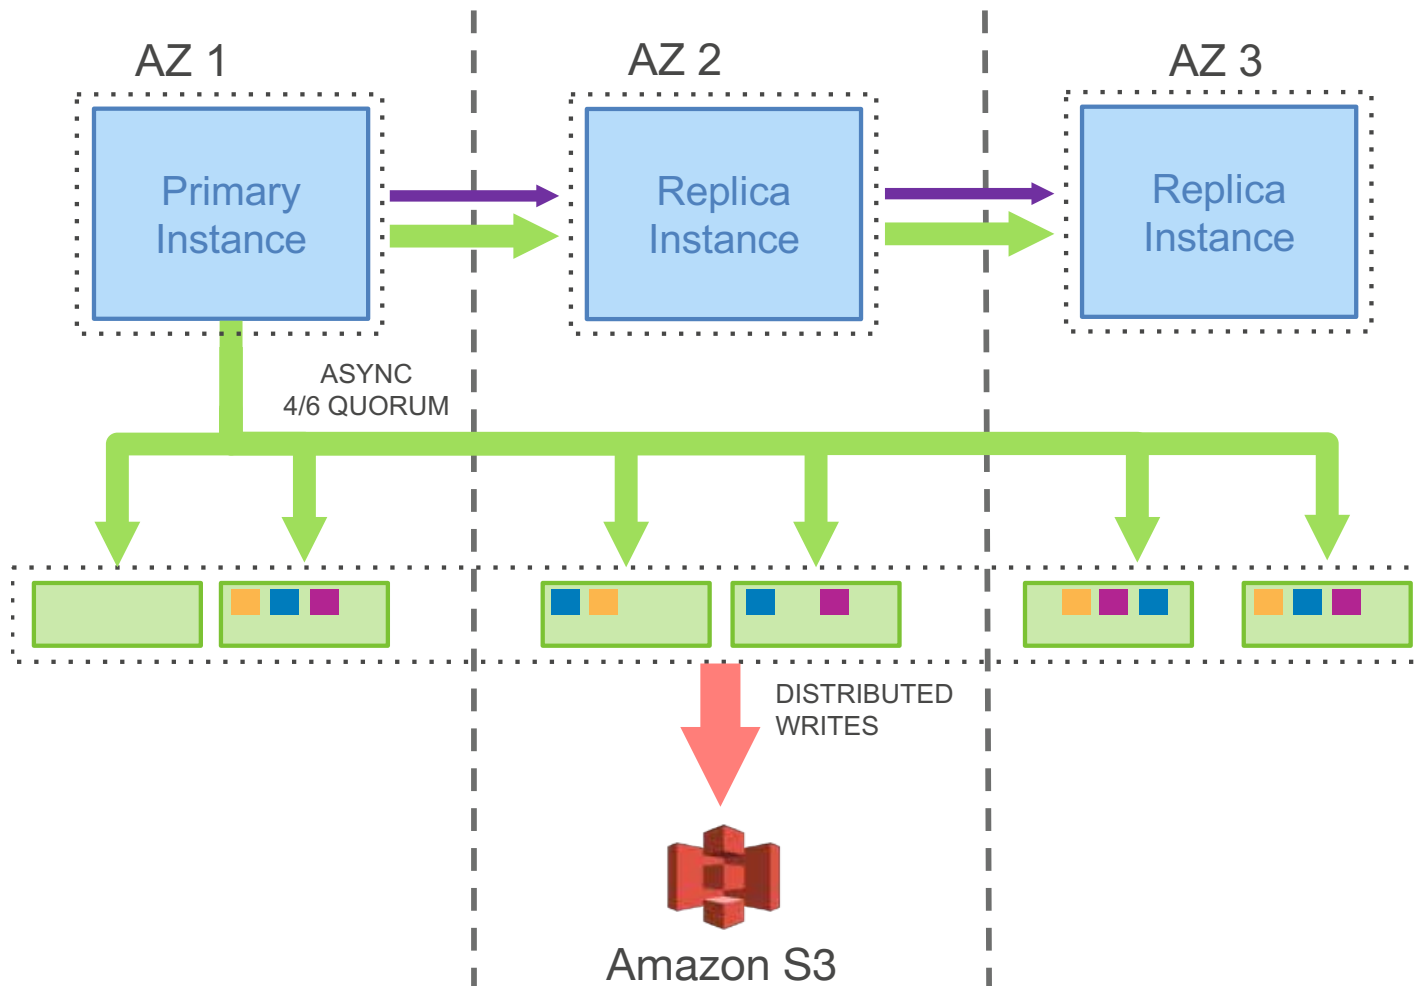
Creates potential downtime

Horizontal Scaling



Consistent reads

Amazon Aurora



Clicker:
Why does Amazon do 4/6 replication?

The log stream generated by the writer and sent to the storage nodes is also sent to all read replicas. Each replica typically lags behind the writer by a short interval (20 ms or less)

Based on MySQL. Changed InnoDB version

TAKEYOURNICKELBACK.COM



Everyone thinks they have impeccable musical taste and emails wanting to buy anti-nickelback merch.



TAKEYOURNICKELBACK.COM



- Has a forum for people complaining about Nickelback
- Has a store for great Anti-Nickelback merch
- Allows users to design their own merch and the top voted designs will be sold on the website.
- Has a blog where I talk about how much better my music tastes are than my friends. Allows users to like and comment.
- Maybe more in the future.


- Maybe I could build on top of an existing embedded database key value store like RocksDB or SQLite

Choice of DBMS



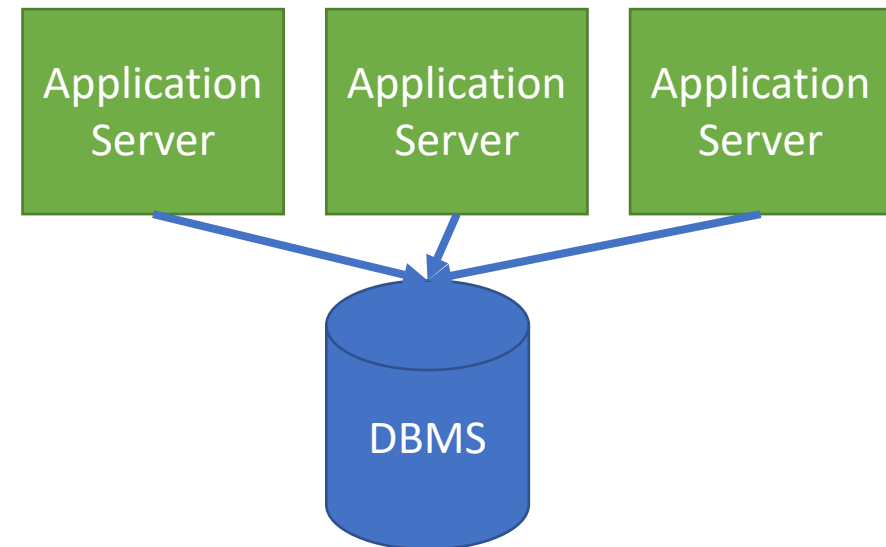
Embedded DBMS

- If you application server fails, you lose data
- Do not easily scale to application workload
- Usually runs in the same process




Independent DMBS

- Scale Application and DB independently
- Designed for many concurrent users



RocksDB



- You get a key-value database (like a hash table) that stores data on disk, supports transactions and is optimized for write-heavy workloads, and can support range scans.
- Is used as a storage engine for other systems (MySQL)
- Disadvantages:
 1. Lacks a full query language, and many common database features. If you need foreign keys, columns, indexes etc you have to build it yourself.
 2. Single User
 3. Doesn't scale beyond a single process (what do we do when our website grows to more users?)

SQLite



- An embedded SQL database that supports transactions. A row-store optimized for transactions.
- Disadvantages: Single user and not built for not built for large applications.

What are my options for data management?



- A full featured transaction processing SQL system like Postgres, MySQL (MariaDB), Oracle or Microsoft SQL Server (but not open source). Optimized for many concurrent users, and offers failover features.
- Used by the largest companies for their most important operational data. If tuned well, you can trust them to handle your data properly.
- Since we are running these ourselves, we need to hire a database administrator to ensure things are running smoothly and safely.



What does it mean to host your own database on-premise or on EC2



- Increasing reliability requires buying new machines
- Scale-out/Scale-up requires buying new machines
- They have to be configured correctly
 - Independent drives
 - Backups
 - ...
- Disaster recovery needs to be automated (how to recover from power-outage)
- Load-balancing etc. needs to be implemented
- Hire administrators
-

What about NoSQL?



- Stores JSON Documents in key-value pairs
- JSON only API
- Distributed across multiple nodes
- Has indexes for fast access
- Tend to store data in a denormalized way (instead of customer, orders table, each customer document stores all their orders)
- Started with no multiple transactions, no joins. Have since added this functionality.
- Have to give up SQL functionality, query optimization and other common DB features.
- ~~Is Web Scale~~

The rapid growth of Nickelback fans causes a backlash, and my store is growing!

- I bought a ton of new machines and put them in my basement!
- My ISP is starting to wonder what I'm doing and my hardware costs are through the roof.



DIMES+STORE

The official unofficial Nickelback Fan-Store



NICKELBACK


FEED THE MACHINE TOUR






**Sales are dropping and we need to find new revenue.
Data Science to the rescue**

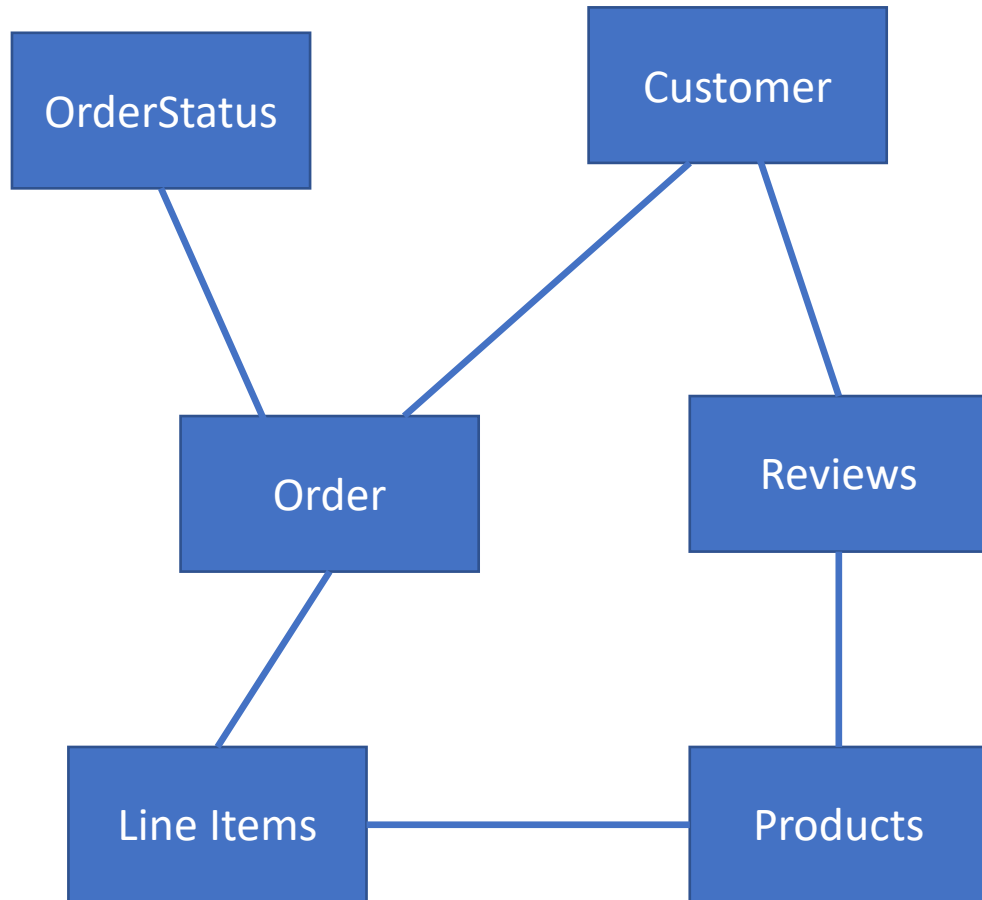
Data Science to the rescue



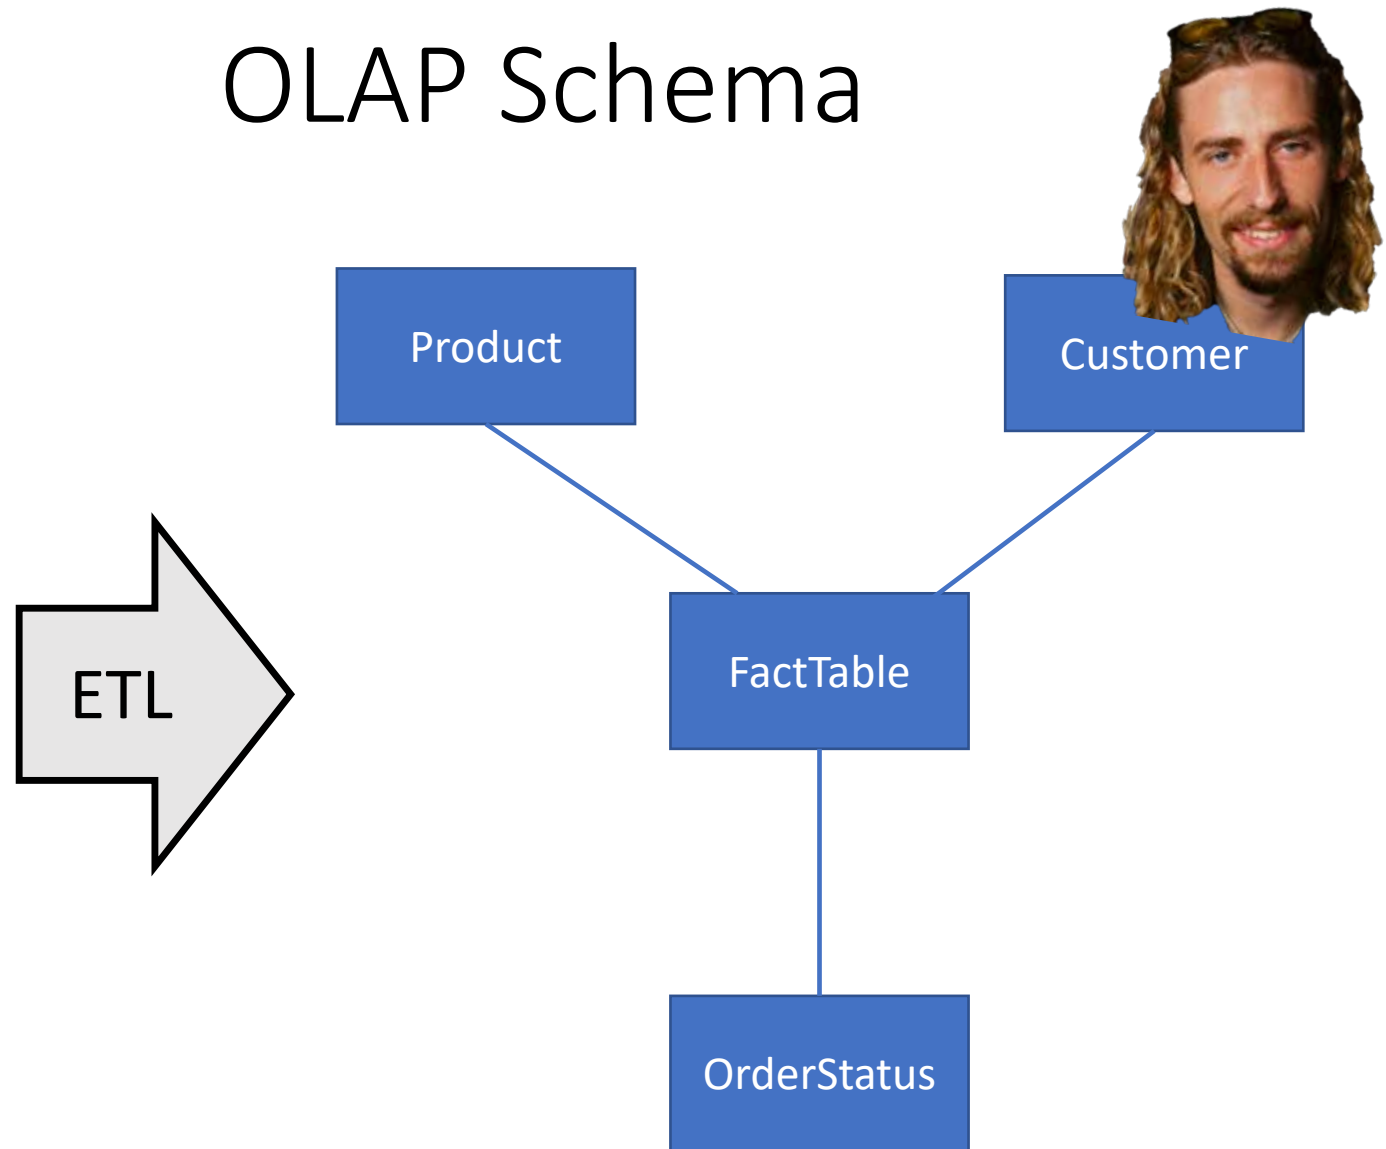
Data Science to the rescue



OLTP Schema

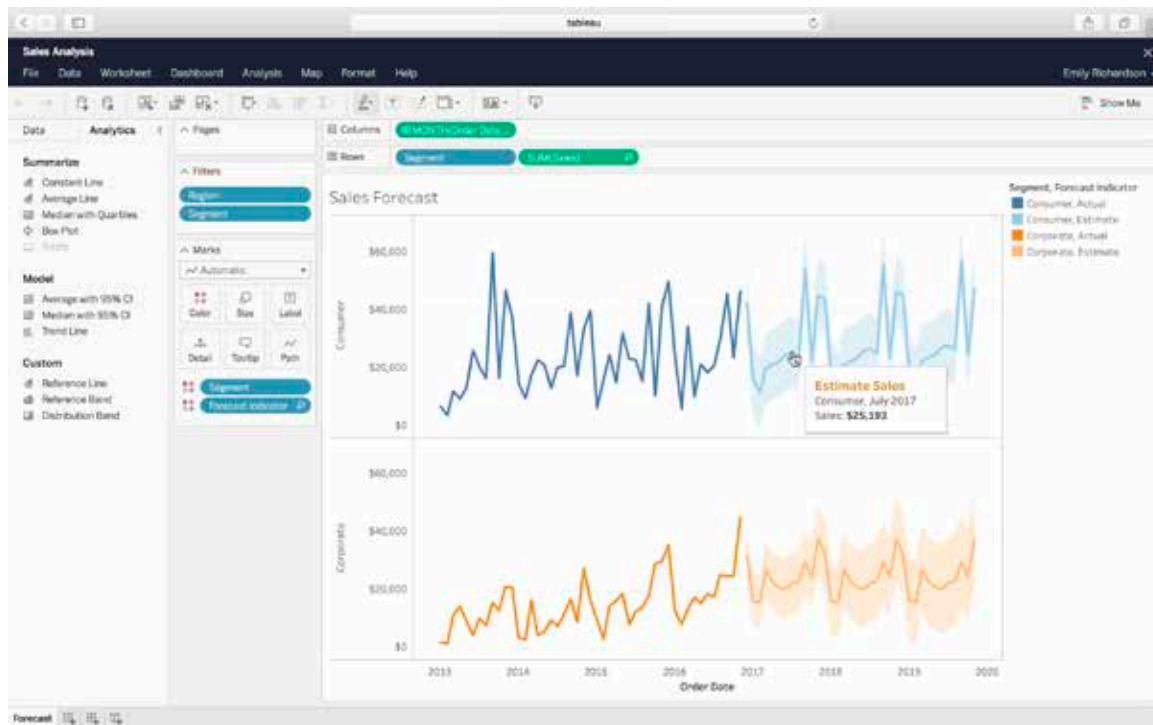


OLAP Schema

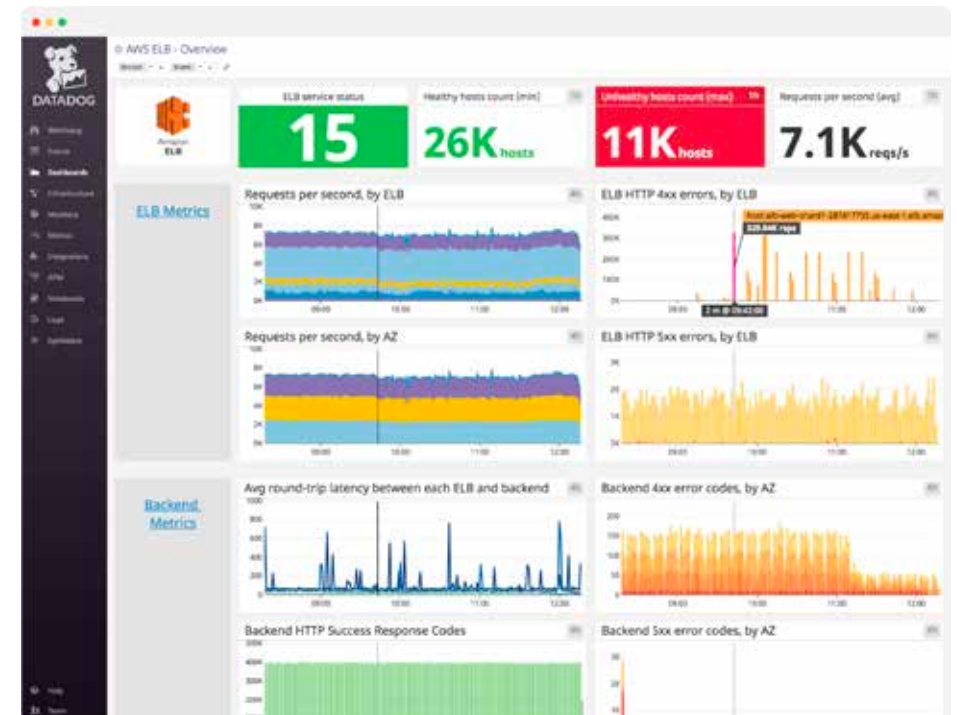


AWS Glue

What are the most common use cases for a datawarehouse?



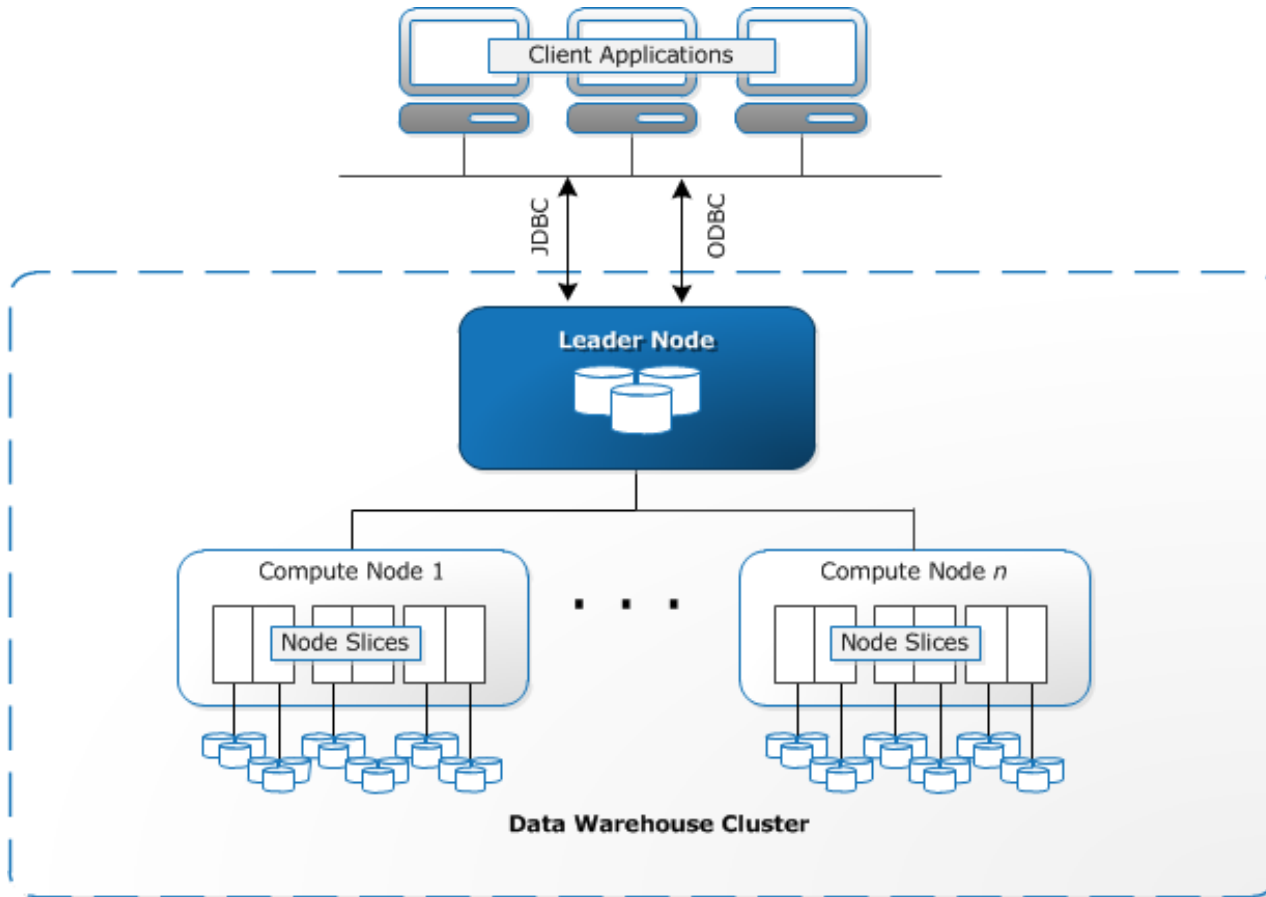
Data Exploration



Dashboards

(AWS QuickSight)

Amazon Redshift




The leader node accepts connections from client programs, parses requests, generates & compiles query plans for execution on the compute nodes, performs final aggregation of results when required

Amazon Redshift



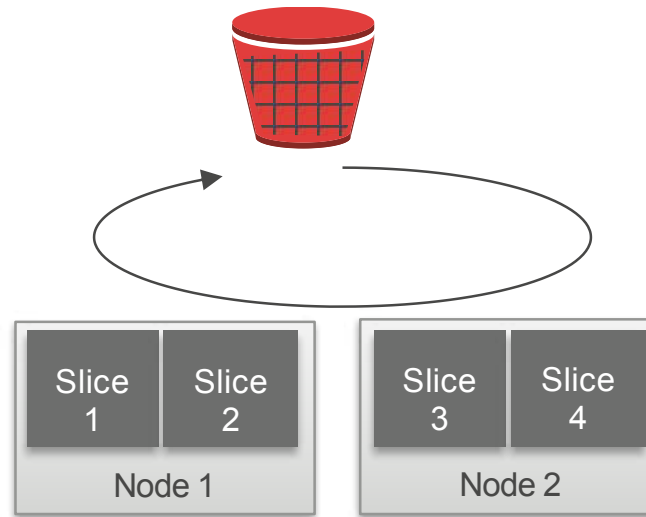
Key

Same key to same location




Even

Round robin distribution



All

All data on every node



The user can specify whether data is distributed in a round robin fashion, hashed according to a distribution key, or duplicated on all slices

Amazon Redshift Spectrum



Life of a query


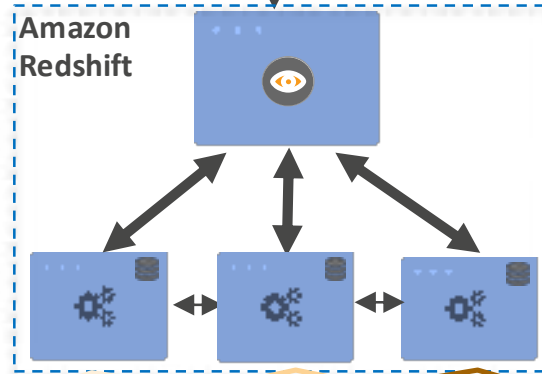


1



```
Query  
SELECT COUNT(*)  
FROM S3.EXT_TABLE  
GROUP BY...
```



JDBC/ODBC

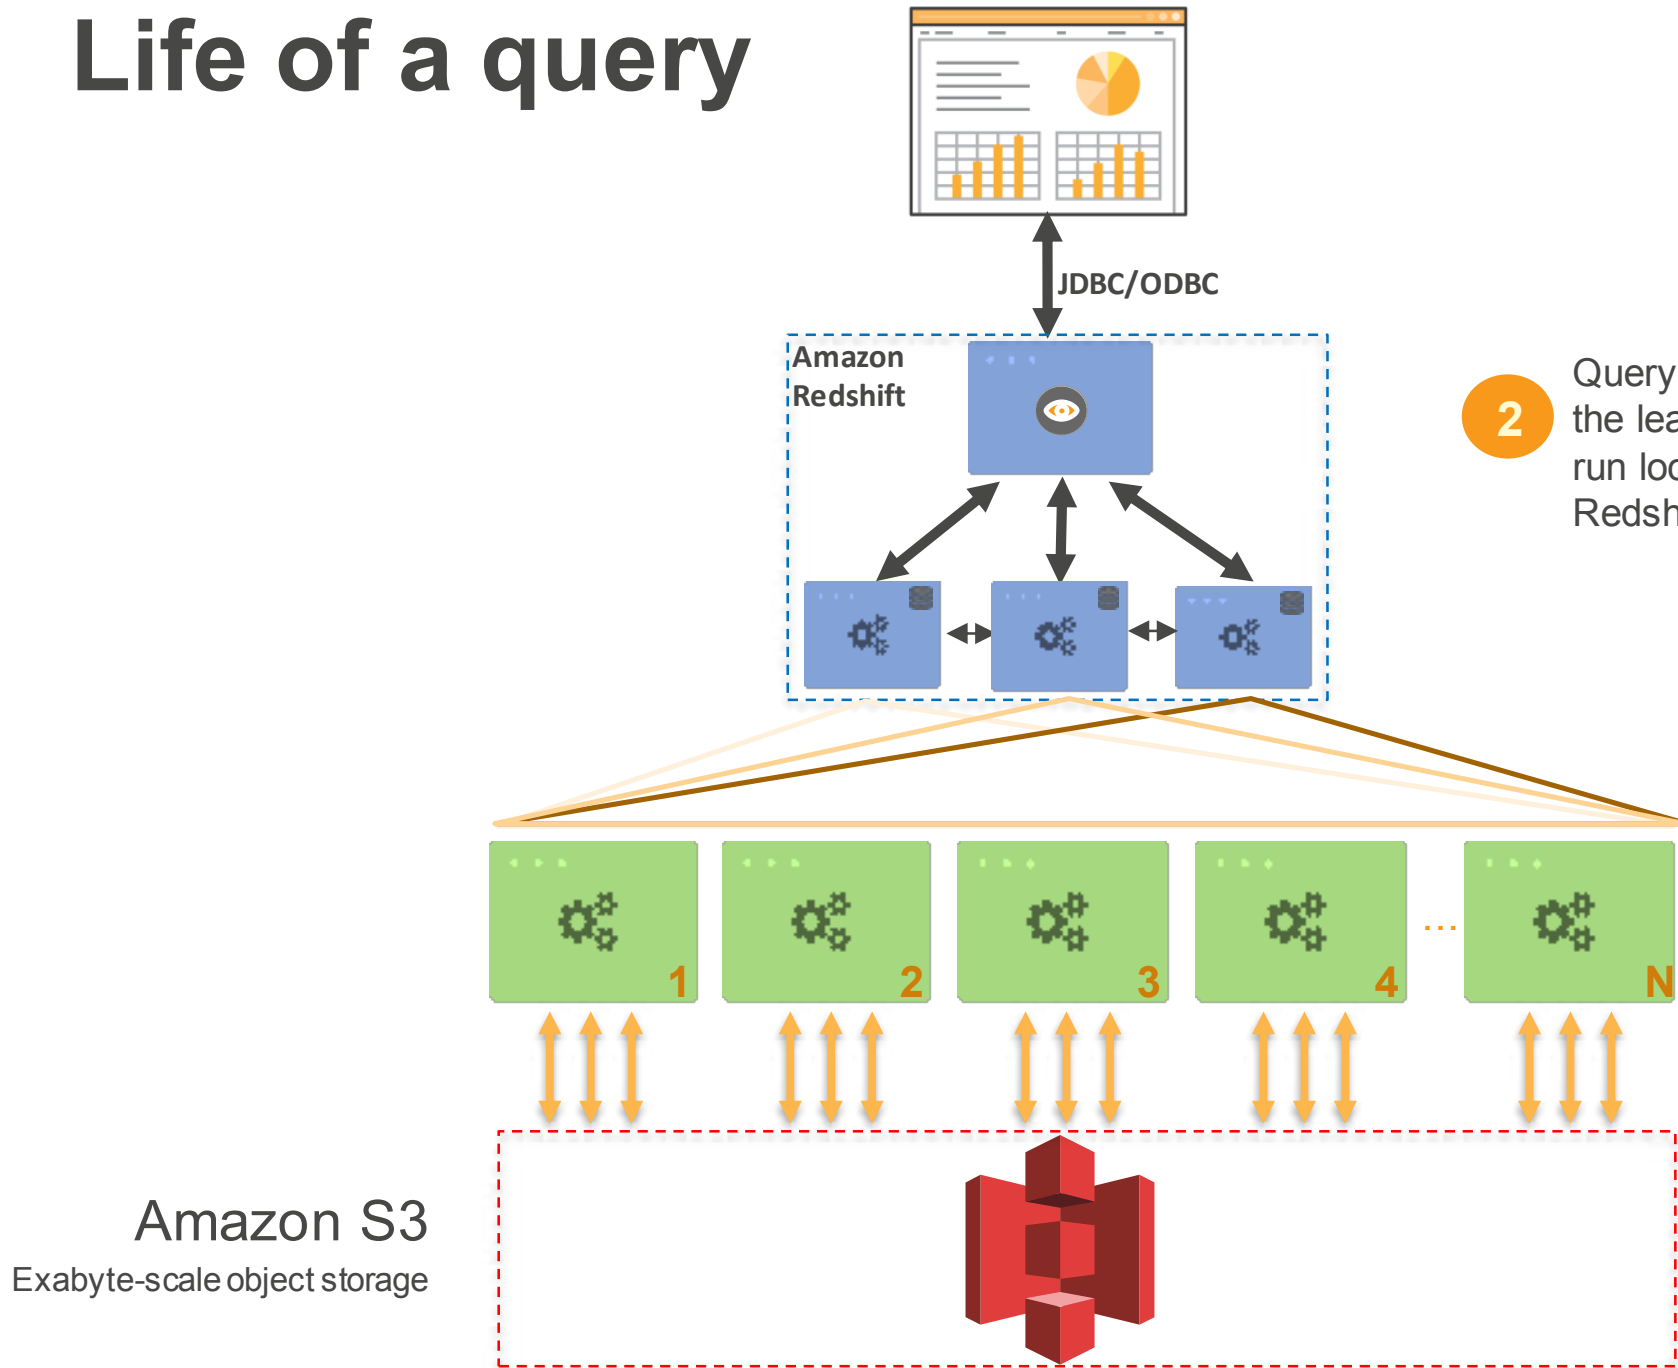


Amazon S3
Exabyte-scale object storage



Data Catalog
Apache Hive Metastore


Life of a query



2

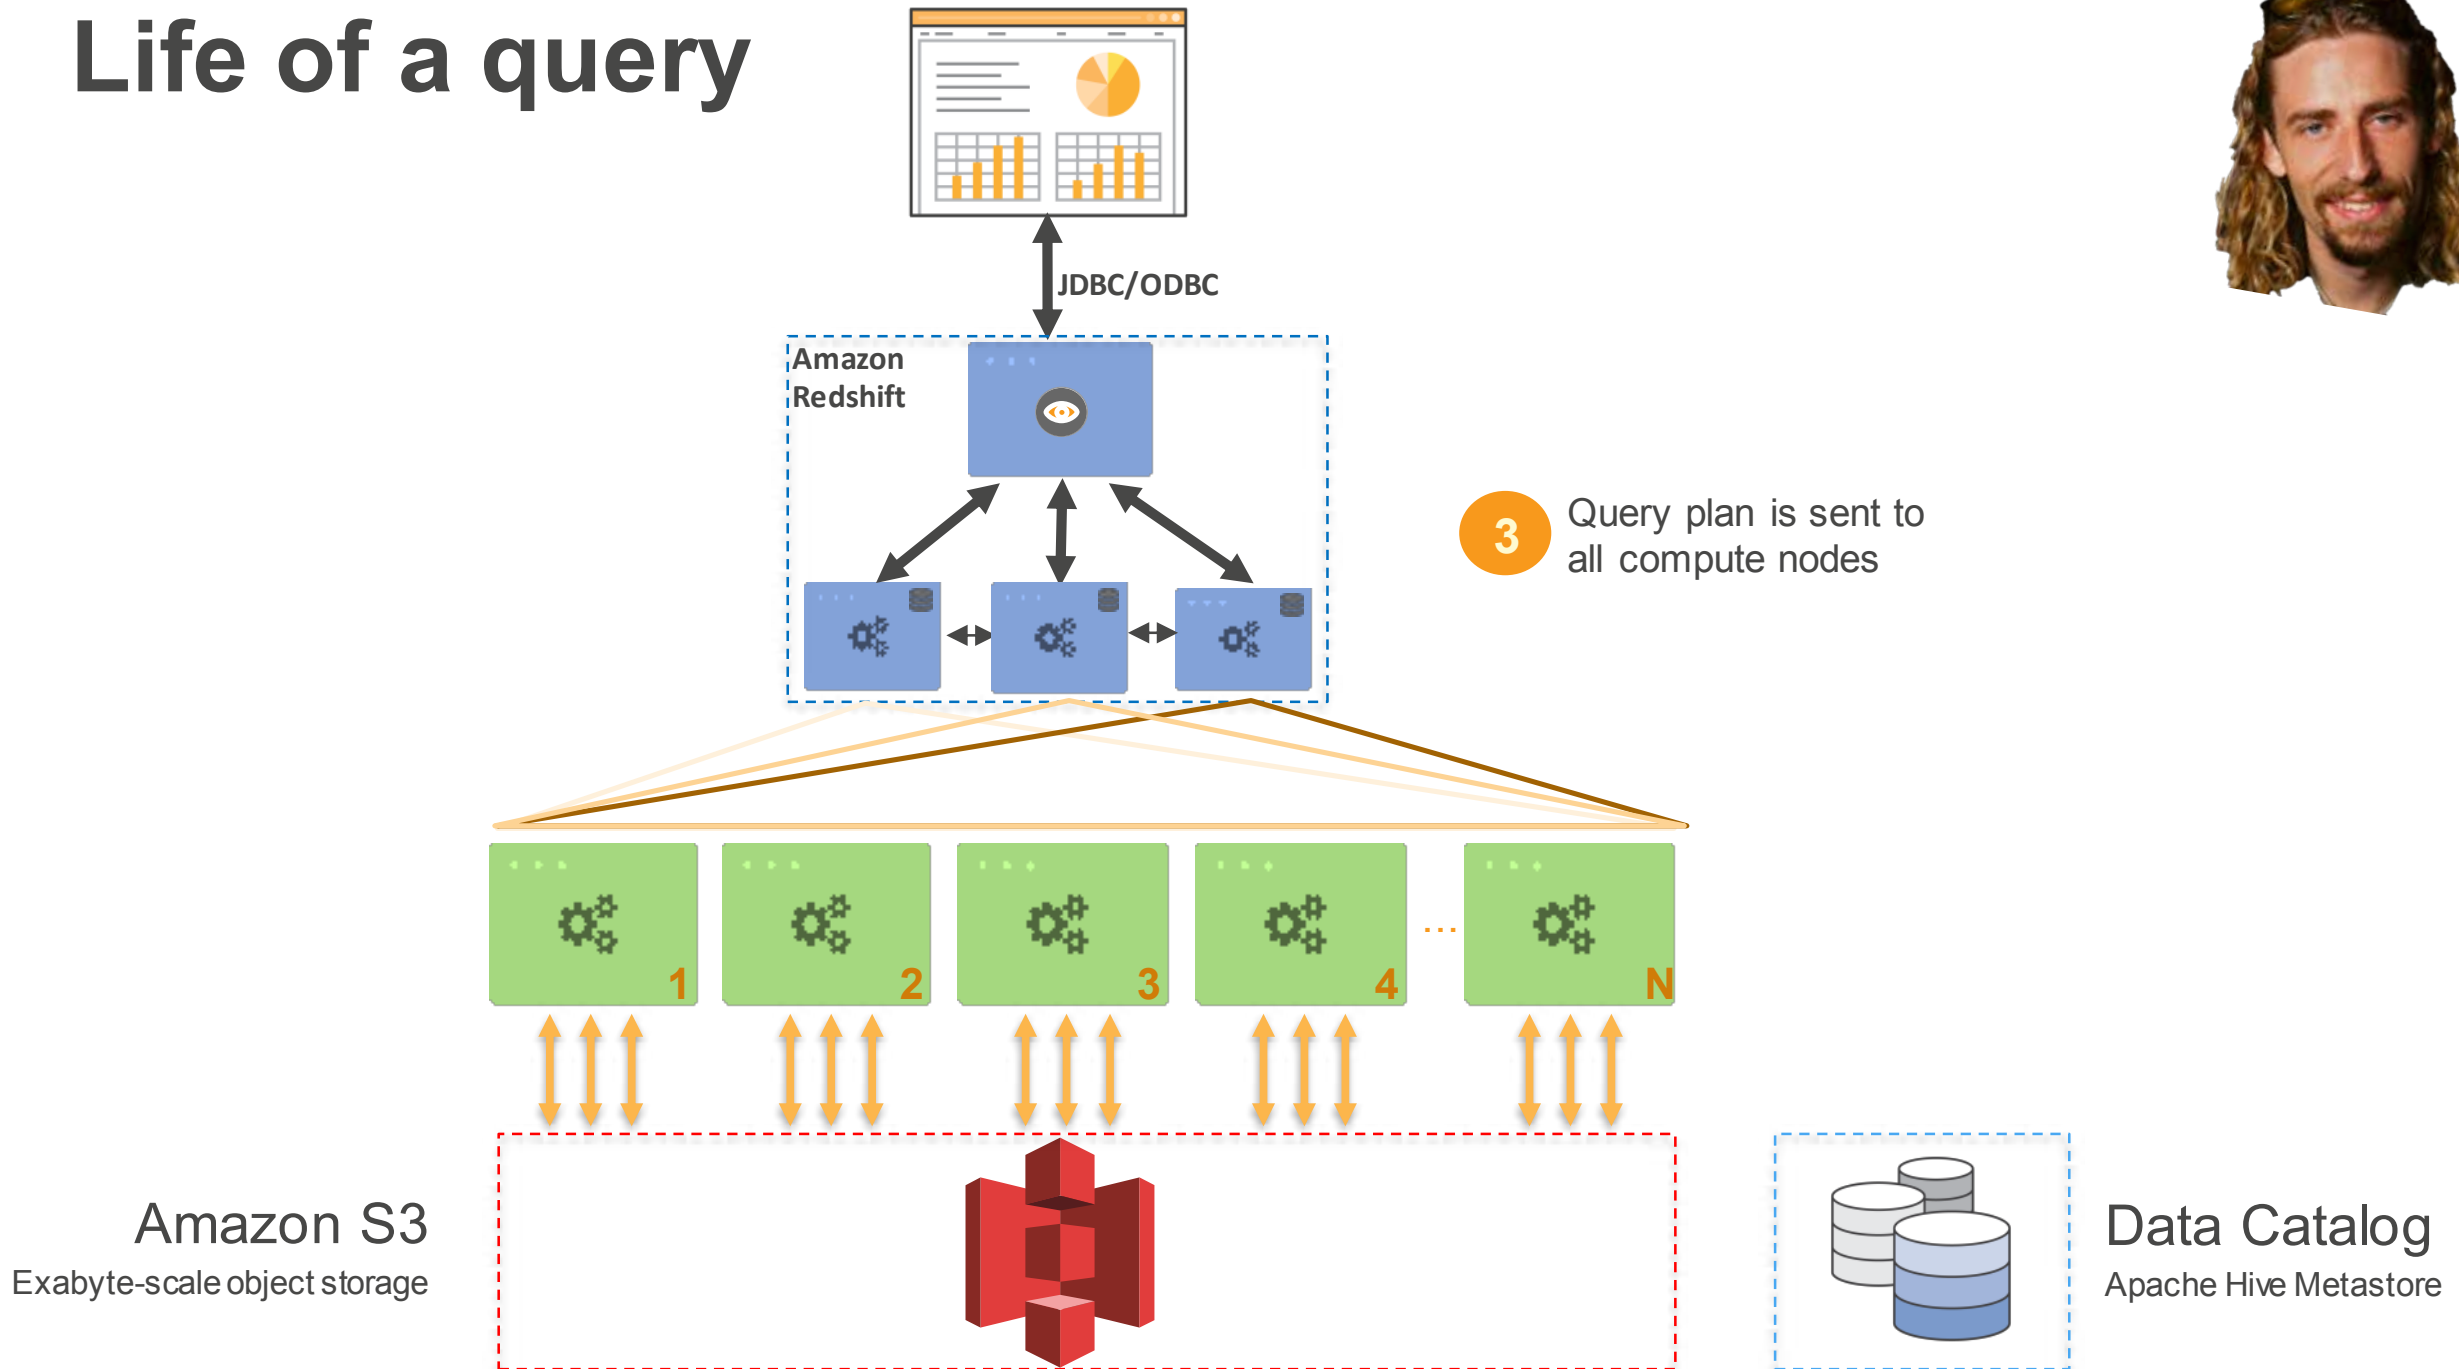
Query is optimized and compiled at the leader node. Determine what gets run locally and what goes to Amazon Redshift Spectrum

Amazon S3
Exabyte-scale object storage



Data Catalog
Apache Hive Metastore

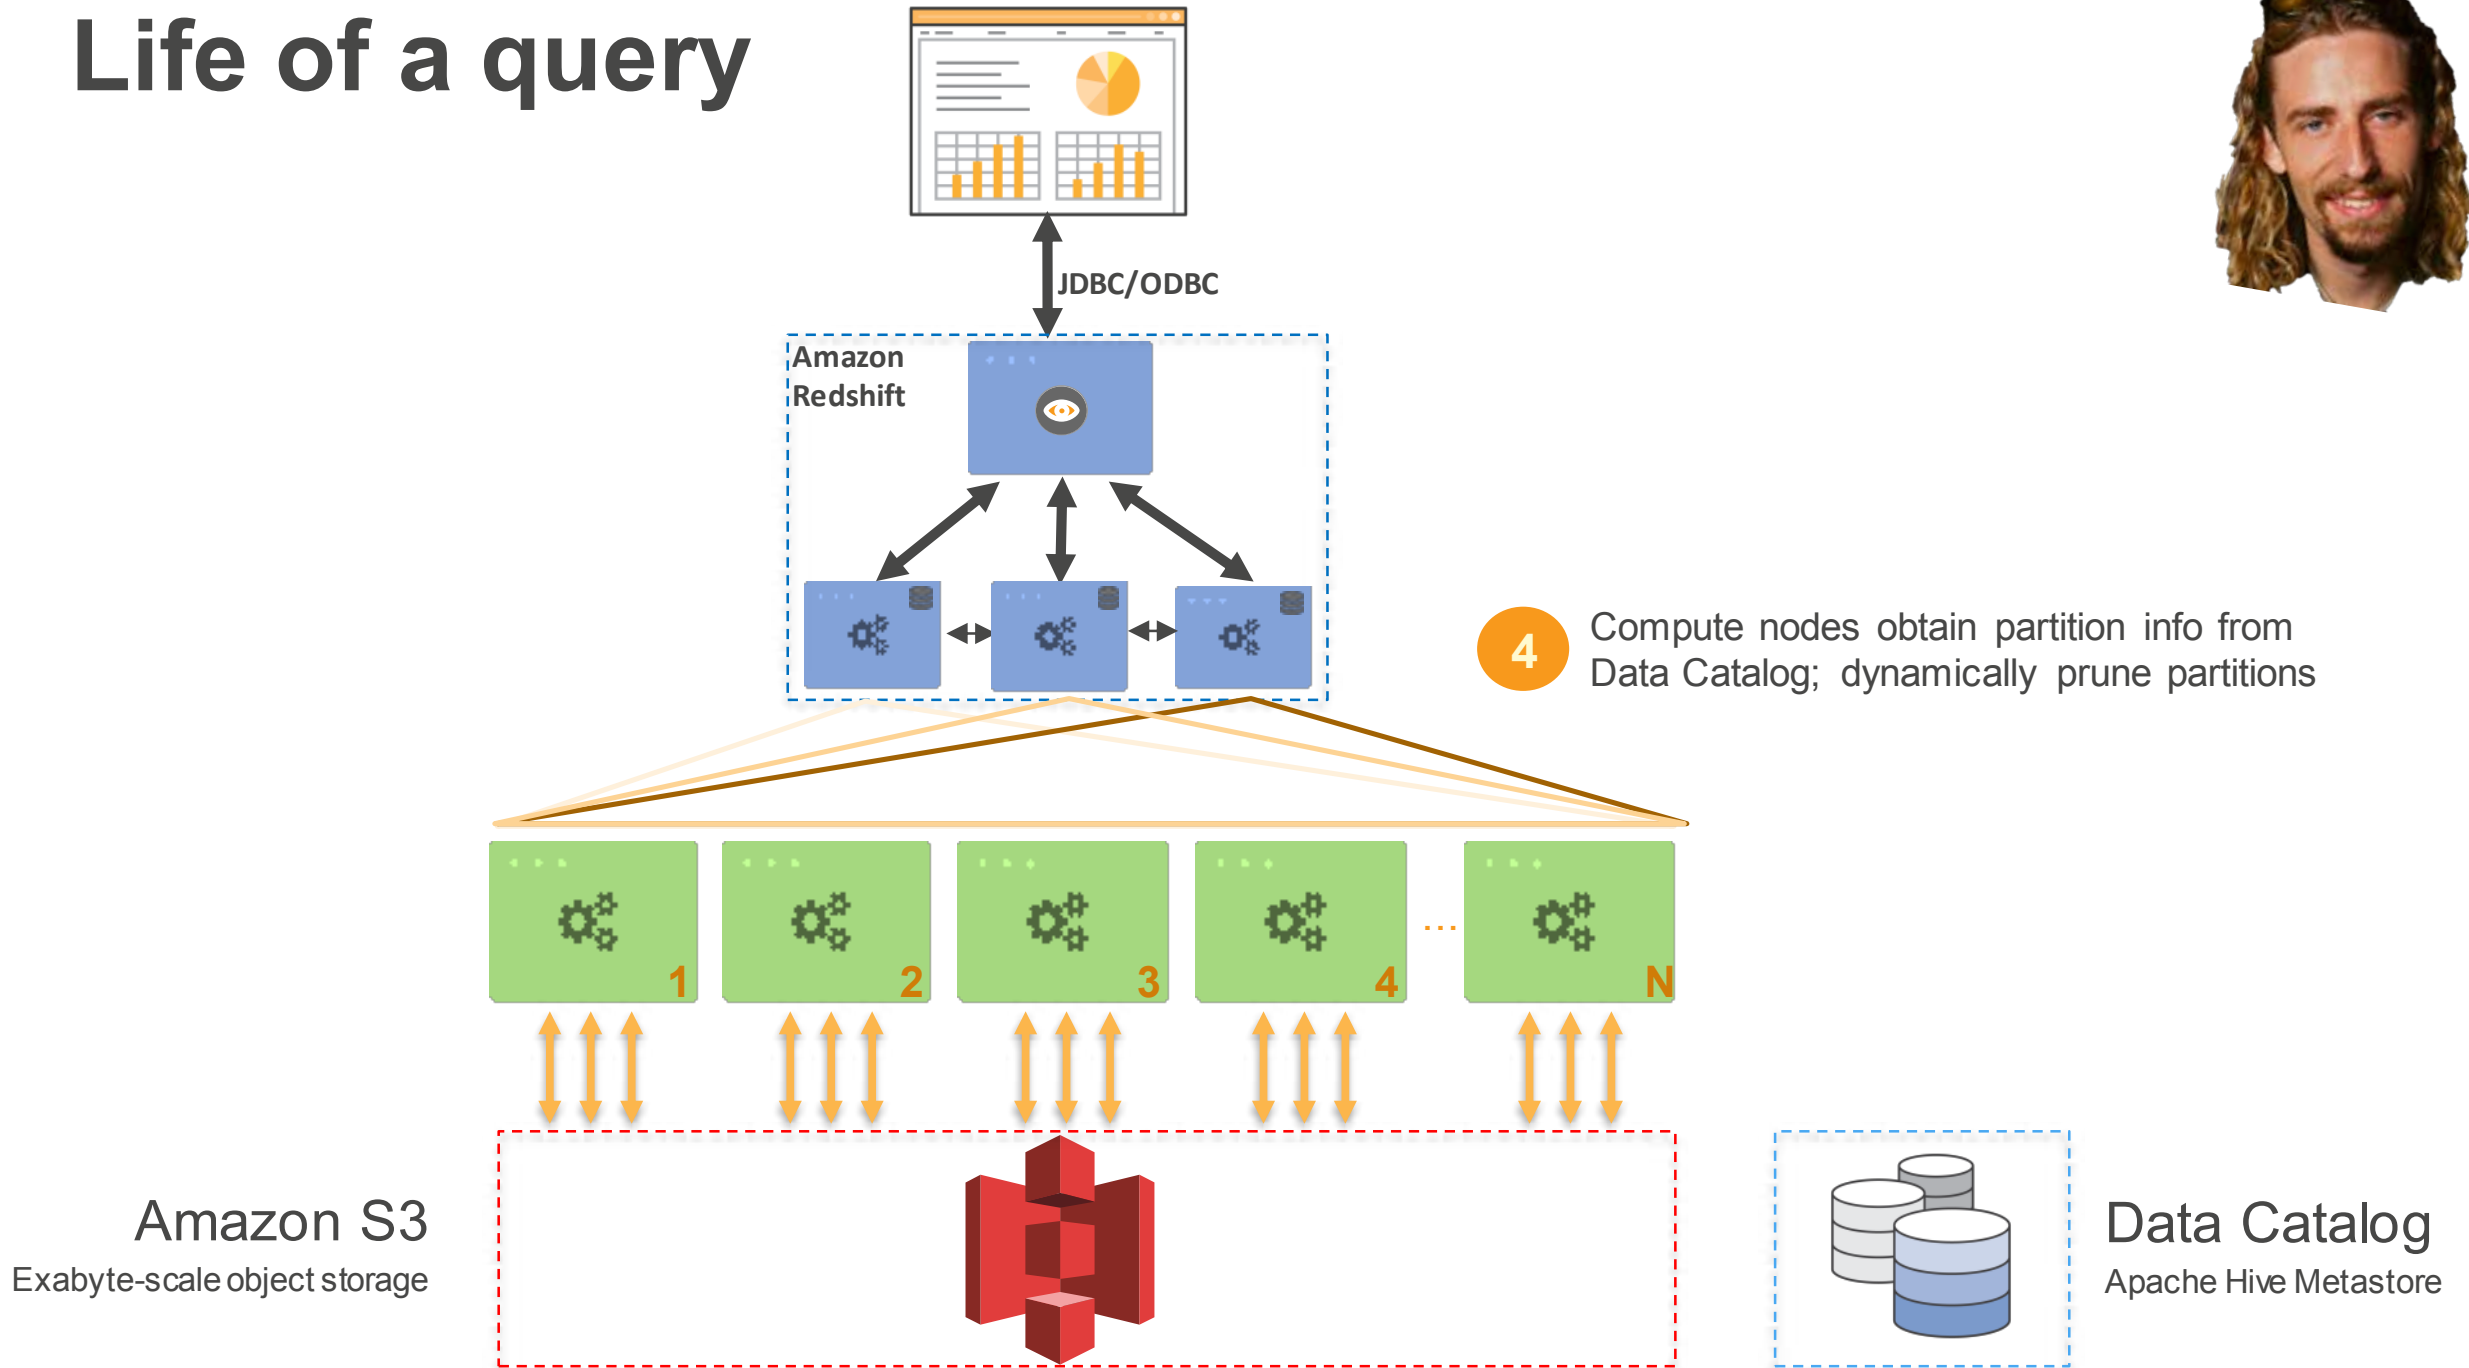
Life of a query



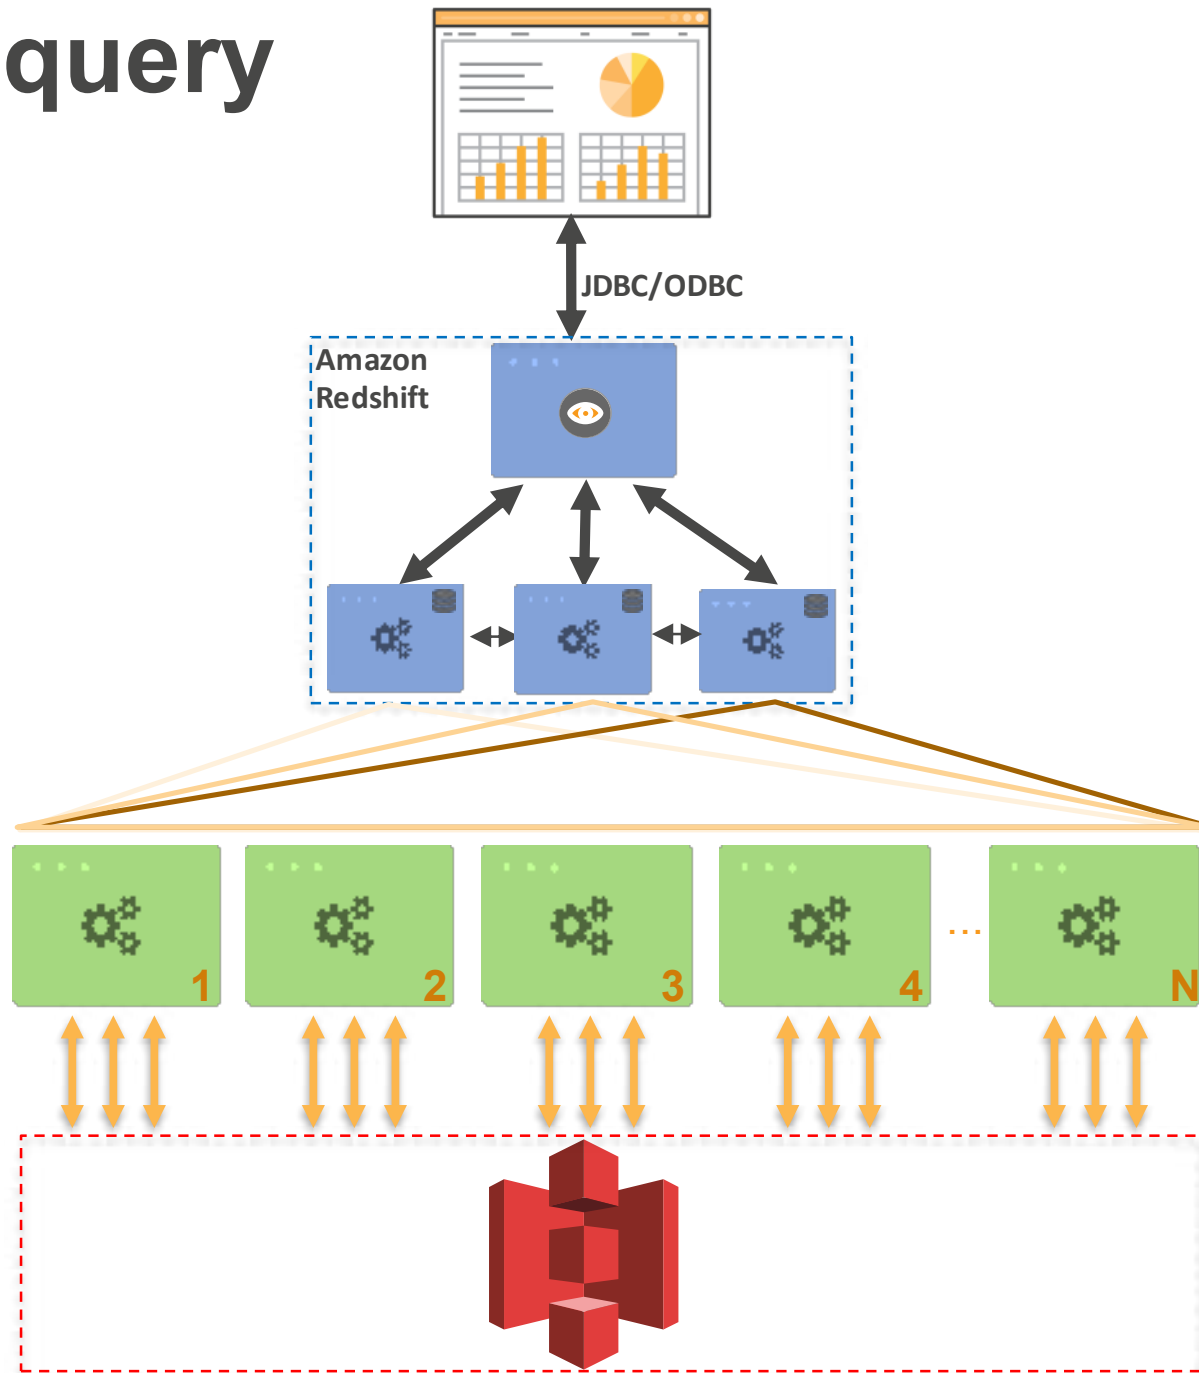
Amazon S3
Exabyte-scale object storage

Data Catalog
Apache Hive Metastore

Life of a query

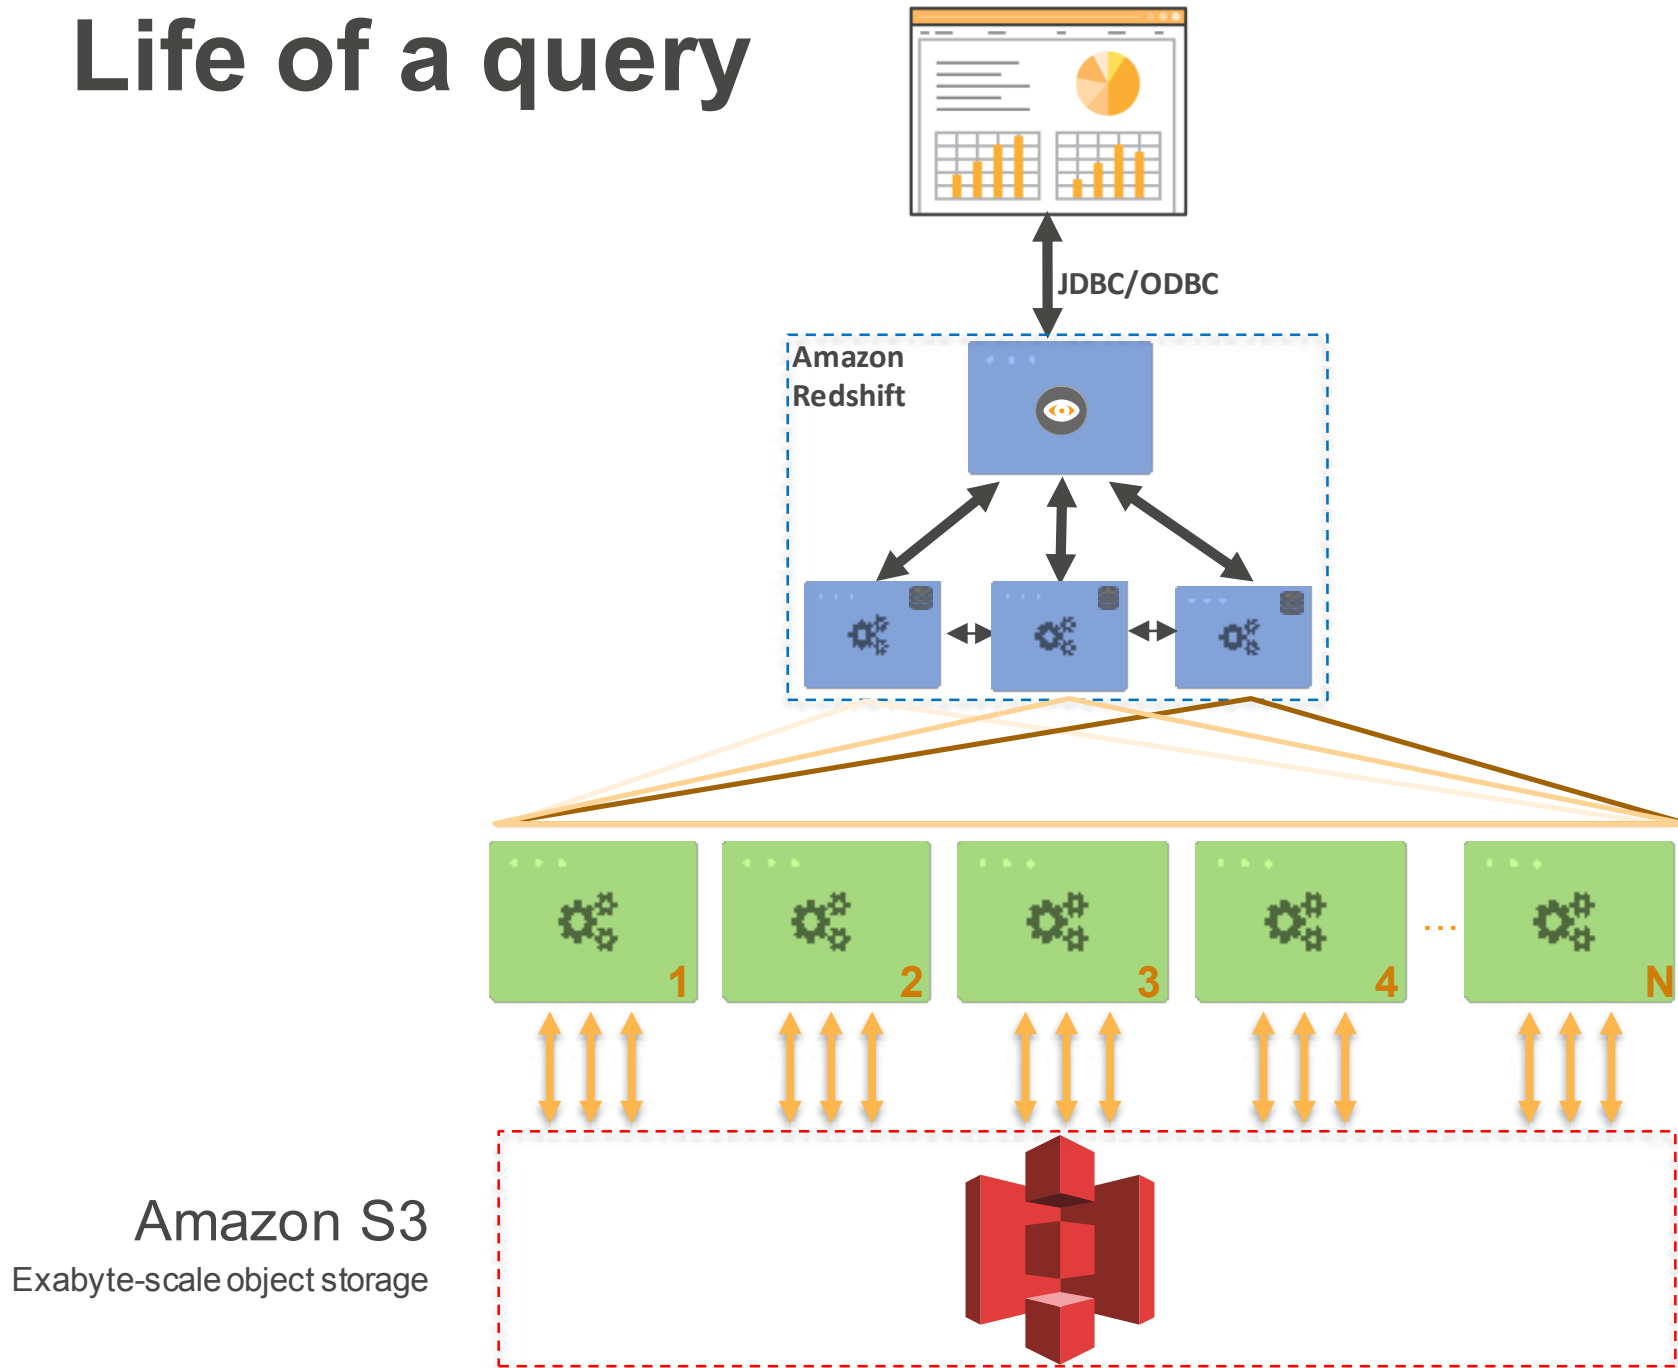


Life of a query




5 Each compute node issues multiple requests to the Amazon Redshift Spectrum layer

Life of a query




6 Amazon Redshift Spectrum nodes scan your S3 data


Life of a query



7 Amazon Redshift Spectrum projects, filters, joins and aggregates

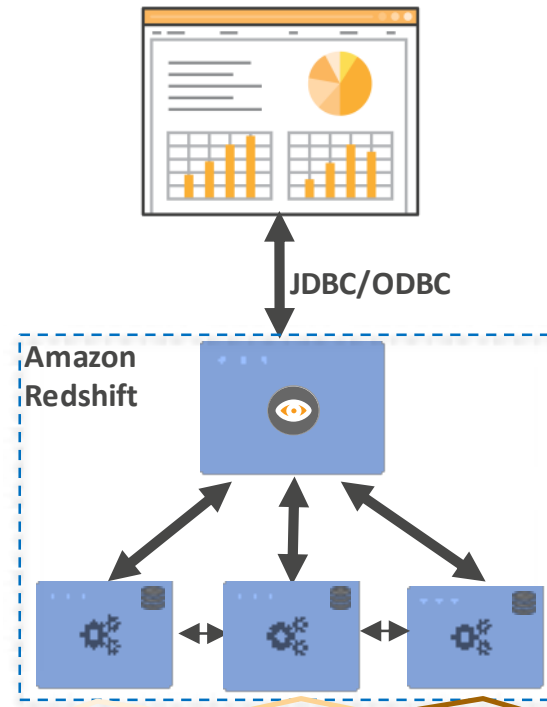


Amazon S3
Exabyte-scale object storage




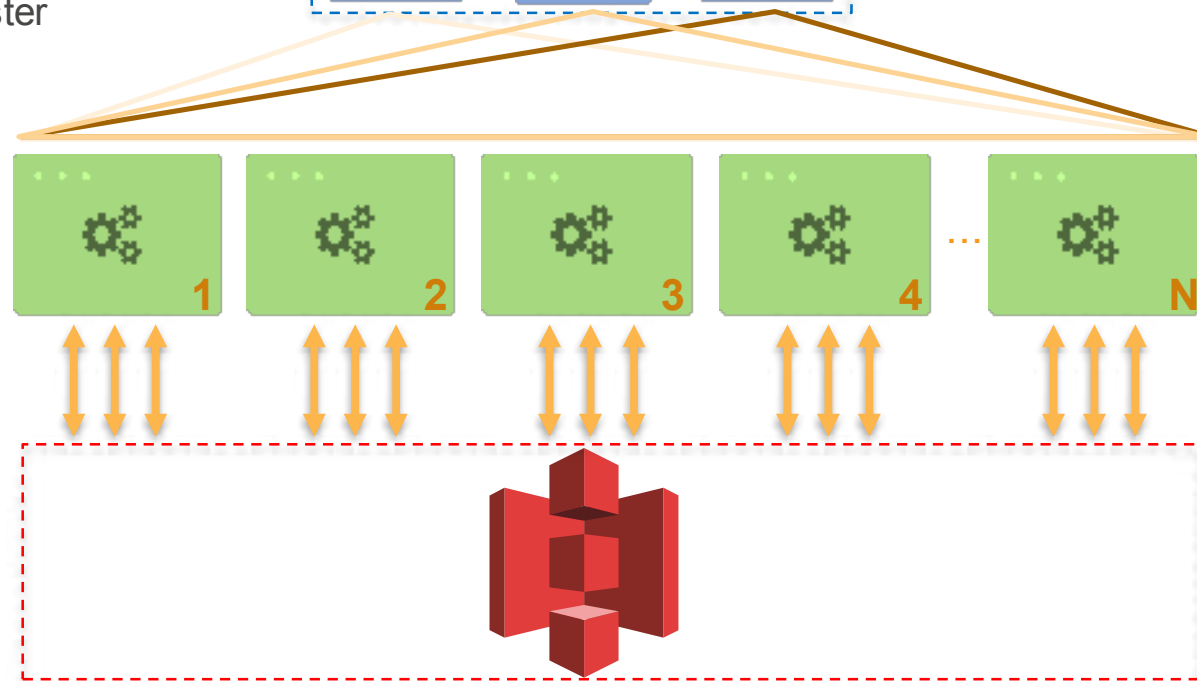
Data Catalog
Apache Hive Metastore

Life of a query



8

Final aggregations and joins with local Amazon Redshift tables done in-cluster

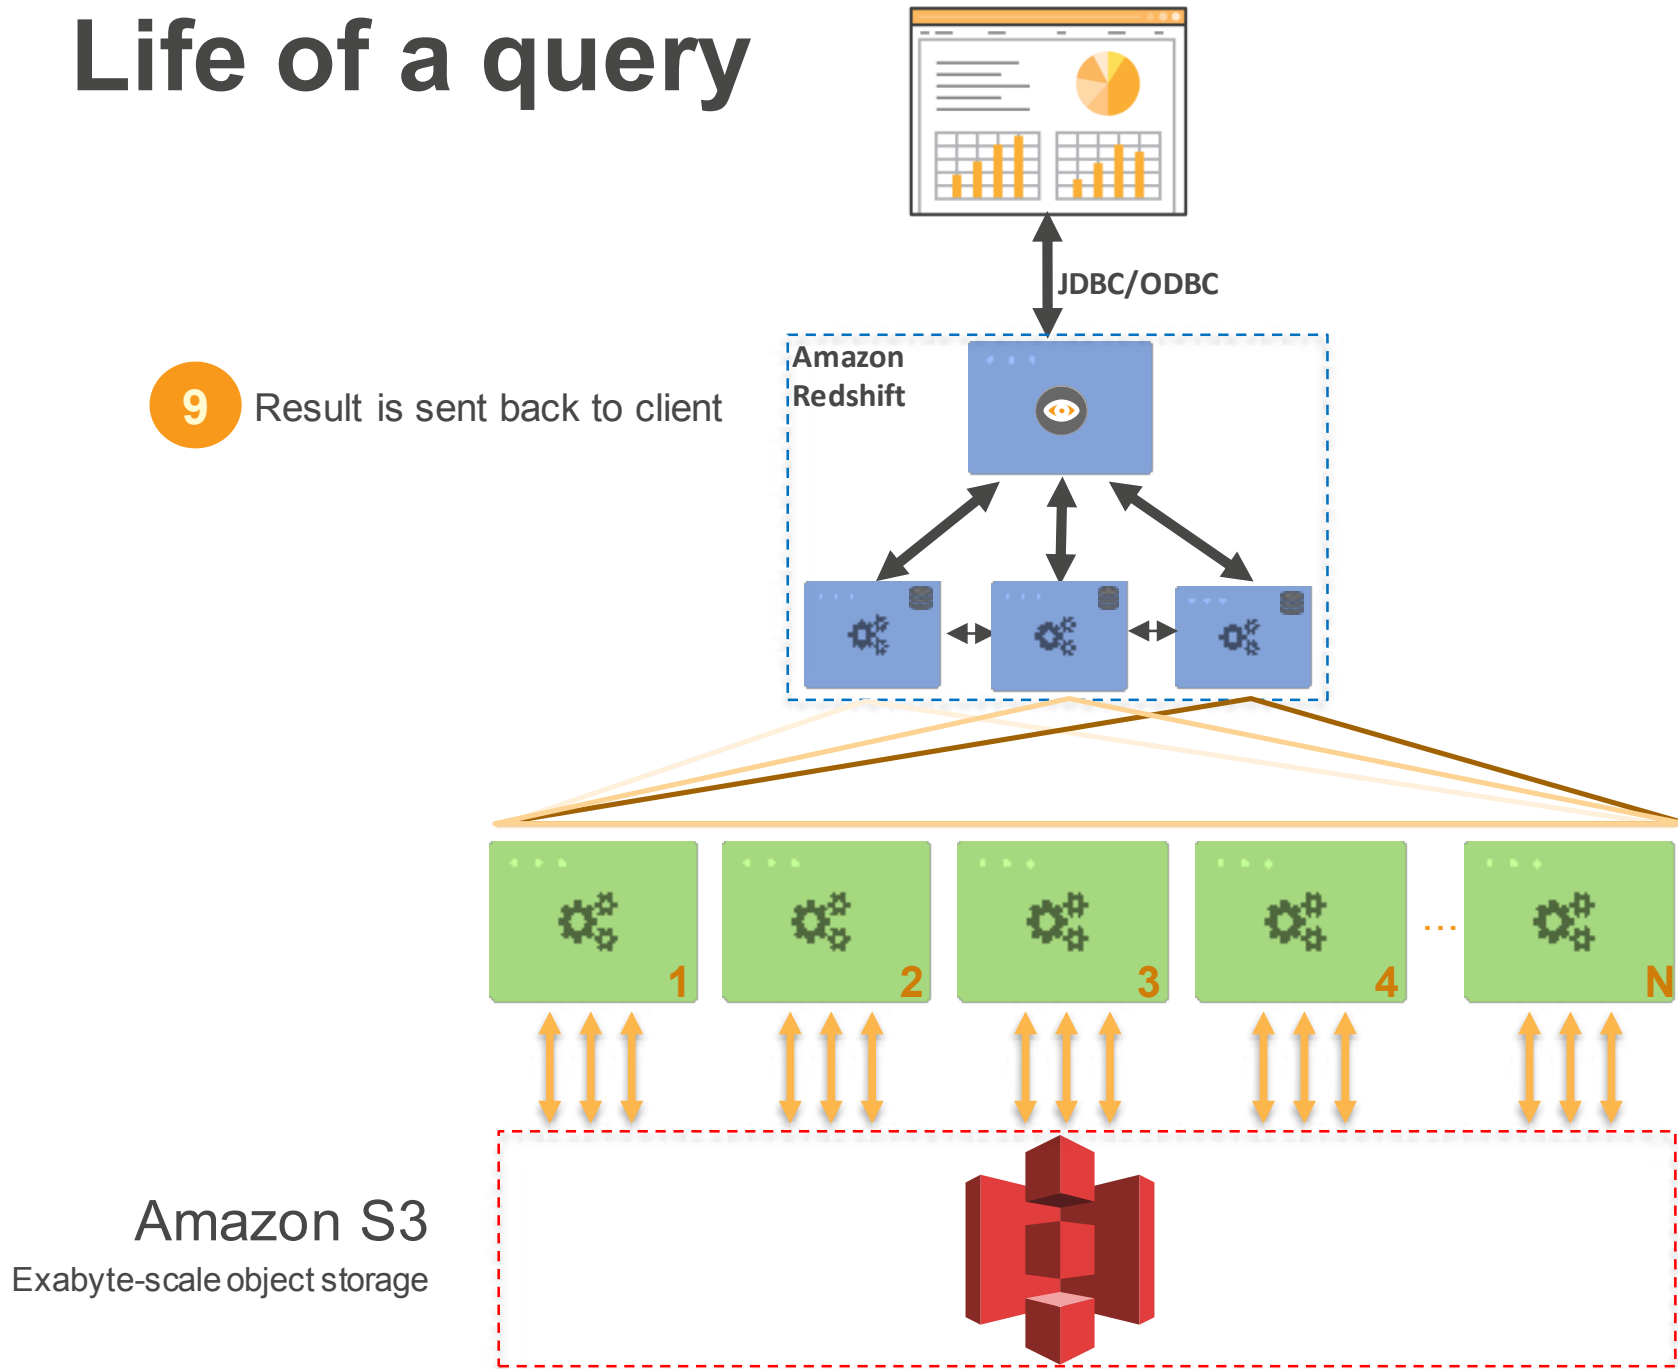


Data Catalog
Apache Hive Metastore


Life of a query



9 Result is sent back to client

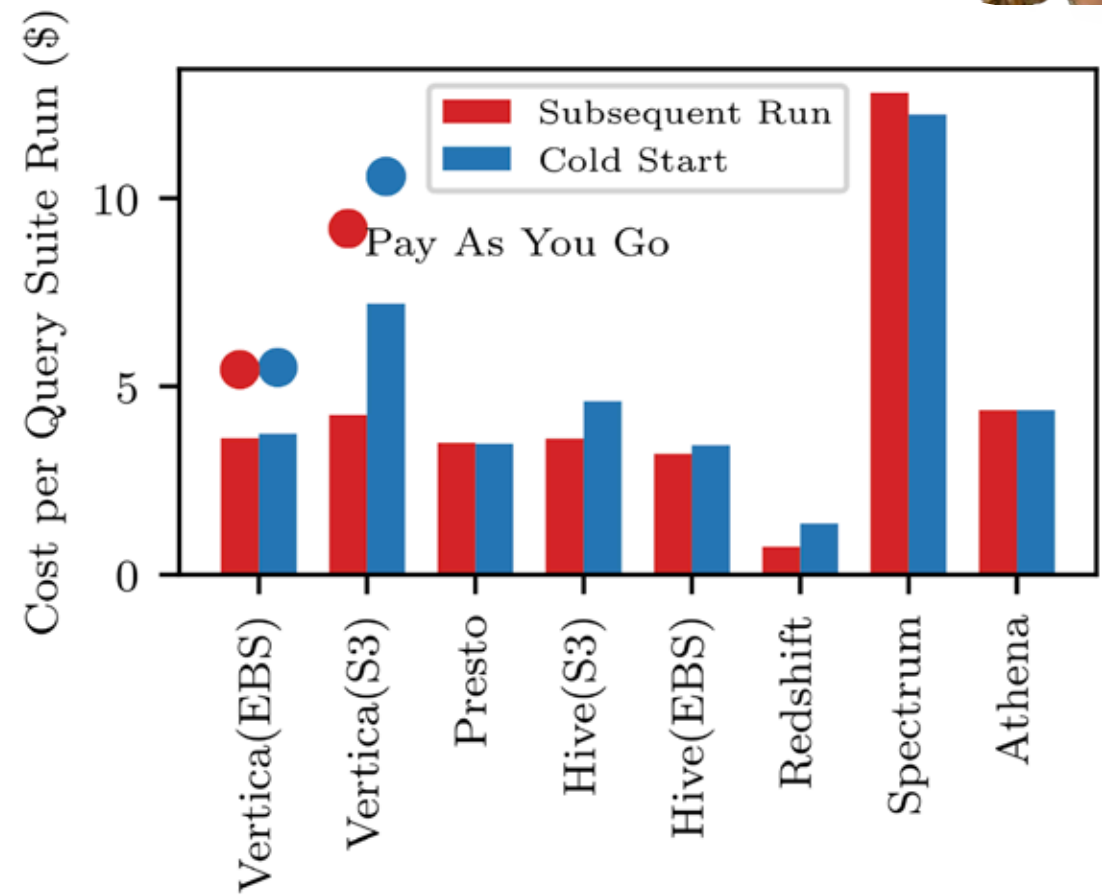
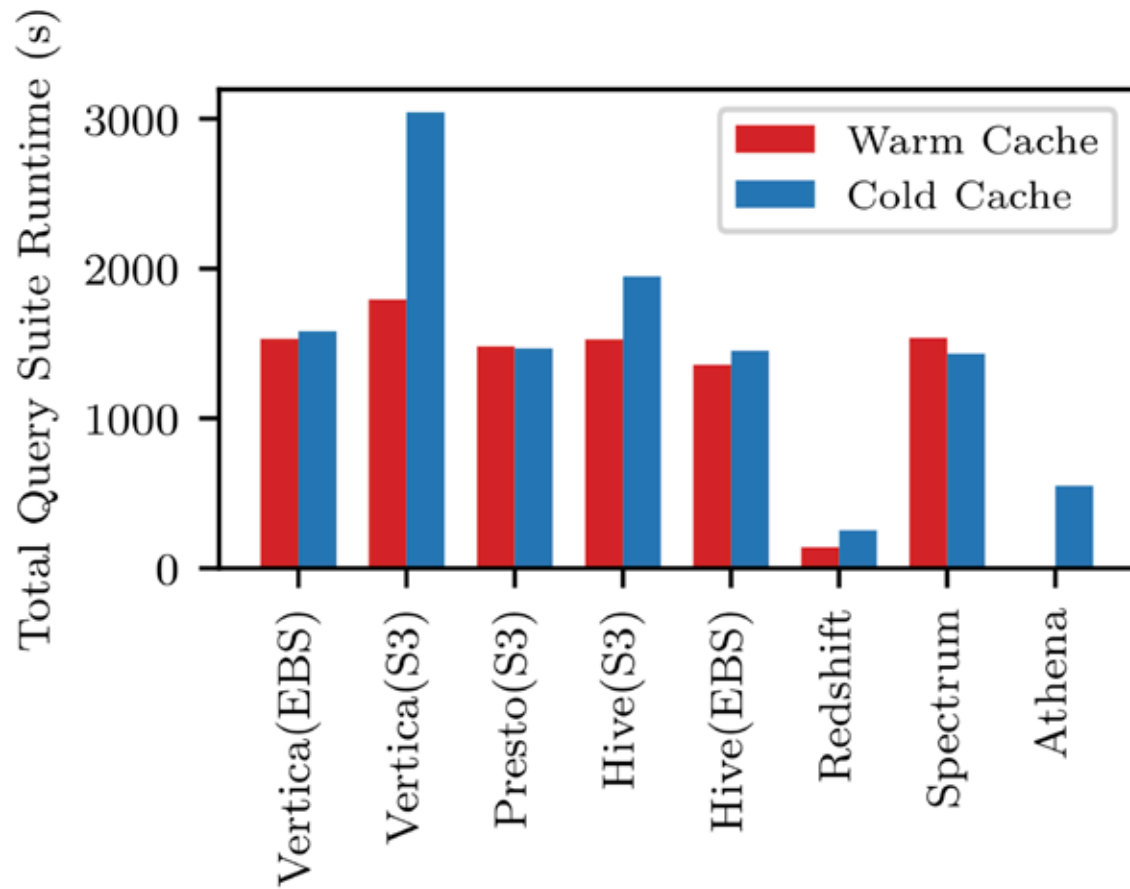


Amazon S3
Exabyte-scale object storage



Data Catalog
Apache Hive Metastore

Comparison of DW



[Junjay Tan](#), [Thanaa Ghanem](#), [Matthew Perron](#), [Xiangyao Yu](#), Michael Stonebraker, [David J. DeWitt](#), [Marco Serafini](#), [Ashraf Aboulnaga](#), [Tim Kraska](#): Choosing A Cloud DBMS: Architectures and Tradeoffs. [PVLDB 12\(12\)](#): 2170-2182 (2019)

TakeYourNickelBack.com



TAKEYOURNICKELBACK.COM



- As we get more cash hungry, we decide to do analytics on our companies data. We purchase data from Cambridge Analytica to find more potential customers (shh)
- Now we need an analytical system that can store tens of terabytes of data and query it quickly (for our new Marketing team)



Data Warehousing systems

- Not as many options for Open source systems
- Mostly have columnar storage and are optimized for analytical queries on TBs of data
- Run SQL but are not optimized for small transactions
- Usually ingesting new data from operational database at a granularity of minutes to days



Extract Transform Load



- Need a way to get our data from operational system to the analytical systems.
- Extract data from the operational system (where we are making sales)
- Transform it as necessary for our analytical system
- Load it into the analytical system
- Could use Spark or Hadoop to write these jobs (and many others)



DIMES+STORE

The official unofficial Nickelback Fan-Store



NICKELBACK

FEED THE MACHINE TOUR



Unfortunately it only got worse



Last attempt



Make Money by working on HITs

HITs - *Human Intelligence Tasks* - are individual tasks that you work on. [Find HITs now.](#)

As a Mechanical Turk Worker you:

- Can work from home
- Choose your own work hours
- Get paid for doing good work



Get Results from Mechanical Turk Workers

Ask workers to complete HITs - *Human Intelligence Tasks* - and get results using Mechanical Turk. [Get started.](#)

As a Mechanical Turk Requester you:

- Have access to a global, on-demand, 24 x 7 workforce
- Get thousands of HITs completed in minutes
- Pay only when you're satisfied with the results





Microtasking – Virtualized Humans

- Current leader: Amazon Mechanical Turk
- Requestors place Human Intelligence Tasks (HITs)
 - Minimum price: \$0.01
 - Other parameters: #of replicas (assignments), expiration, **User Interface**,...
 - API-based: “createHit()”, “getAssignments()”, “approveAssignments()”, “forceExpire()”
 - Requestors approve jobs and payment
- Workers (a.k.a. “turkers”) choose jobs, do them, get paid

AWS Lambda



<https://aws.amazon.com/blogs/architecture/ten-things-serverless-architects-should-know/>

Aftermath

Nickelback ranked one of the worst bands ever



According to [rollingstone.com](http://www.rollingstone.com)



Readers' Poll: The Ten Worst Bands of the Nineties

1. Creed. It's no surprise that Creed won this poll.
2. Nickelback. It's hard not to feel a little bad for Nickelback.
3. Limp Bizkit. If you think that Limp Bizkit fans are a bunch of bullies, you aren't alone. ...
4. Hanson. ...
5. Nirvana. ...
6. Hootie and the Blowfish. ...
7. Bush. ...
8. Spin Doctors. ...



DimeStore was banned on AWS because of UA violations



<https://www.rollingstone.com/music/music-lists/readers-poll-the-ten-worst-bands-of-the-nineties-13654/5-nirvana-241095/>

TakeYourNickelBack.com



Identity Crisis



- As our profits fall to zero, I suffer an identity crisis as I become the very thing that I hated most: the mainstream establishment. Also I'm millions of dollars in debt after buying hundreds of servers and my power bill is through the roof.

