



Building a Customer 360 Graph Application on Amazon Neptune

Justin Thomas

Principal Amazon Neptune Specialist

Agenda


- What is a Customer 360 Graph?
- Why a Graph Database and Amazon Neptune?
- Customer examples and how to get started

What is a Customer 360 Graph?



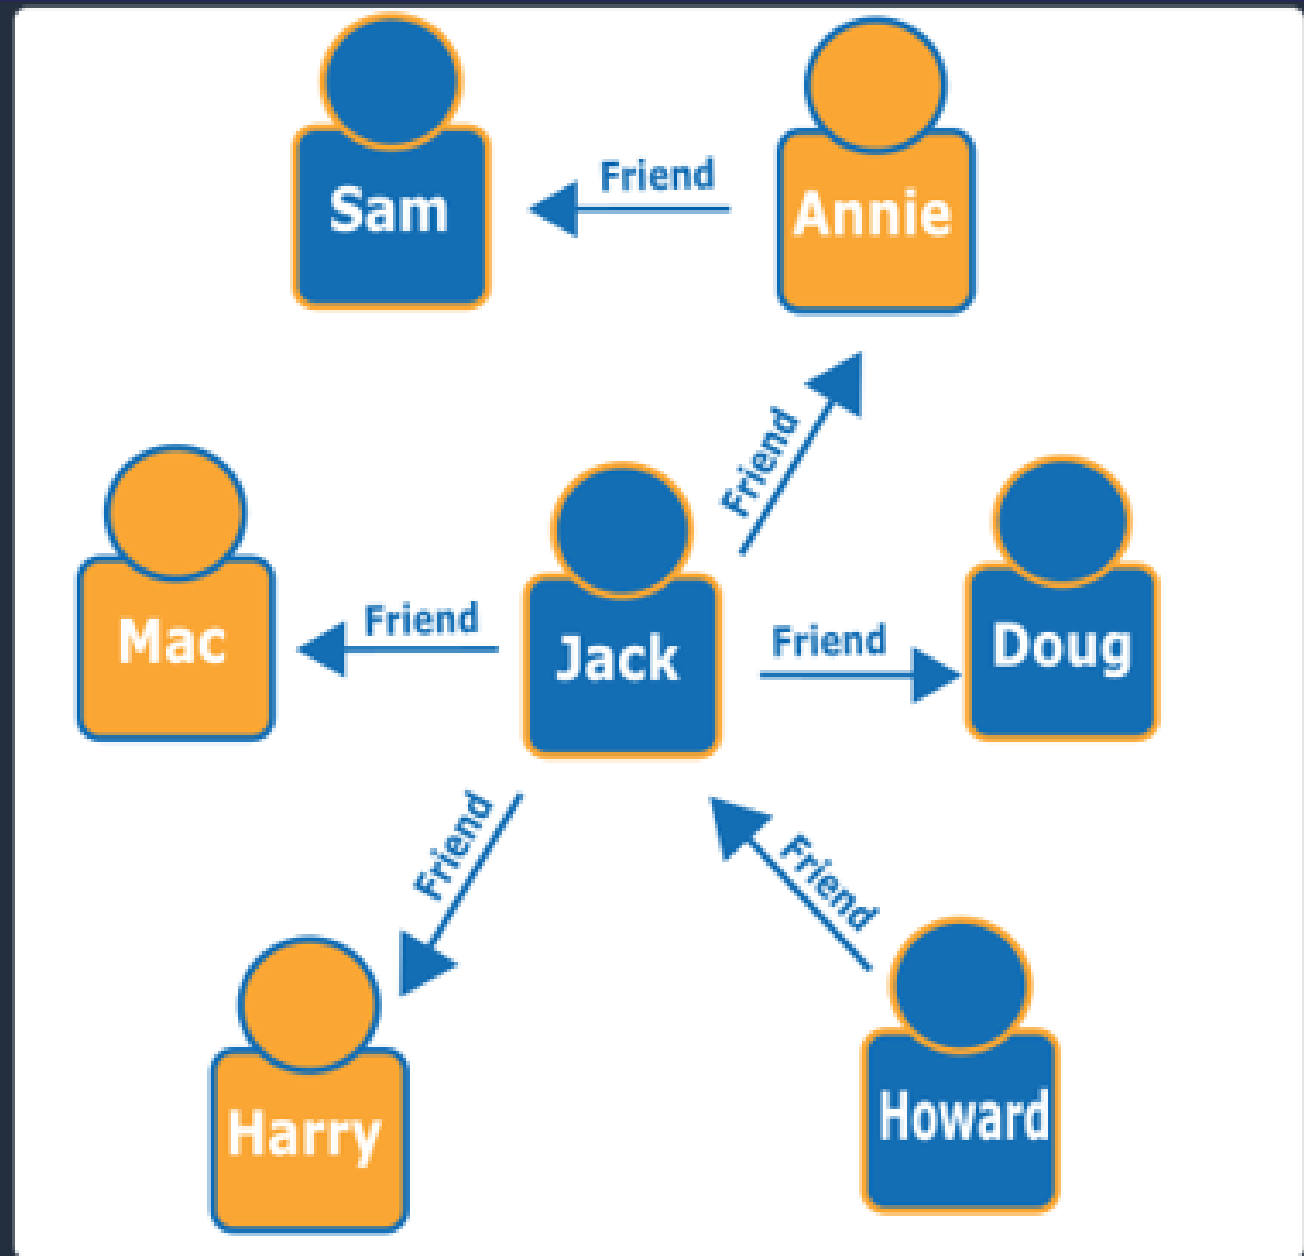
What is a Graph?

- Graph databases are optimized to store and query relationships
- Nodes represent real-world objects
- Edges store relationships between objects
- Properties and Labels can be added to both Nodes and Edges



A Social Network is a Graph

- Nodes: People
- Edges: Friend Relationships
- Edges can have direction



Consumer 360 Graph – Unified view across channels and devices



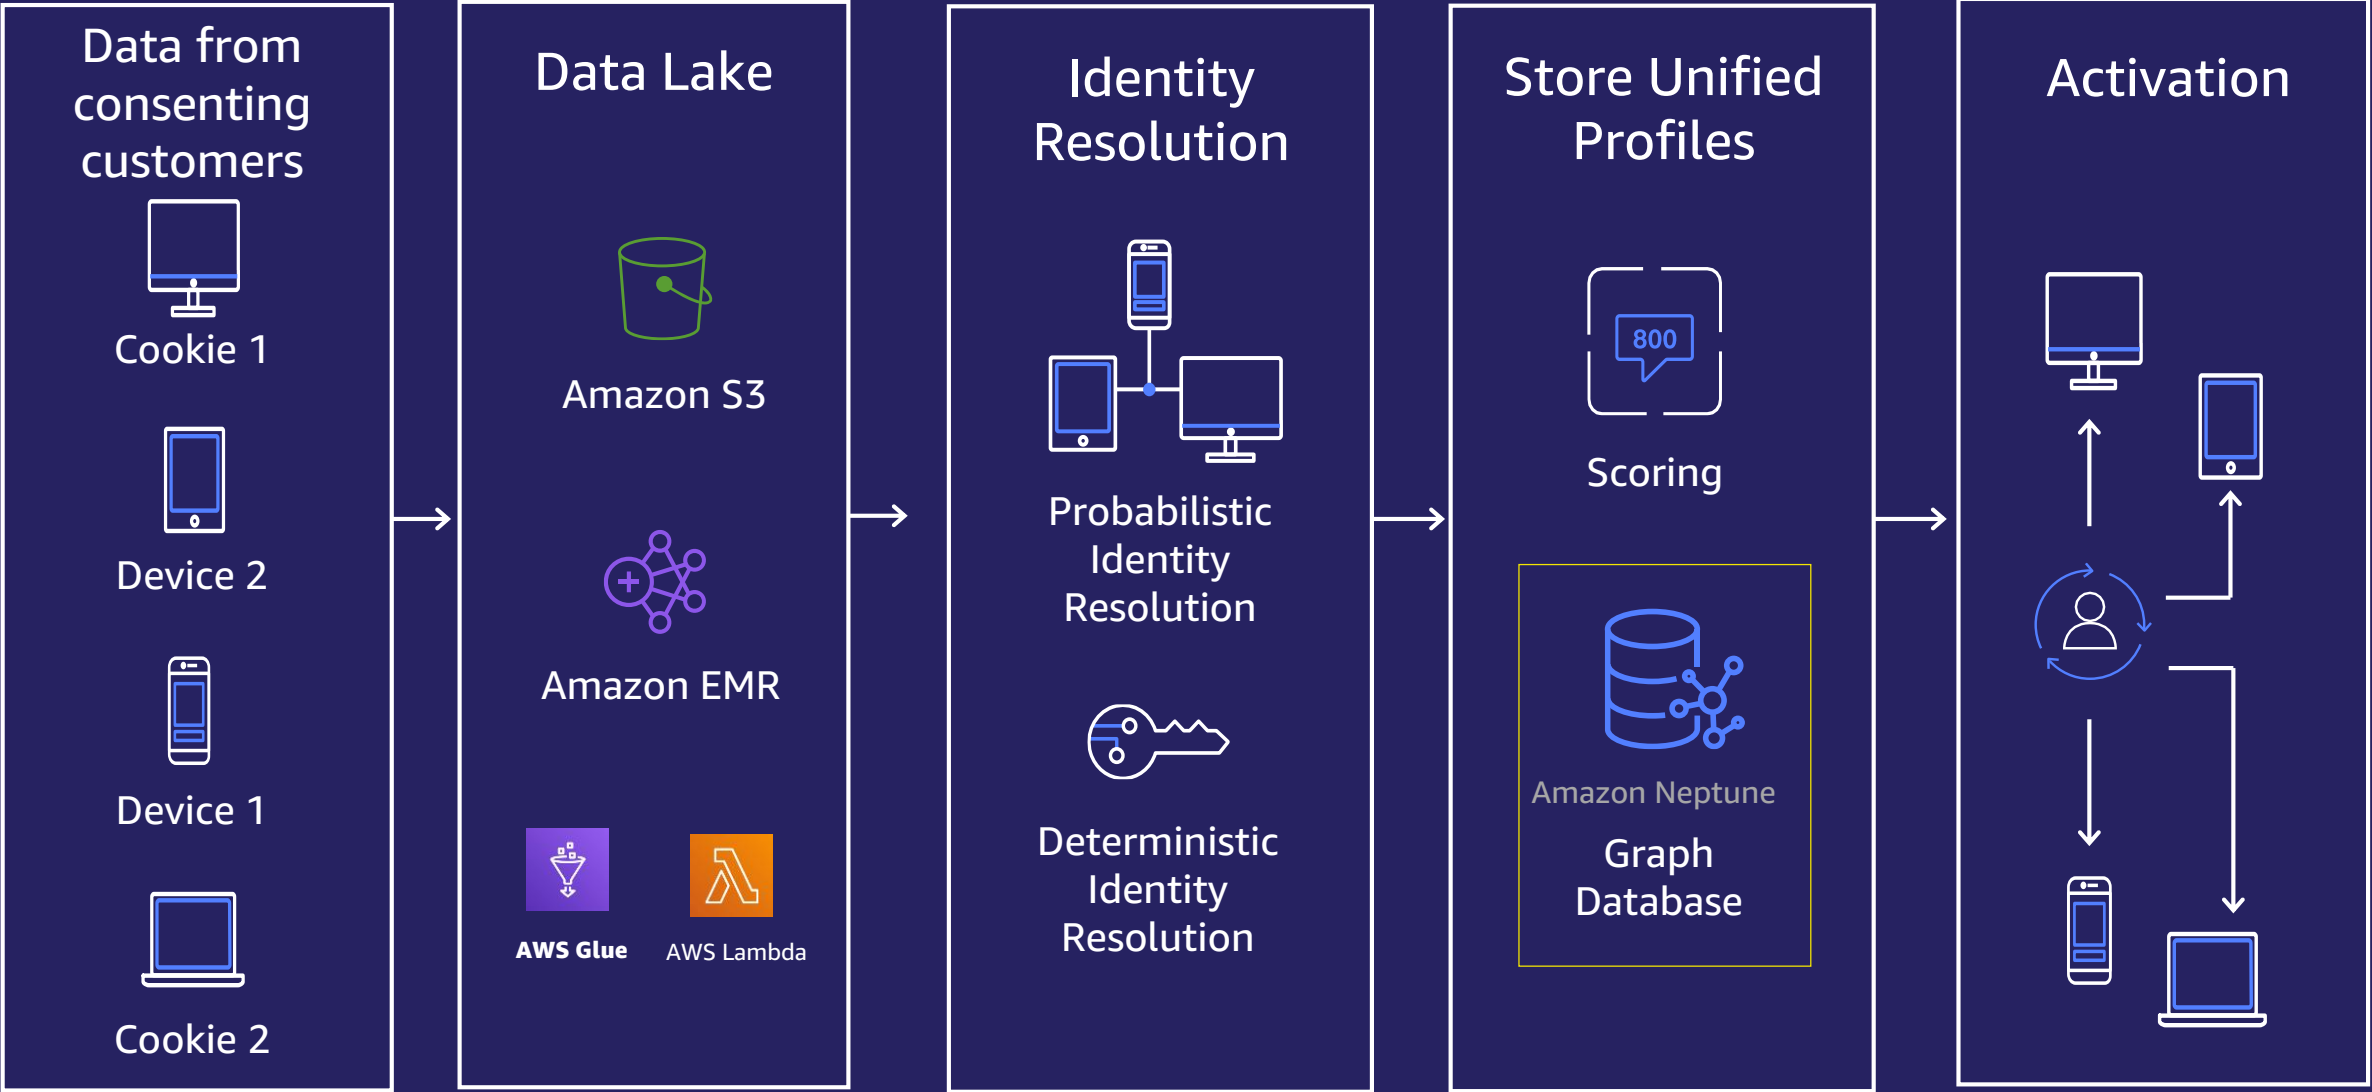
Customers interact with brands and products across multiple channels, devices, and identifiers

Lack of a common identifier linking customer interactions across devices and platforms

Consumer 360 enables capture and linking of identities, devices, interests and behaviors to develop a unified

Marketers use Consumer 360 to personalize content, products, offers and advertising

Logical Flow — Customer 360 Graph



Cox Automotive - connected consumer journey

A connected consumer journey, combined with machine learning, resulting in reliable predictive insights



LEASE/FINANCE PROPENSITY
Lease



CURRENT OWNERSHIP
Compact Sport
Sedan



ONLINE RESEARCH
Full Size SUV



LIFE EVENT
New Baby



HOUSEHOLD INCOME
\$70K-\$80K

COX AUTOMOTIVE SOURCES

Autotrader Kelley Blue Book DEALER.COM VinSolutions Dealertrack ∞xtime

EXTERNAL SOURCES

LiveRamp Adobe acxiom.



Identity resolution to power personalization



“We have addressed the immediate goals of reducing dependence on third-party cookies and building a 360-degree view of the consumer household.”

– Carlos Rendon, Principal Technical Architect,
Cox

Amazon Neptune
Amazon EMR
Amazon S3



Challenge

- Loss of third-party cookie expected to highly impact ability to personalize website content
- Build a 360-degree view of households that can be utilized across business units, many of which were the result of M&A



Solution

- Simplified data modeling that will support new use cases using a graph database
- Created an Identity Graph to define identify and combine shoppers or prospects of a “household” and leads
- Graph consists of approximately 0.5 billion edges and 0.4 billion vertices running on a Neptune cluster



Results

- Yielded twice as much browsing history per household compared to using individual cookies resulting in higher engagement, better click through rates, and higher email open rates
- Able to integrate with downstream applications (ad segmentation and vehicle recommendations)
- 180% increase average online interactions per consumer
- 380% increase average cookies per customer
- Six use cases that benefit from identity graph today



Audible scales up for business customers with help from Amazon Neptune

Challenge:

Audible for Business needed to enable its enterprise customers' administrators to maintain their own sets of end users and required a database that would scale to seamlessly manage the complex network of relationships.

Solution:

The company used Amazon Neptune, a managed graph database, to provide automated reporting and an increased self-service experience that could scale to support hundreds of thousands of end users.

Results:

- Scalable to thousands of customers
- 5K transactions per second
- 4 million edges and vertices total
- A single-source solution



Why a Graph Database?



Technical Challenges with Customer 360



Traditional Databases

- Relational databases are **not built to capture and store relationships** between billions of interconnected entities
- Challenging to deliver **low latency required for real-time cross-device** advertising targeting, personalization and other customer experience use-cases



Other Graph Solutions

- Other graph options are often **expensive, inflexible**, and may be difficult to optimize and scale globally
- Require **hardware management, provisioning**, manual scaling for open-source models



Identity Resolution & Privacy

- Identity resolution is an ongoing process due to growth of devices and new digital touchpoints—presenting challenges with schema changes, indexing and scaling
- Challenging to support privacy requirement with audit trail

Graph vs. RDBMS

Easy to model for highly connected data



Complex models

Relationships are first class citizens



Relationships are represented by Foreign Keys

Imperative Gremlin



Declarative SQL

Graphs are flexible



Rigid Schema

Graph Traversal Performance



SQL Query Performance

Why Amazon Neptune

Build and run graph applications with highly connected datasets



Cost-effective



No hardware management



Instant provisioning



Scaling



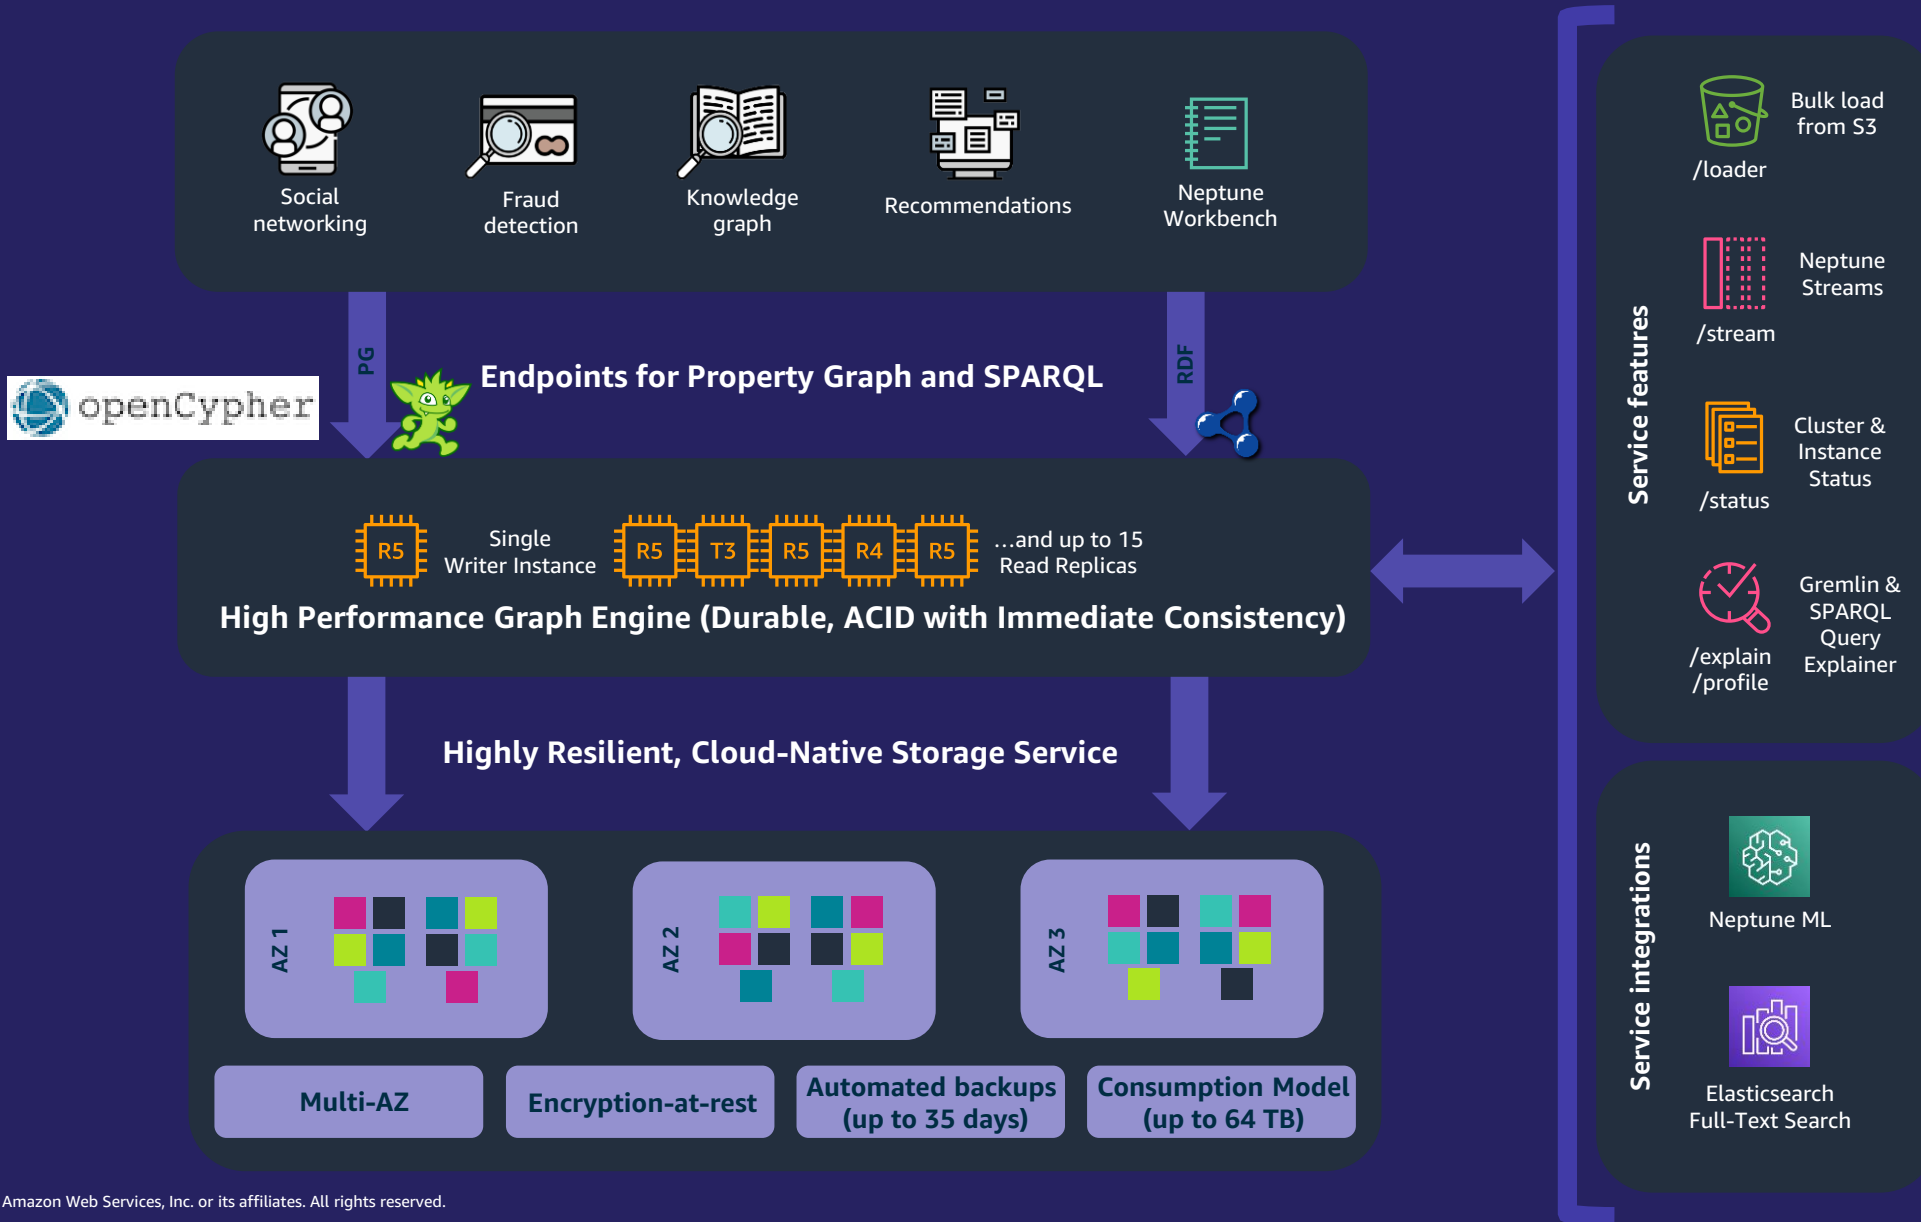
Security & compliance

- Optimized to **store and map hundreds of billions of relationships** between customer profiles, identifiers, cookies, devices, products and channels
- Enables **real-time personalization, ad targeting** with **millisecond query** response time
- Deploy high performance graph applications with popular open-source APIs such as **openCypher, Gremlin, SPARQL**, and easily migrate applications

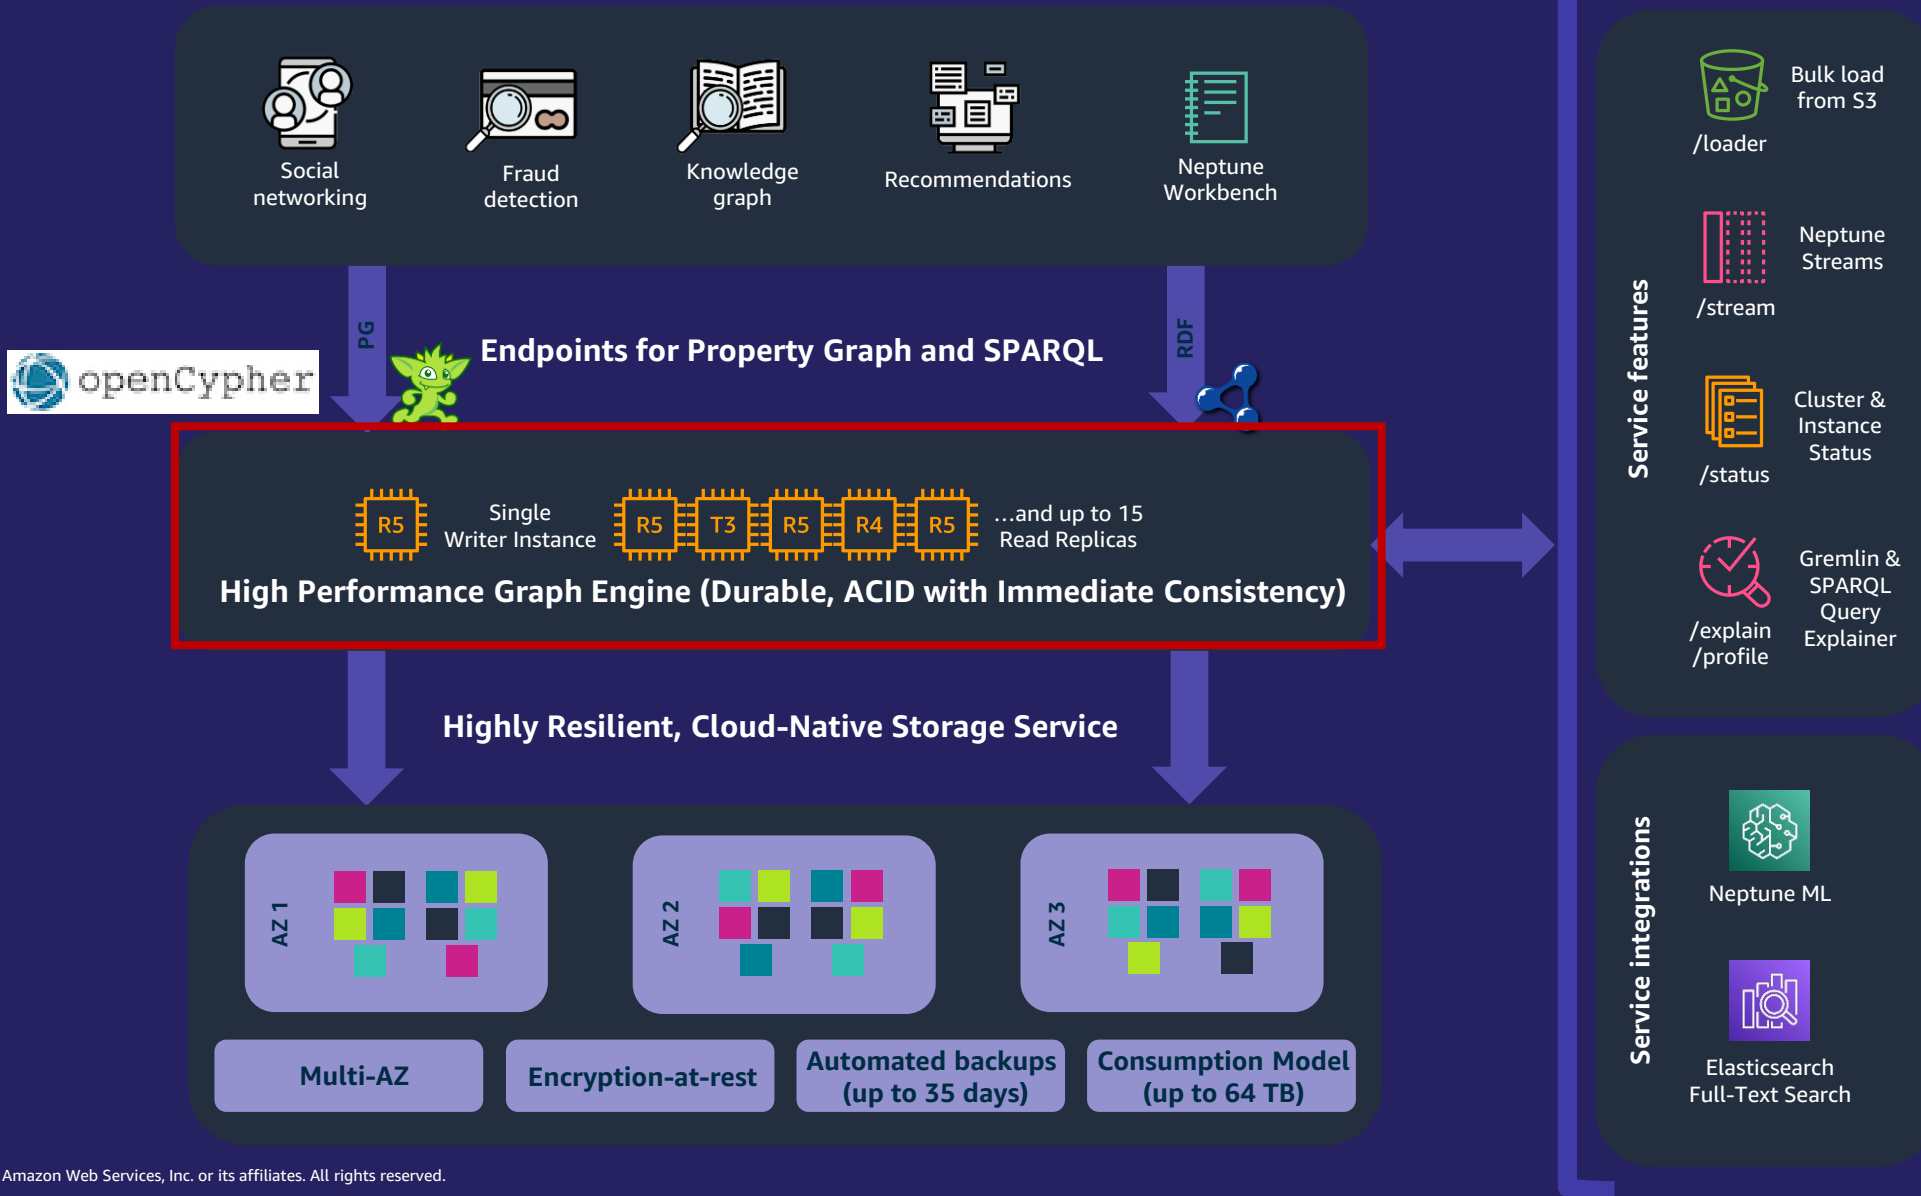
Neptune Customers



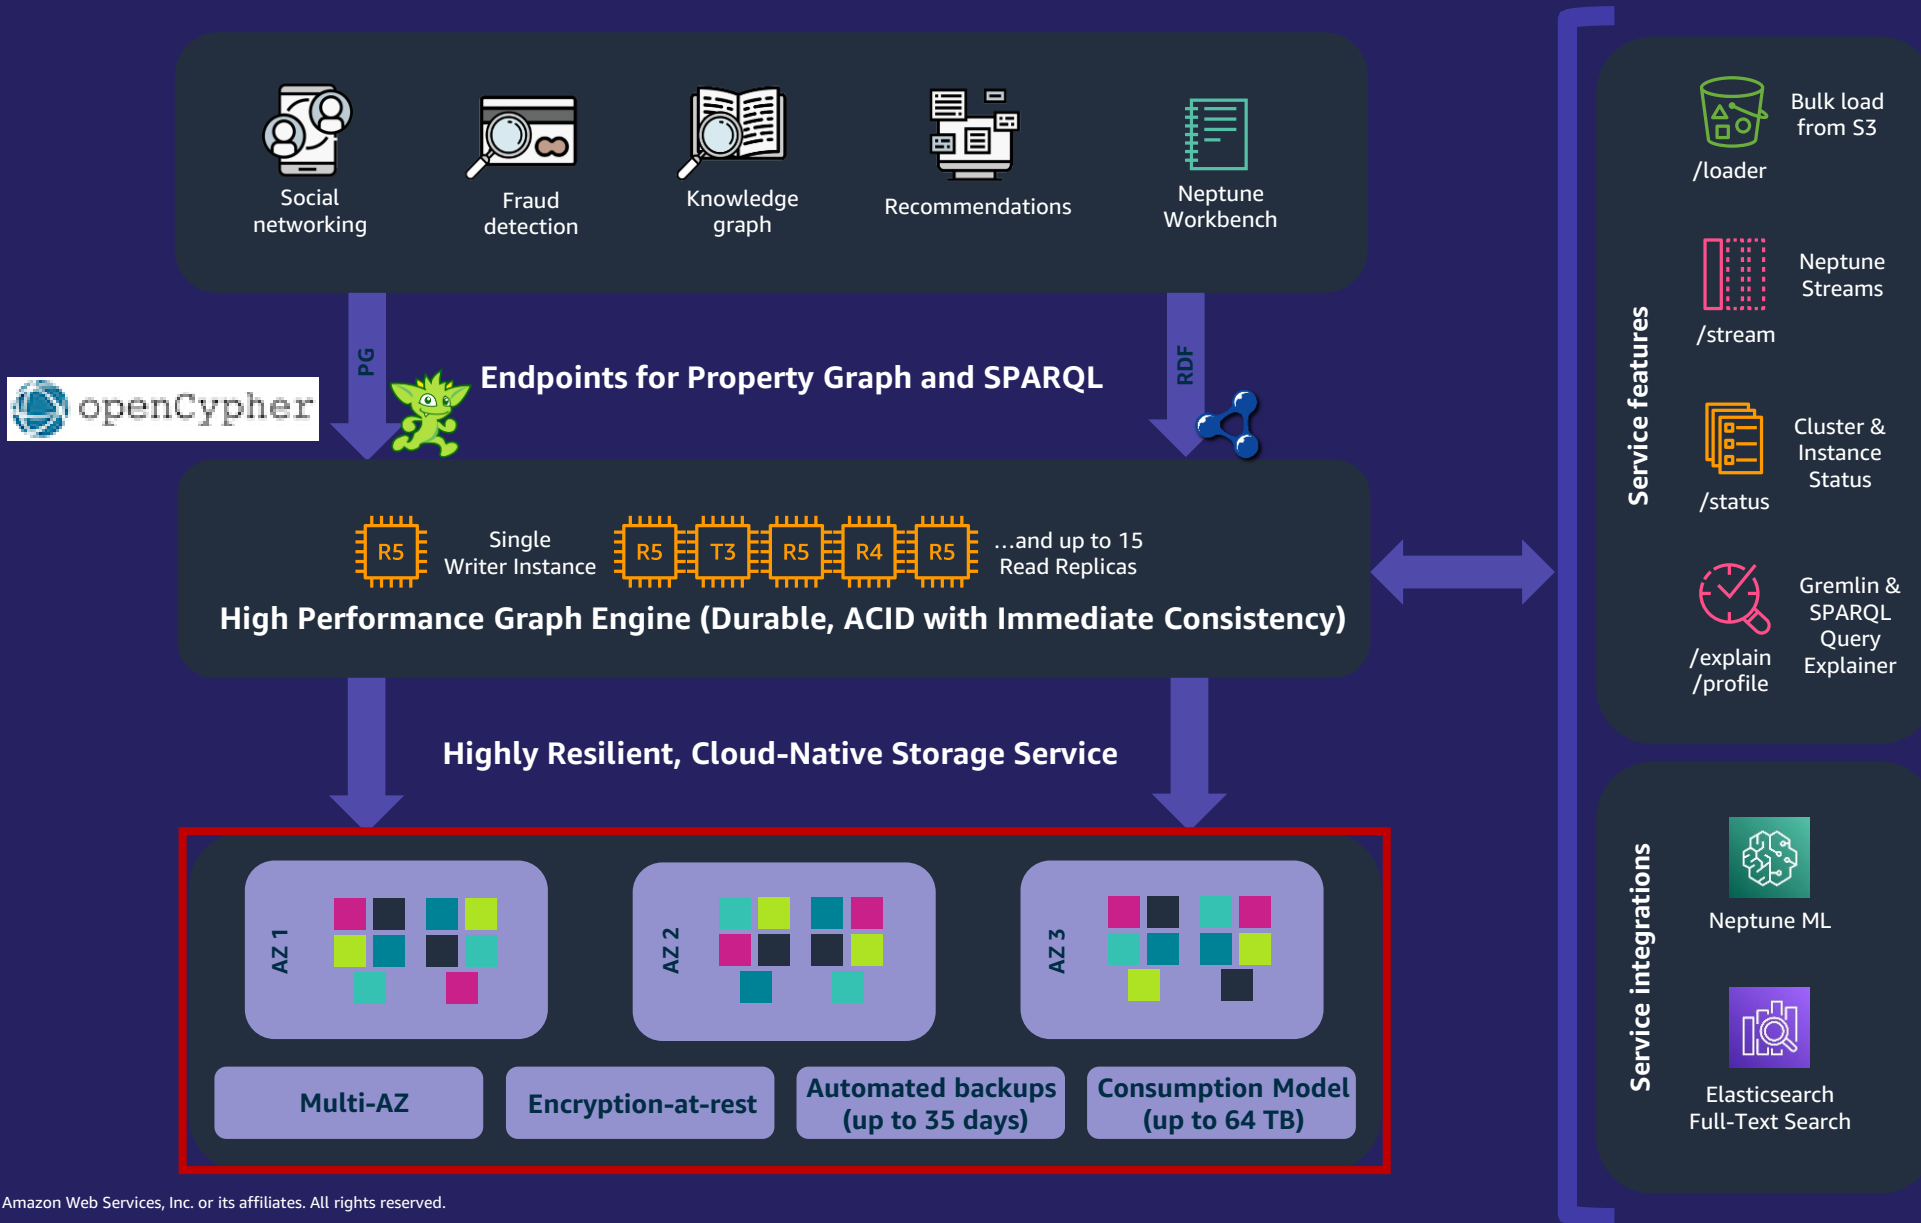
Amazon Neptune High Level Architecture



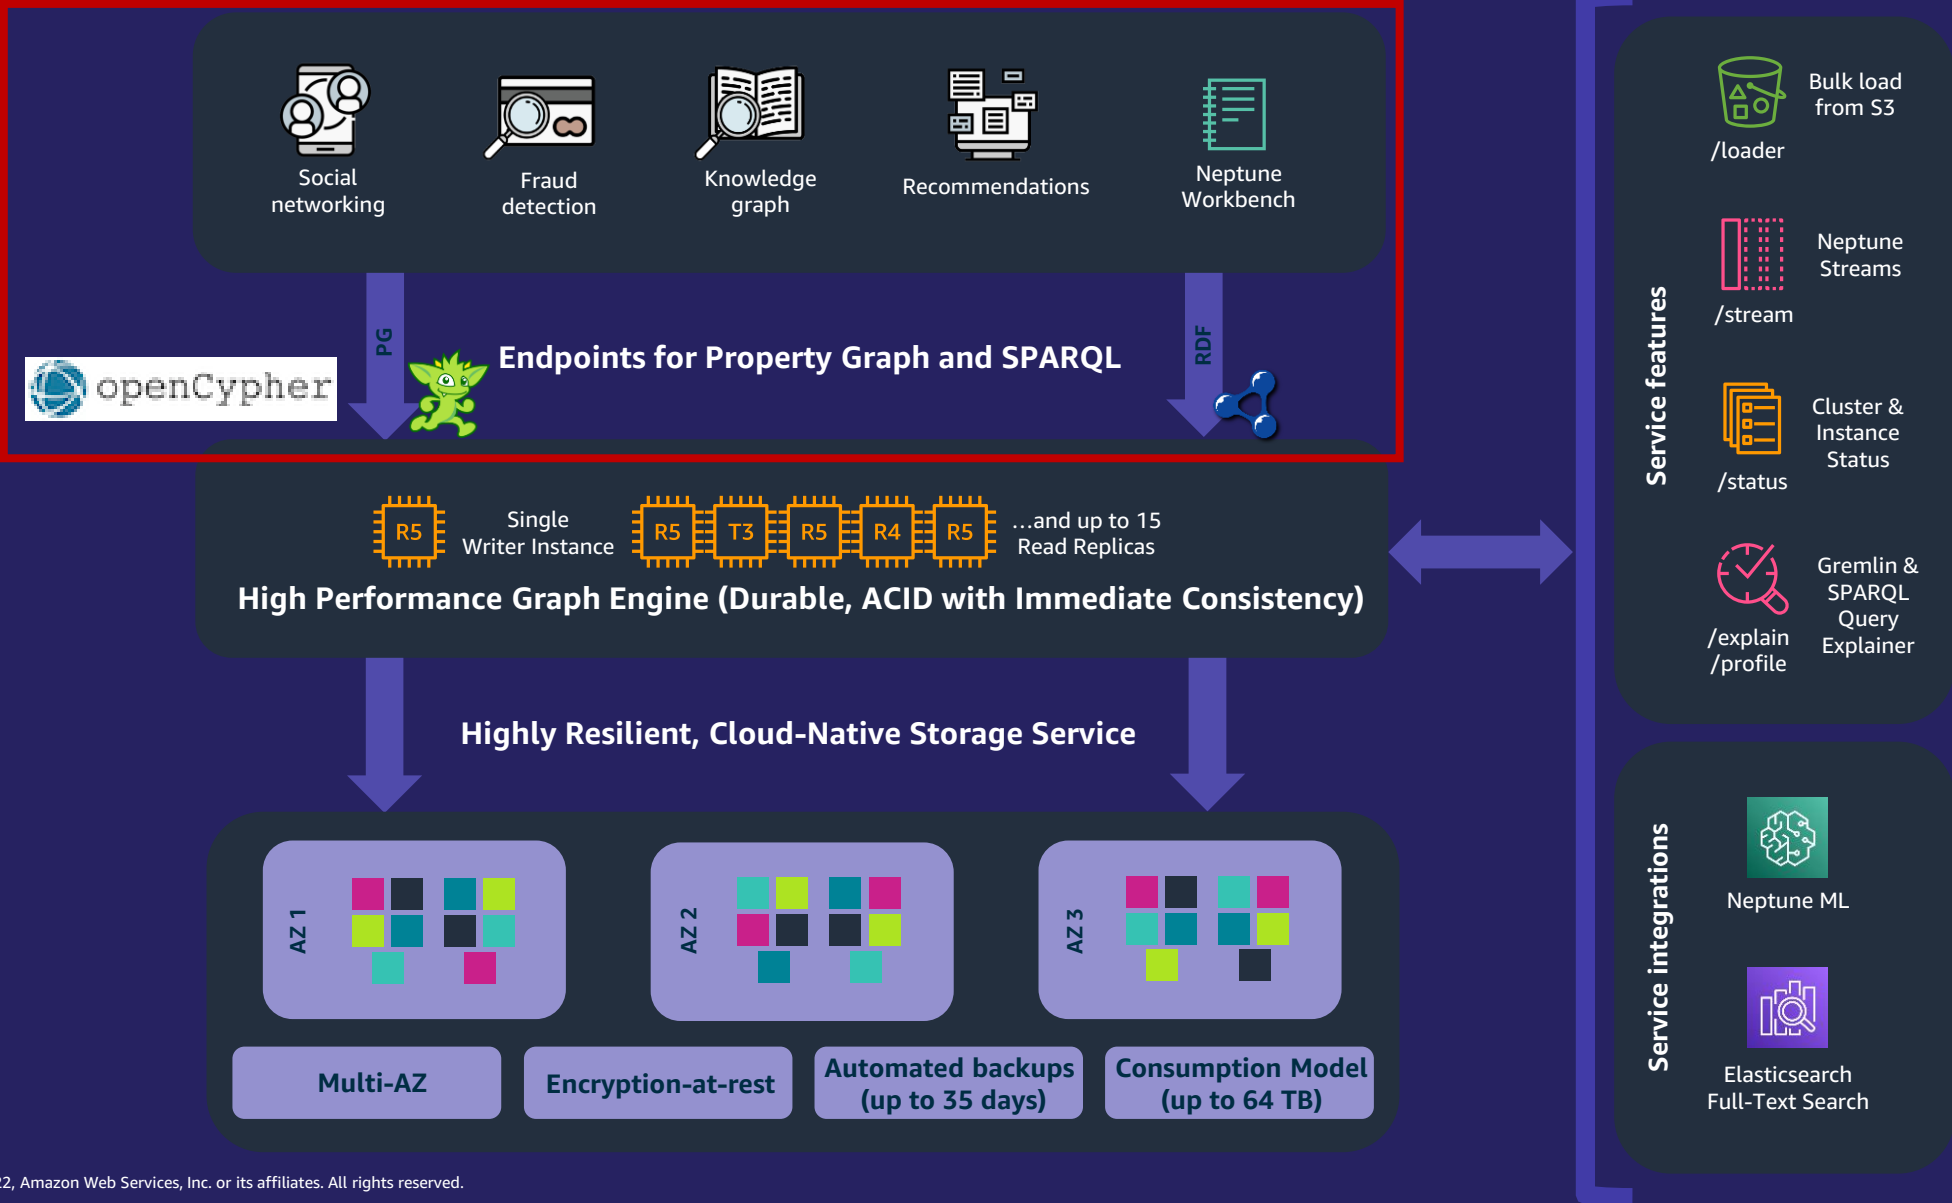
High Performance Graph Engine



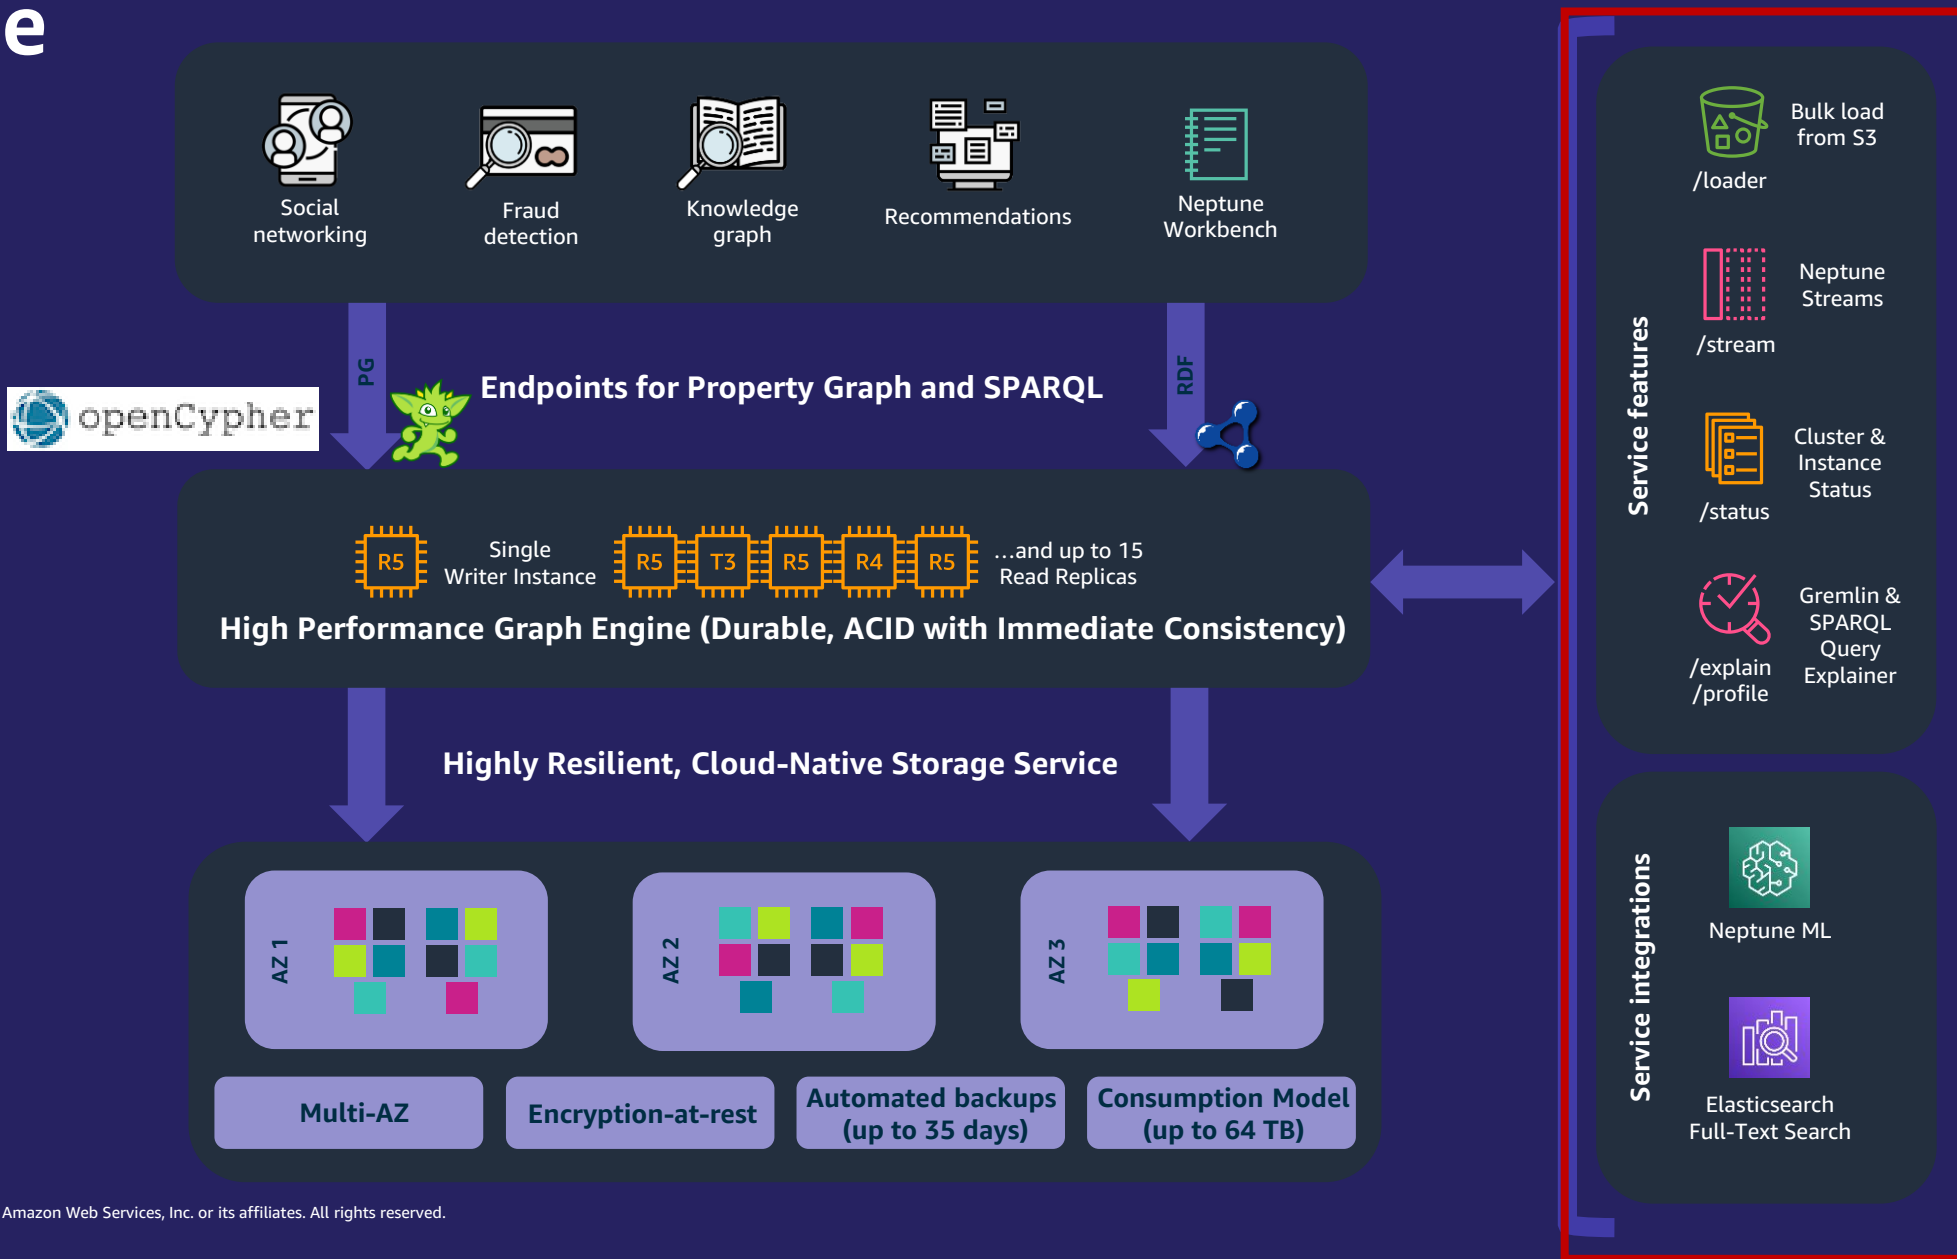
Cloud-Native Storage Service – Multi AZ up to 128TiB



Graph Models and Query Languages



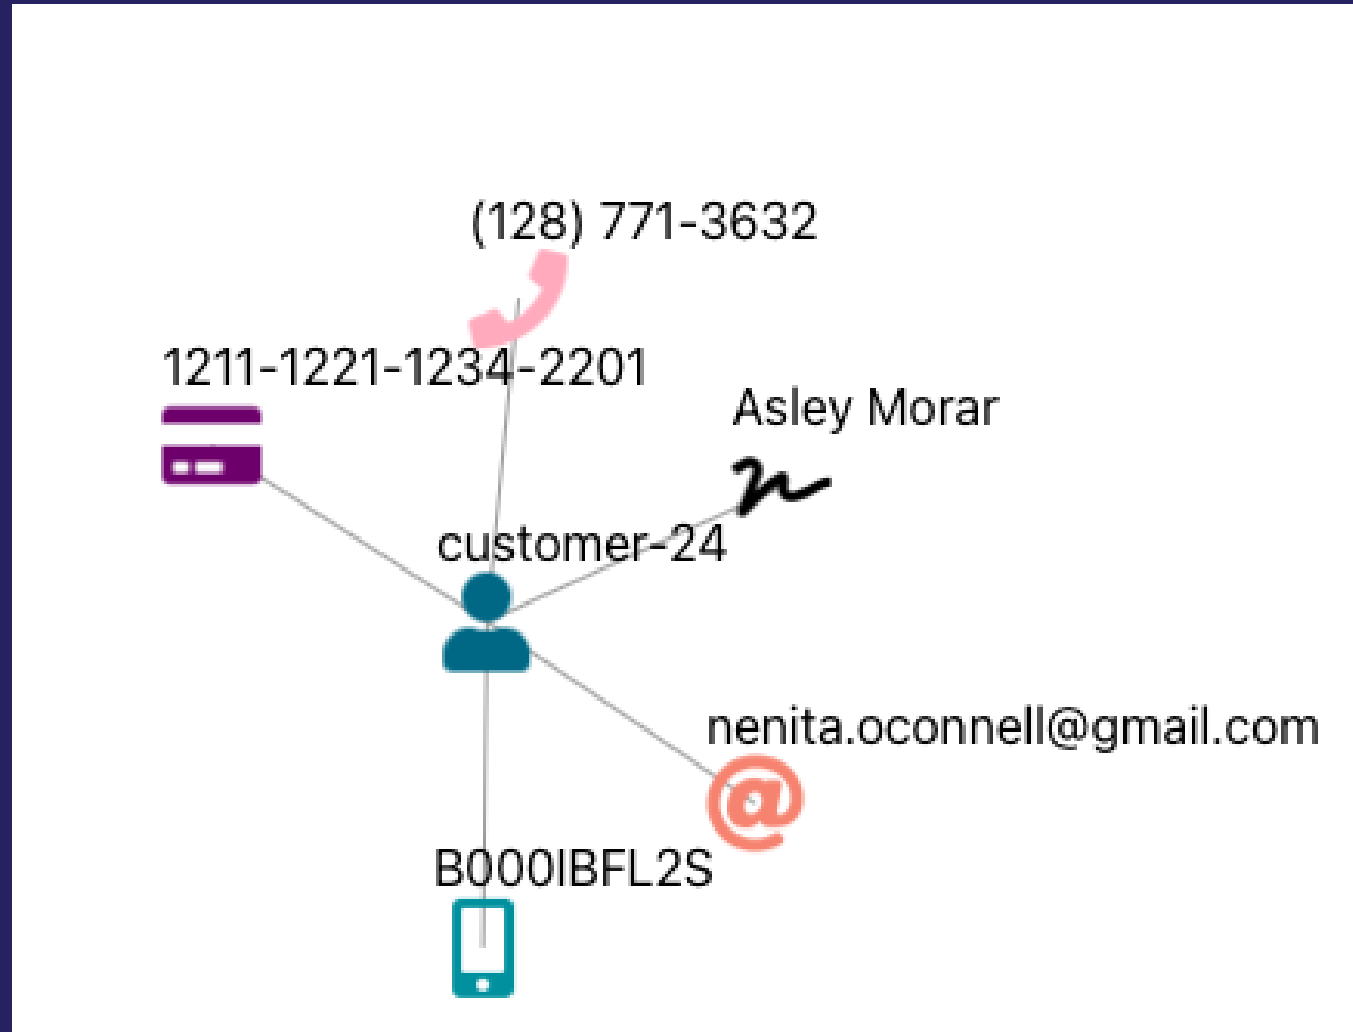
APIs for Load, Streams, ML, Search, Status, Explain, Profile



Demo



Customer 360 Demo

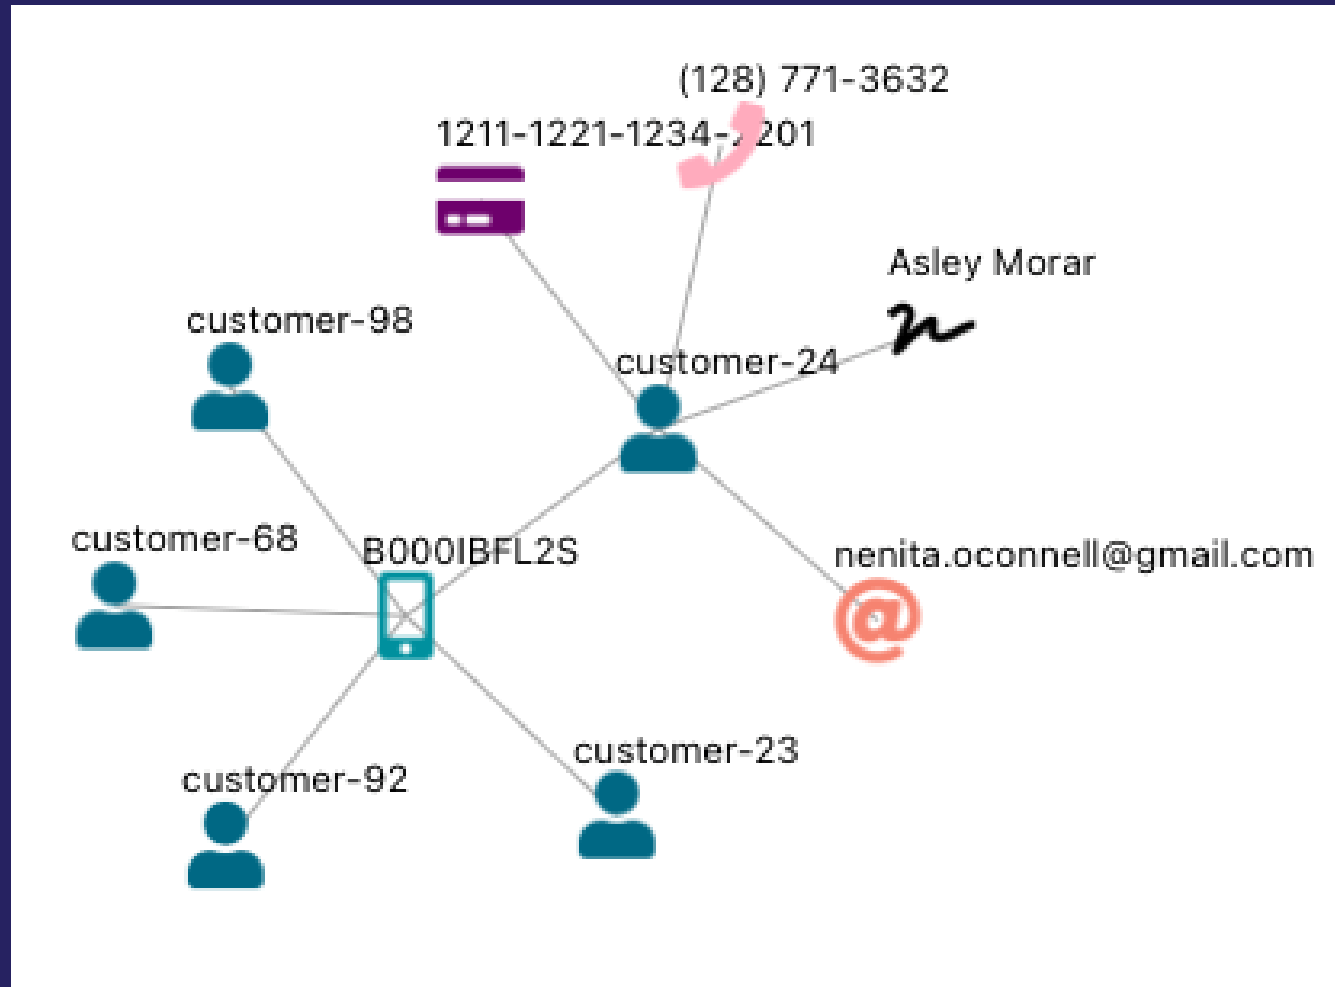


Customer 360 Demo

Connected Identity Results

Number	Customer 1	Customer 2	Shared Attributes		
1	customer-24	customer-92	phone-210	device-105	email-143
2	customer-92	customer-24	phone-210	device-105	email-143
3	customer-28	customer-100	phone-225	name-163	token-184
4	customer-100	customer-28	phone-225	name-163	token-184
5	customer-6	customer-59	phone-201	email-140	token-191
6	customer-59	customer-6	phone-201	email-140	token-191
7	customer-10	customer-91	device-115	token-182	
8	customer-14	customer-30	phone-209	device-110	
9	customer-15	customer-77	device-118	token-184	
10	customer-16	customer-71	device-107	email-132	

Customer 360 Demo

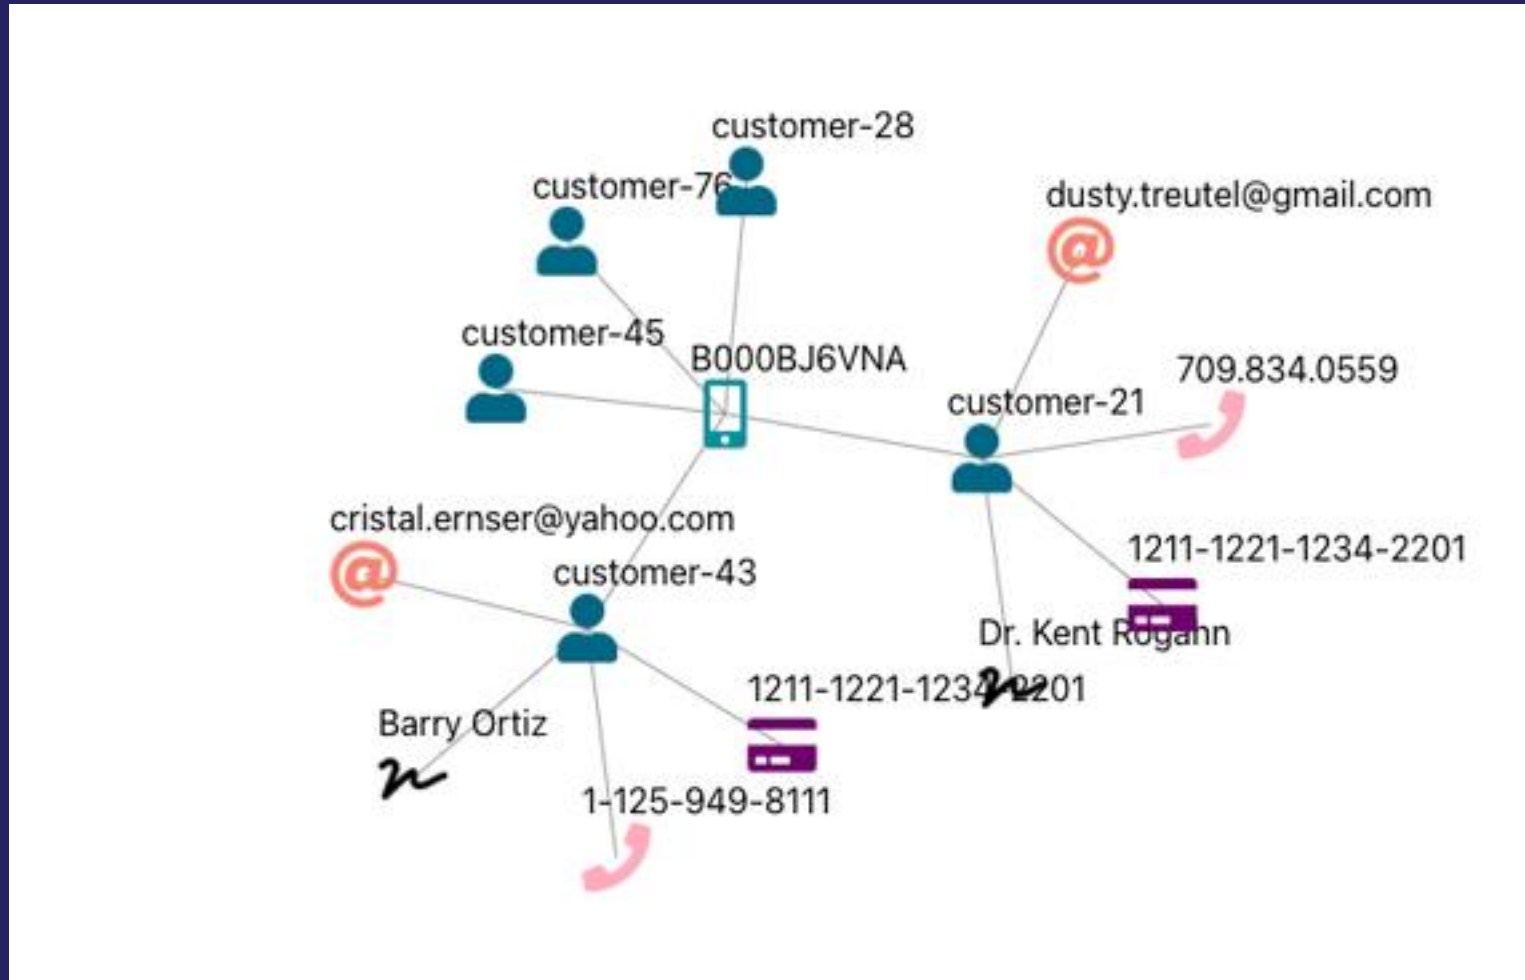


Customer 360 Demo

Connected Identity Results

Number	Customer 1	Customer 2	Shared Attributes		
1	customer-24	customer-92	phone-210	device-105	email-143
2	customer-92	customer-24	phone-210	device-105	email-143
3	customer-28	customer-100	phone-225	name-163	token-184
4	customer-100	customer-28	phone-225	name-163	token-184
5	customer-6	customer-59	phone-201	email-140	token-191
6	customer-59	customer-6	phone-201	email-140	token-191
7	customer-10	customer-91	device-115	token-182	
8	customer-14	customer-30	phone-209	device-110	
9	customer-15	customer-77	device-118	token-184	
10	customer-16	customer-71	device-107	email-132	

Customer 360 Demo - Householding



Resources for Next Steps

- Neptune Best Practices, Training, and Developer Resources
<https://aws.amazon.com/neptune/developer-resources/>
- Neptune Reference Architectures
<https://github.com/aws-samples/aws-dbs-refarch-graph>
- Neptune Graph Notebook & Sample Applications
<https://github.com/aws/graph-notebook/>

Summary

- AWS customers are building Customer 360 Graphs
- The benefits of a Graph Database and Amazon Neptune
- Customer examples and how to get started



Thank you!

Justin Thomas
Principal Amazon Neptune Specialist