Reducing the costs

of your openCypher application

Dave Bechberger

Principal Graph Architect
Amazon Web Services

Graphs are awesome!

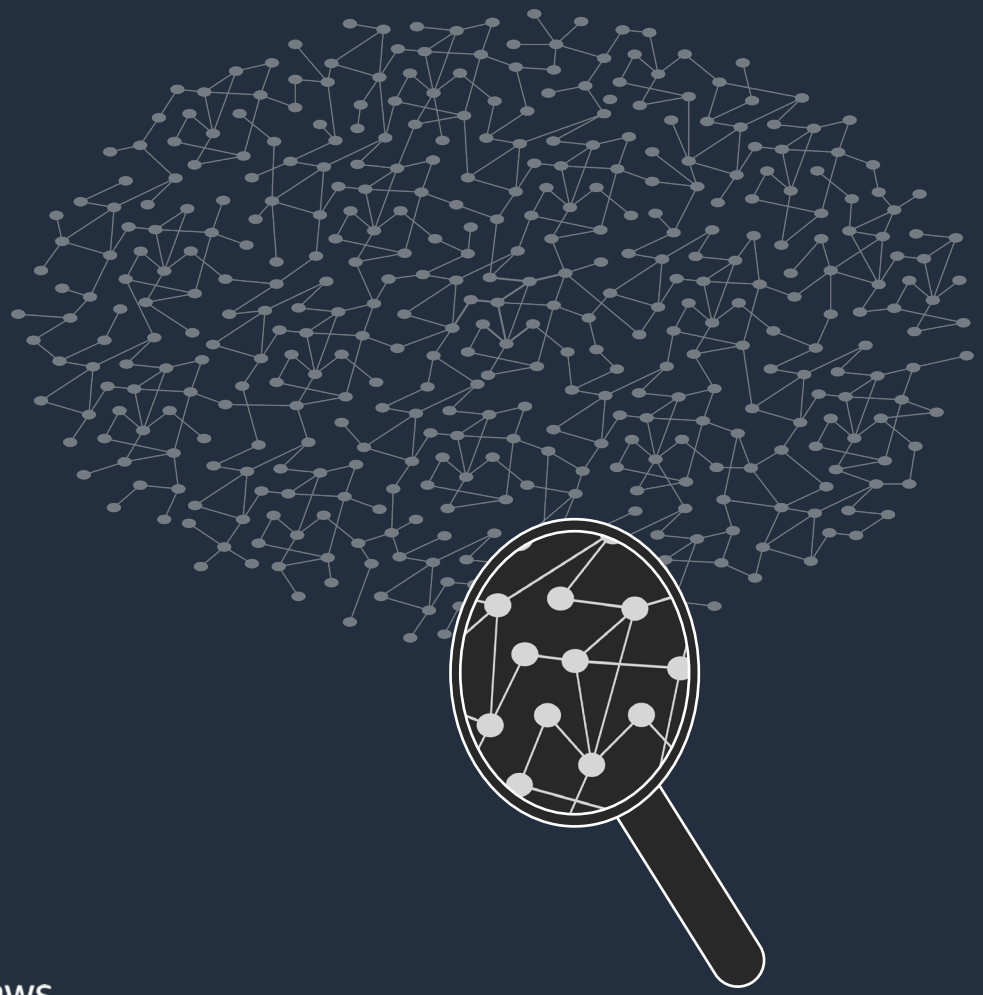

1. Model data based on relationships

2. Applications explore connections and patterns in connected data

3. Processing graphs is hard due to random data access

4. Generalized graph operations require purpose-built processing

Amazon Neptune

(Now Serverless and Global too!)

FULLY MANAGED, PURPOSE-BUILT GRAPH DATABASE IN THE CLOUD

Cost-effective

No hardware management

Instant provisioning

Security and compliance

Serverless

Global

- Optimized to **store and map billions of relationships**
- Enables **real-time navigation of connections with millisecond query** response time
- Supports **open standard query languages** openCypher, Gremlin, and SPARQL

aws

Why migrate applications to Neptune

- Remove complexity with a fully managed service
- Cost optimization, pay as you go, no upfront license
- Support for common open source standards
- Compliant with 20+ standards including SOC, FedRAMP, HIPAA

Customers told us they loved graphs, but they want …

Optimized Developer Experience

openCypher

```
(db:Neptune)<-[:listens_to]-(:Customers)
```

Cost Reduction/Optimization

Amazon Neptune (Serverless)

What is openCypher?

Declarative query language

Focuses on what, not on how

CLAUSE-based language

Uses keywords: MATCH, WHERE, CREATE, MERGE

Pipelined language

Results of each statement are input into the next

What are the benefits of openCypher?

Flexible

Load once, use either query language (support for HTTPS and client drivers)

(Neptune)-[:LOVES]->

(openCypher)

Familiar

Provides a SQL-inspired query syntax purpose-built for property graph data

Fast

Runs on an engine that efficiently scales DB resources such as CPU cores, memory, and I/O

Amazon Neptune Serverless

The first serverless graph database that automatically scales database capacity up or down to optimize cost and performance.

Amazon Neptune Serverless

Scale Instantly

- Instantly scale capacity in a fraction of a second to meet workload demands.

Optimize performance for demanding workloads

- Scales capacity in fine-grained increments. Eliminate the complexity of configuring capacity for unpredictable or variable workloads.

Save up to 90% on database costs

- Reduce costs by up to 90% compared to provisioning for maximum database capacity.

How is Serverless capacity managed?

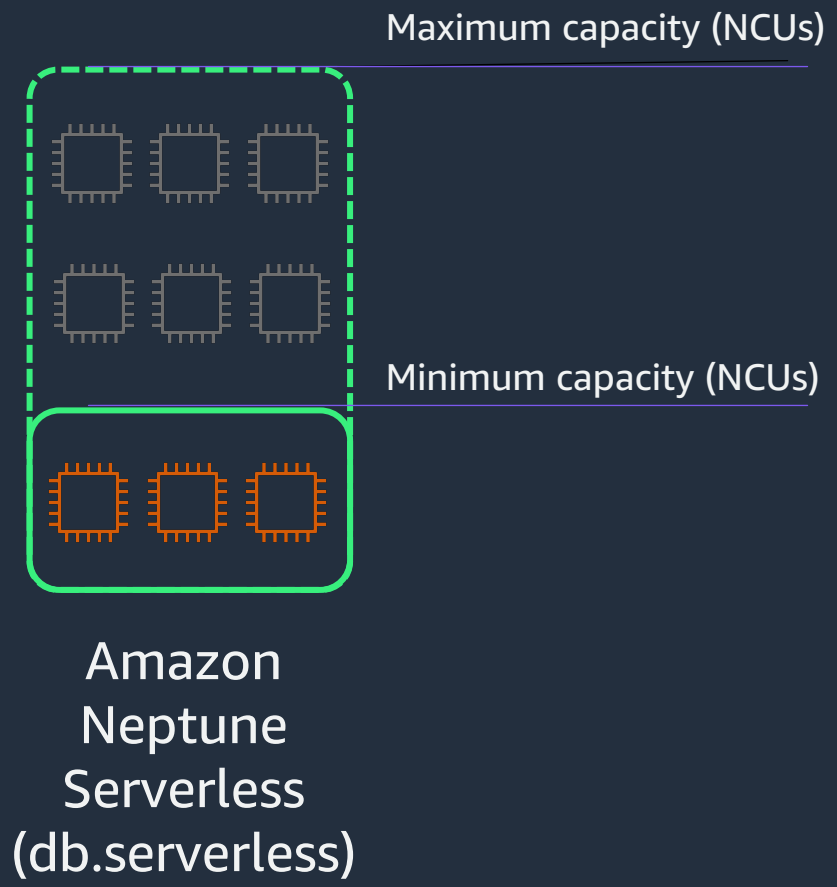

- A Neptune Capacity Unit (NCU) is the measure of scaling
  - 1 NCU = 2GB RAM of capacity and proportionate CPU and network bandwidth
- User specified min. and max. (1-128 NCU)
  - Minimum capacity determines the starting capacity of the instance
  - Maximum capacity is a budget control

Neptune Serverless customer scaling behavior

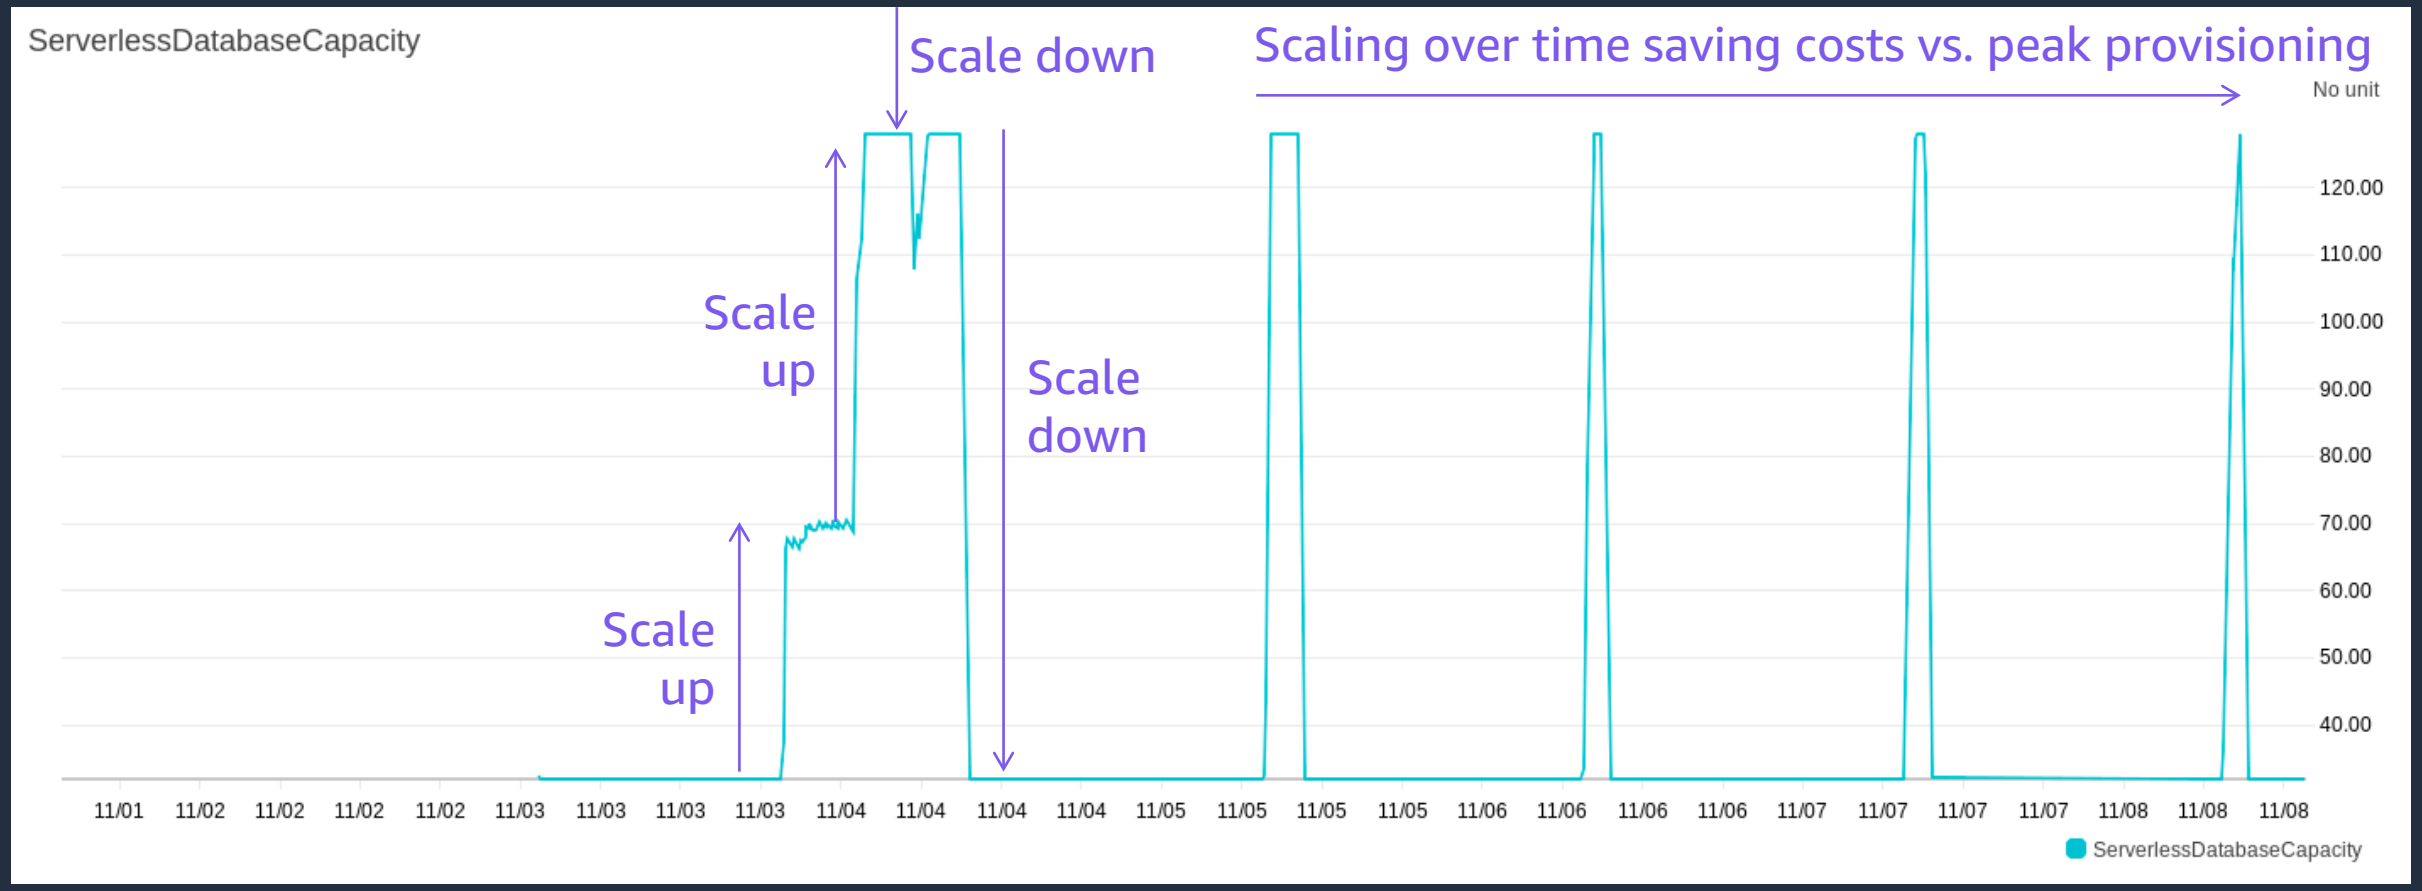

How do we know when and how to scale?

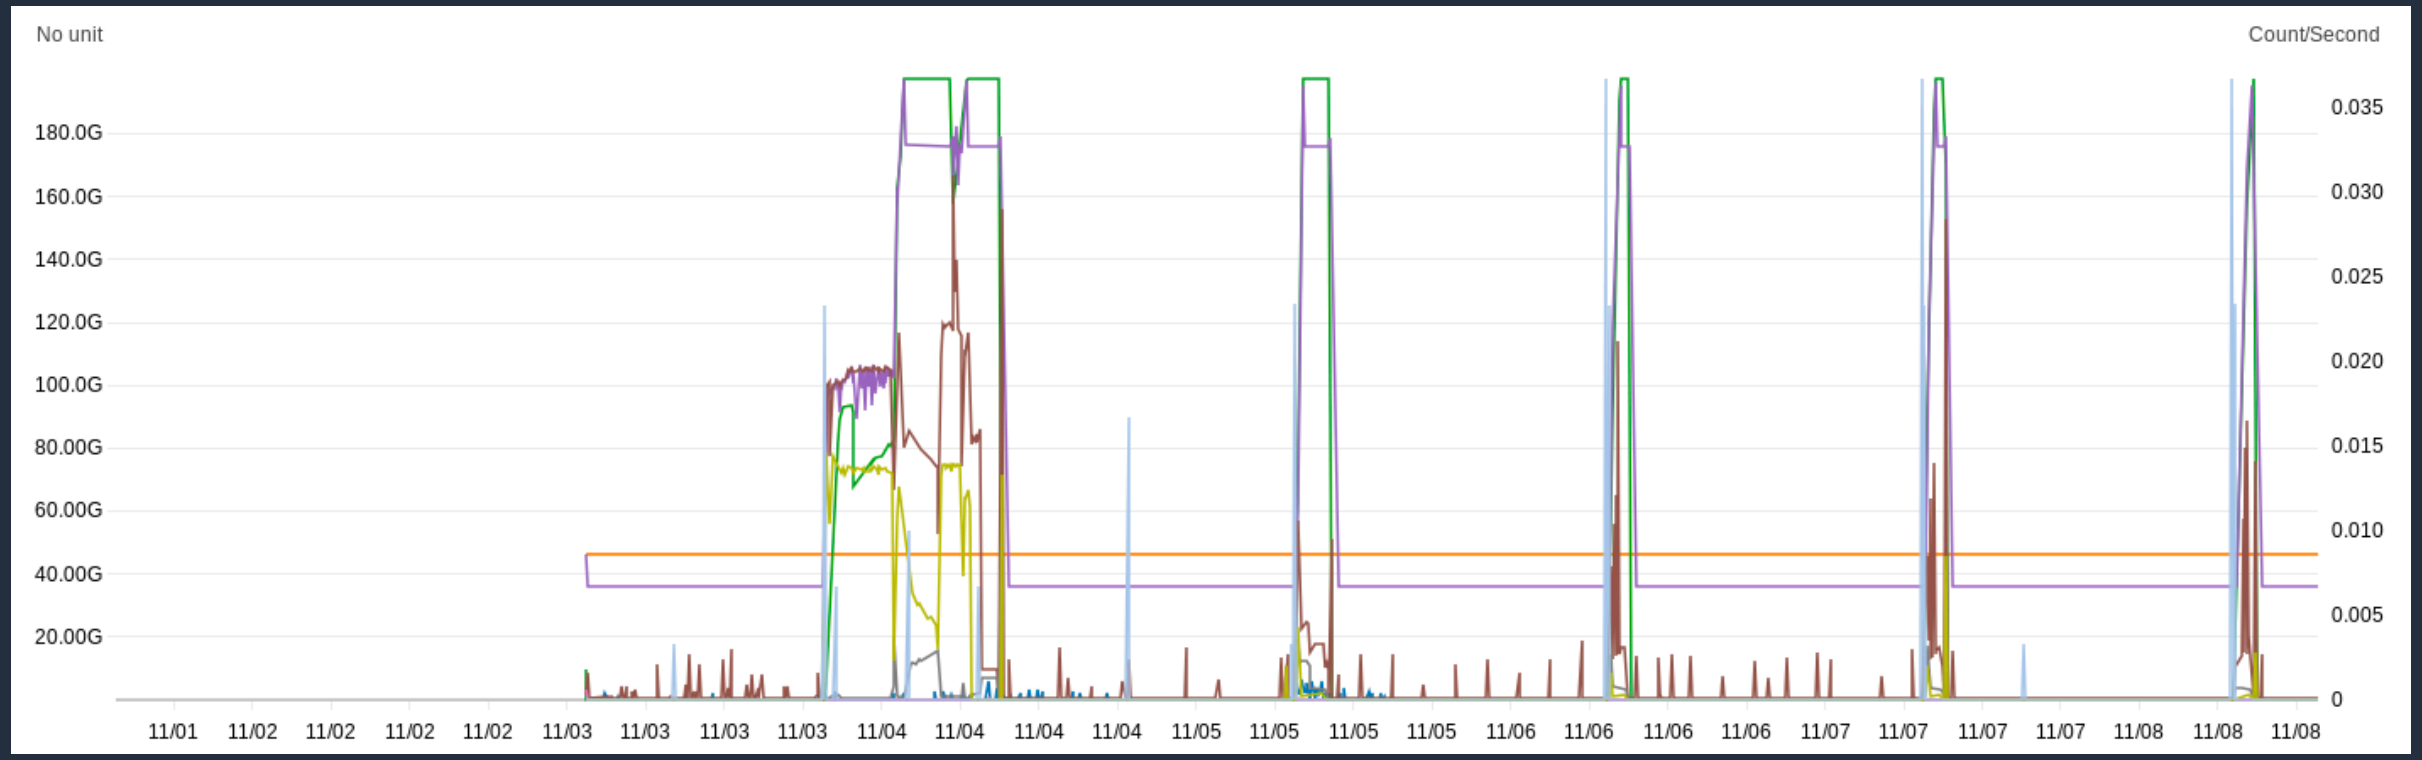

Colored lines represent different decision making processes that are monitoring memory, CPU, network utilization, min/max NCU.

Demo: openCypher workloads on a Serverless Graph Database

Recap and Resources

1. Graphs are awesome. Thousands of customers use Neptune for use cases like knowledge graphs, identity graphs, fraud, and security.
2. Neptune Serverless automatically scales so you don’t have to.
3. openCypher simplifies new application development and migration of existing applications

Thank you!

Dave Bechberger

dbechbe@amazon.com

bechbd