



What's New in Amazon Neptune

Fast, fully managed graph database service


Navtanay Sinha
Sr. Product Manager, Neptune

Why graph?

Relationships enable new applications. Explore **connections**, **paths** and **patterns** in connected data



Knowledge
Graphs



Fraud Detection



Recommendations



Social
Networking



Life Sciences



Network & IT
Operations

Connected Data Queries

Navigate (variably) connected structure

Filter or compute a result based on *strength*, *weight* or *quality* of relationships

Amazon Neptune

Fully managed graph database

Fast



Query billions of relationships with millisecond latency

Reliable



Six replicas of your data across three AZs with full backup and restore

Easy



Build powerful queries easily with Gremlin and SPARQL

Open



Supports Apache TinkerPop & W3C RDF graph models

Amazon Neptune Availability

- Available 19 AWS Regions including GovCloud.
- Encryption-at-rest with Amazon Key Management Service (KMS)
- Encryption-in-transit with TLS 1.2 client connections
- ISO, HIPAA, SOC, PCI/DSS compliant

What's New in Neptune

13 releases in the last 6 months

New features in Amazon Neptune

SPARQL 1.1 Federated Query	<ul style="list-style-type: none">• Use SPARQL to express queries across diverse data sources
Transaction semantics	<ul style="list-style-type: none">• Formalized semantics to help you avoid data anomalies
Gremlin/SPARQL Explain	<ul style="list-style-type: none">• Gain insights into the query plan and evaluation order
Gremlin sessions	<ul style="list-style-type: none">• Queries run during the session are committed as part of a single transaction
Database cloning	<ul style="list-style-type: none">• Create multiple clones of a DB cluster using copy-on-write semantics
Neptune Streams	<ul style="list-style-type: none">• Complete sequence of change-log entries: record every change made to graph
Elasticsearch integration	<ul style="list-style-type: none">• Full-text search using Elasticsearch with graph data in Neptune
Neptune Workbench	<ul style="list-style-type: none">• In-console notebook experience to query your graph
Low cost T3 instances	<ul style="list-style-type: none">• Next generation burstable general-purpose instance as low as 10 cents/Hr
Delete Protection	<ul style="list-style-type: none">• Configure a cluster with deletion protection to prevent accidental deletes
Start/Stop cluster	<ul style="list-style-type: none">• Stop databases when it is not required to be running all of the time
Cross region snapshot copy	<ul style="list-style-type: none">• Copy snapshots across regions for testing and disaster recovery
Enforce SSL connections	<ul style="list-style-type: none">• Enforce SSL connections with option to disable SSL in regions where both are supported.
Simplified Console Management	<ul style="list-style-type: none">• Simplified console experience to manage cluster, instances and their properties.
Alternative query engine	<ul style="list-style-type: none">• Better query performance with DFE.

SPARQL 1.1 Federated Query

- Use SPARQL to express queries across diverse data sources
- Execute a portion of the query across different SPARQL endpoints within the VPC and combine results.
- Distribute data across multiple Neptune clusters and use a single query to access that data across those clusters.
- Use federation to combine data in Neptune with data from external SPARQL endpoints.

Transaction Semantics

- Formalized semantics to help you avoid data anomalies
 - SPARQL and Gremlin do not define transaction semantics for concurrent query processing.
 - Avoid data anomalies and provide well defined transaction guarantees
- Separate isolation read queries and queries with mutation
 - Snapshot isolation for read queries
 - Read committed transactions

Gremlin/SPARQL Explain

- Self-service tool to understand execution approach taken by Neptune
 - Provides information about the logical structure of query execution plans
 - Use this information to identify potential evaluation and execution bottlenecks.
 - Use query hints to improve your query execution plans.

Gremlin sessions

- Queries run during the session are committed as part of a single transaction
 - Queries committed together when session closed
 - If queries fail, entire session can be rolled back
 - Use CLI or APIs to create sessions

```
gremlin> :remote connect tinkertop.server conf/neptune-remote.yaml session
...
...
gremlin> :remote close
```

Database cloning


- Quickly and cost-effectively create clones of all your databases.
 - Clones require only minimal additional space when created.
 - Uses a *copy-on-write* protocol: Data is copied at the time that it changes
 - Make multiple clones from the same DB cluster.
- Used In:
 - Test and assess schema or parameter group changes.
 - Workload-intensive operations on clones
 - Dev-test on clones instead of production DB

Neptune Streams

- Generate a complete sequence of change-log entries that record every change made to graph
- Notify processes as changes occur in your graph.
- Use streams to maintain a current version of your graph in a different region or service such as the Amazon Elasticsearch Service, Amazon ElastiCache, or Amazon S3.

Elasticsearch integration

- Full-text search queries using Elasticsearch
- Use match, fuzzy, prefix, query_string options
 - Easily run full-text search on graph data
 - Access Elasticsearch's built-in text indexing and query capabilities from Neptune.
 - Support for both SPARQL and Gremlin data in Elasticsearch.



Neptune Workbench

- Query and Visualize your graph
- In-console experience
 - Quickly query Neptune databases with Jupyter notebooks.
 - Create notebooks from the AWS Console for Neptune.
 - Invoke the bulk loader, run query plans and profile queries.
 - Samples to help get started with Neptune quickly.

Low cost T3 instances

- Launch a cluster with T3 instances
 - Development or test use cases
 - Next generation burstable general-purpose instance type
 - Provides baseline level CPU performance with the ability to burst CPU usage.
 - T3 instances accumulate CPU credits when operating below baseline threshold.
 - Each credit provides an opportunity to burst with the performance of a full CPU core for one minute when needed.
 - Neptune supports T3.medium instance type, configured for Unlimited mode.
 - Instances can burst beyond the baseline over a 24-hour window

Delete Protection

- Enable delete protection for your Amazon Neptune database clusters.
- When configured, the database cannot be deleted by any user.
- Enabled by default for databases created through the AWS Management Console.
- Use the AWS Management Console, AWS CLI or Neptune Management API.

Start/Stop cluster

- Stop a database when not in use, lowers Neptune cost where DB isn't required to run continuously.
- Stop a database cluster for up to 7 days at a time
- Stop/start cluster from the AWS Management Console, AWS CLI or Neptune Management API.

Cross region snapshot copy

- Copy a database snapshot from one region to another for DR
- Simplify and streamline data manipulation operations associated with building and running global graph applications
- Initiate a copy from the AWS Management Console, AWS CLI or Neptune Management API.

Enforce SSL connections

- SSL connections are now enforced in all regions
 - When new cluster is created, the newly introduced *neptune_enforce_ssl* parameter is enabled.
 - Databases that have this parameter enabled will only accept SSL connections.
 - Modifiable from the AWS Management Console, AWS CLI or Neptune Management API.

Simplified Console Management

- Updates to AWS Management Console experience
 - Single page database creation
 - Simplified cluster and instance management using a single list view
 - Navigate through cluster properties easily

Alternative query engine (DFE)

- Enabled using lab mode parameter for development or test purposes
- Faster query performance
- Uses DB instance resources more efficiently
- Supports a subset of SPARQL and Gremlin query constructs.

Resources

Neptune releases

AWS What's New Feed on Neptune

<https://aws.amazon.com/new/?whats-new-content-all.q=neptune>

Neptune Engine Releases

<https://docs.aws.amazon.com/neptune/latest/userguide/engine-releases.html>

Neptune User Guide

<https://docs.aws.amazon.com/neptune/latest/userguide/intro.html>

Thank you!