

Case Studies

- Neo4J embedded graph context required
- DST twice a year
- Hidden Concurrency issues using Maps and Lists
- Needed a fully automated Kubernetes install for CI/CD
- Deployment Hierarchy not reflecting dependencies - needed quick deploy sequencing fix
- Cost Analysis
 - Cost Savings: AWS SPOT VMs for CI/CD instances
 - Cost Savings: AWS SPOT VMs for Kubernetes Cluster VMs
 - Cost Savings: EFS Infrequent Access

Neo4J embedded graph context required

client not ready for separate JVMs - collocated speed is higher

DST twice a year

Naive 4 level join slowing JPA DB queries

I designed it - I fixed it using a set of default lazy load queries with optional fetch joins

Hidden Concurrency issues using Maps and Lists

use CopyOnWriteArrayList and ConcurrentHashMap when replacing older ForkJoin code with

```
list.parallelStream().forEach(consumer)
```

```
aList = Stream.of(String[]).map(elem -> new String(elem)).collect(Collectors.toList())
```

Needed a fully automated Kubernetes install for CI/CD

Deployment Hierarchy not reflecting dependencies - needed quick deploy sequencing fix

Short term

sh script

Long term

https://github.com/helm/helm/blob/master/docs/helm/helm_dependency.md

Cost Analysis

Cost Savings: AWS SPOT VMs for CI/CD instances

with retry at 90% off

Cost Savings: AWS SPOT VMs for Kubernetes Cluster VMs

Pick us-east-2 for up to 3-6 months - and R4 memory optimized - not in demand

Cost Savings: EFS Infrequent Access

<https://aws.amazon.com/blogs/aws/new-infrequent-access-storage-class-for-amazon-elastic-file-system-efs/>