CACHE ME IF YOU CAN!
GETTING STARTED WITH AMAZON ELASTICACHE

AWS Charlotte Meetup /
Charlotte Cloud Computing Meetup

Bilal Soylu

October 2013

amazon web services
CHARLOTTE

Agenda

- Hola!
- Housekeeping
- What is this use case
- What is Amazon Elasticache
- What can we do with this
- How much is it going to cost me
- Let’s Brake it Down
  - Concepts, Examples, etc.

Hola! Guten Tag! Bonjour!

- Bilal Soylu
  - CTO Verian Technologies LLC (www.verian.com)
  - Of course, we are looking for peeps ! What kind of question is that!
  - Like to play with AWS stuff
  - Open Source Contributor
  - I really, really learn from my mistakes ;o)
- Blog
  - http://BonCode.blogspot.com
- Contact
  - @BmanClt
  - bilal.soylu[at]gmail.com

Housekeeping

- Meeting Place (We are trying new locations!)
- Meeting Time
- Need Speakers
- Vendor Presentations
- **Speakers please! Now is the time to think about change.**
- Communication
  - More / Less
  - What medium to use

Let’s get started

Cache Me!

Cloud Humor

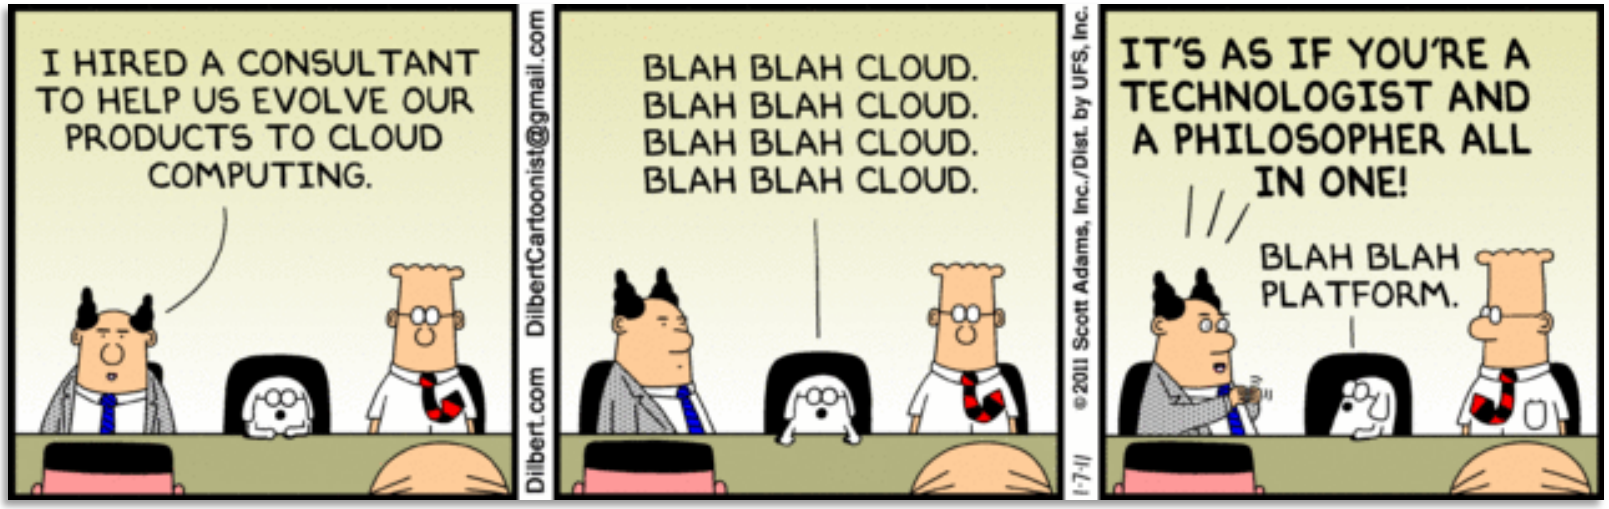
I HIRED A CONSULTANT TO HELP US EVOLVE OUR PRODUCTS TO CLOUD COMPUTING.
BLAH BLAH CLOUD.
BLAH BLAH CLOUD.
BLAH BLAH CLOUD.
BLAH BLAH CLOUD.
IT'S AS IF YOU'RE A TECHNOLOGIST AND A PHILOSOPHER ALL IN ONE!
BLAH BLAH PLATFORM.
Dilbert.com DilbertCartoonist@gmail.com
©2011 Scott Adams, Inc./Dist. by UFS, Inc.
1-7-11

What is the use case

- Use Case Examples
  - Significantly improve latency and throughput for read heavy applications.
    - Social networking, gaming, media sharing
  - Improve throughput for compute-intensive workloads
  - Augment database tier (using Redis)

- Technical Benefits
  - Highly reliable
  - Engine Compatibility (Memcache and Redis)
  - Automatically detect and replace failed nodes
  - STILL MARKED AS BETA
  - *Does not seem to Auto-Scalable

What is AWS Elasticache

"ElastiCache is a web service that makes it easy to deploy, operate, and scale an **in-memory cache in the cloud**. The service **improves** the **performance** of web applications by allowing you to retrieve information from fast, managed, in-memory caches, instead of relying entirely on slower disk-based databases"

Two Flavors

Memcached - a widely adopted memory object caching system. ElastiCache is protocol compliant with Memcached, so popular tools that you use today with existing Memcached environments will work seamlessly with the service.

Redis – a popular open-source in-memory key-value store that supports data structures such as sorted sets and lists. ElastiCache supports Redis master / slave replication which can be used to achieve cross AZ redundancy.

Some other noted features

- Simple to configure
- Low latency, high throughput
- Easy Admin
- Easy Monitoring
- Utility cost model based on underlying instance use
- Engines are compatible with existing interfaces

Basic getting started

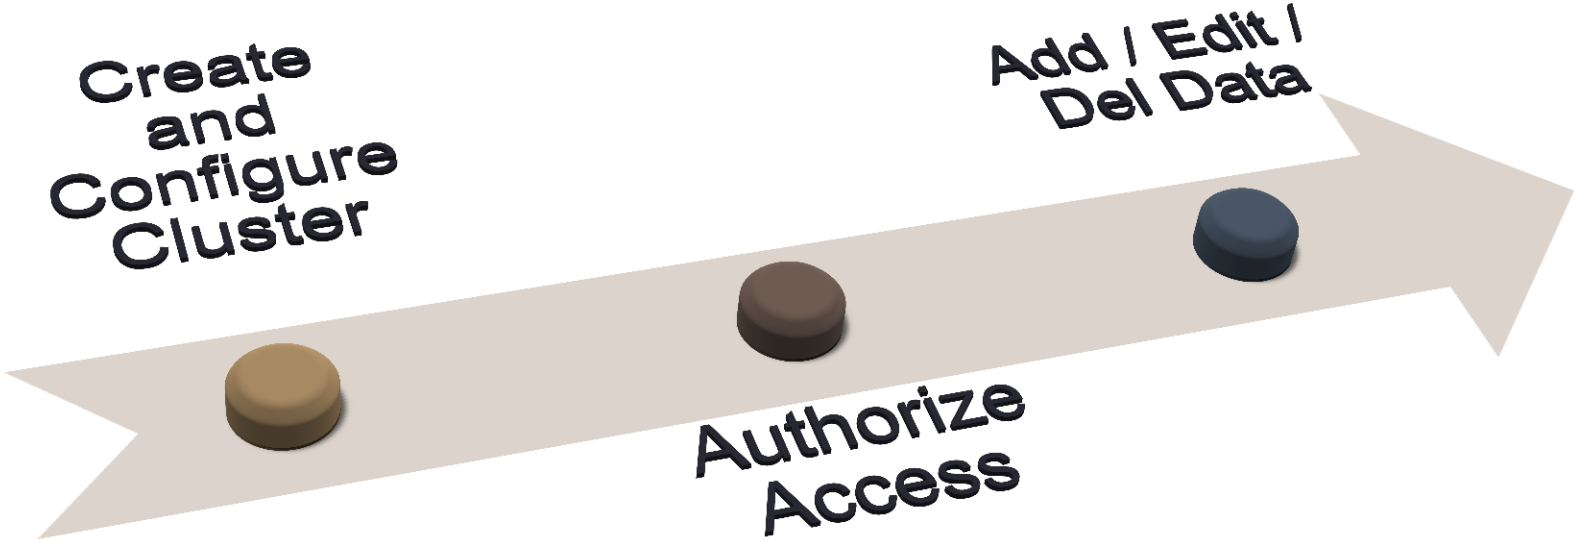

Create Cluster – Step 1

- Use console or API

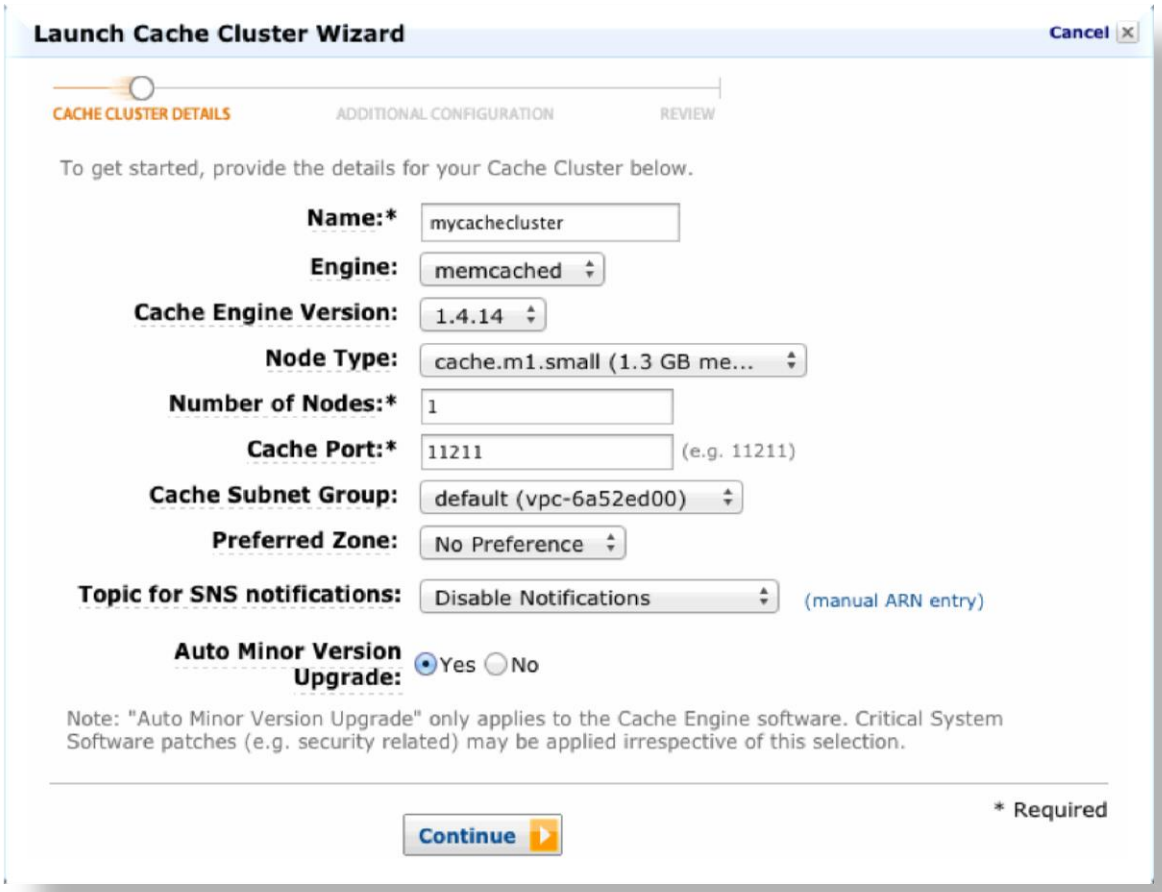

Create Cluster – Step 2

- Additional Configuration

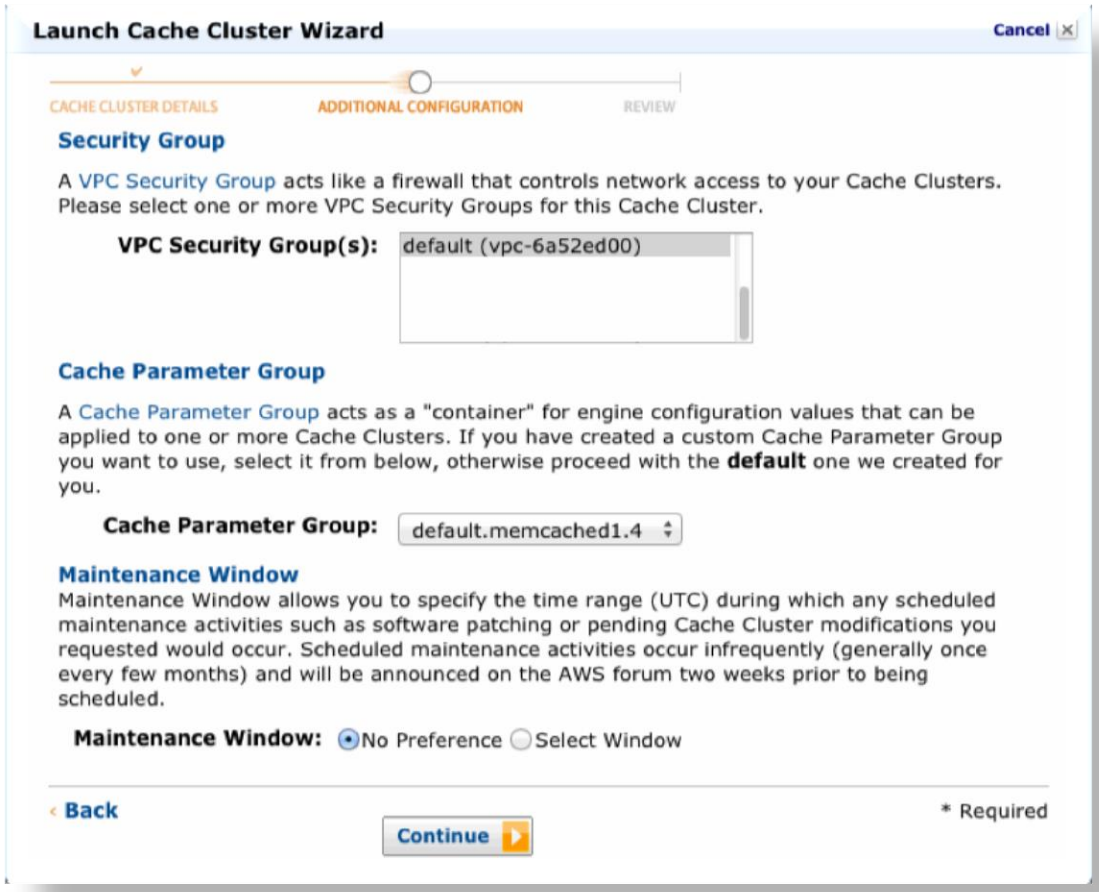

Create Cluster – Step 3

- Launch

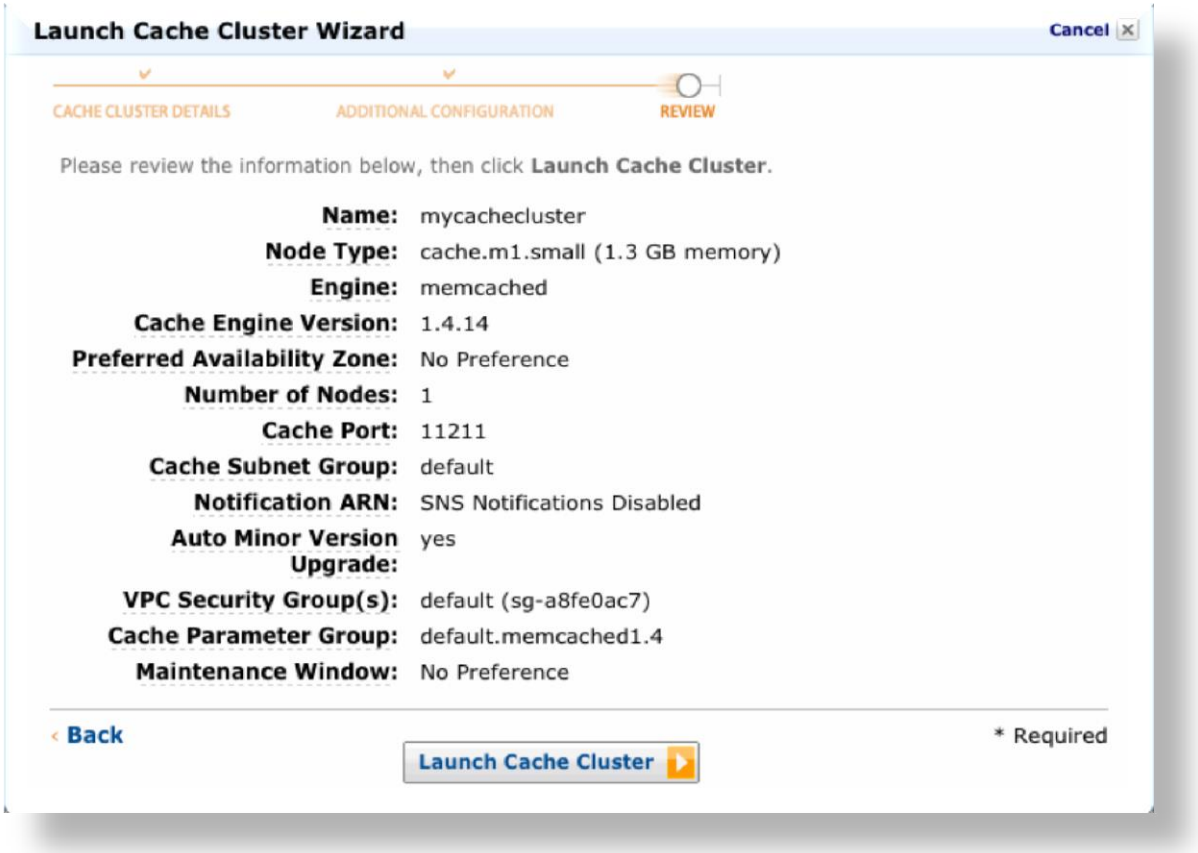

Authorize Access – EC2 Tools

1. Create EC2 security Group or map existing group to Elasticache Group
2. On EC2 navigation pane, under Network & Security, click Security Groups.
3. In the list of security groups, click the security group for your VPC. If you are a new ElastiCache user, this security group will be named default. Consider creating a security grop
4. In the Details pane, click Inbound, and do the following:
    1. In the Create a new rule box, choose Custom TCP rule.
    2. In the Port range field, enter the port number for your cache cluster node. This must be the same number that you specified when you launched the cache cluster.
    3. In the Source field, leave the value at its default (0.0.0.0/0) so that any Amazon EC2 instance that you launch within your VPC can connect to your ElastiCache nodes.
    4. Click Add Rule, and then click Apply Rule Changes.

Connect to Cache via Endpoint (Node)

- Even in multi-node cluster you connect to a node unless you have special client.

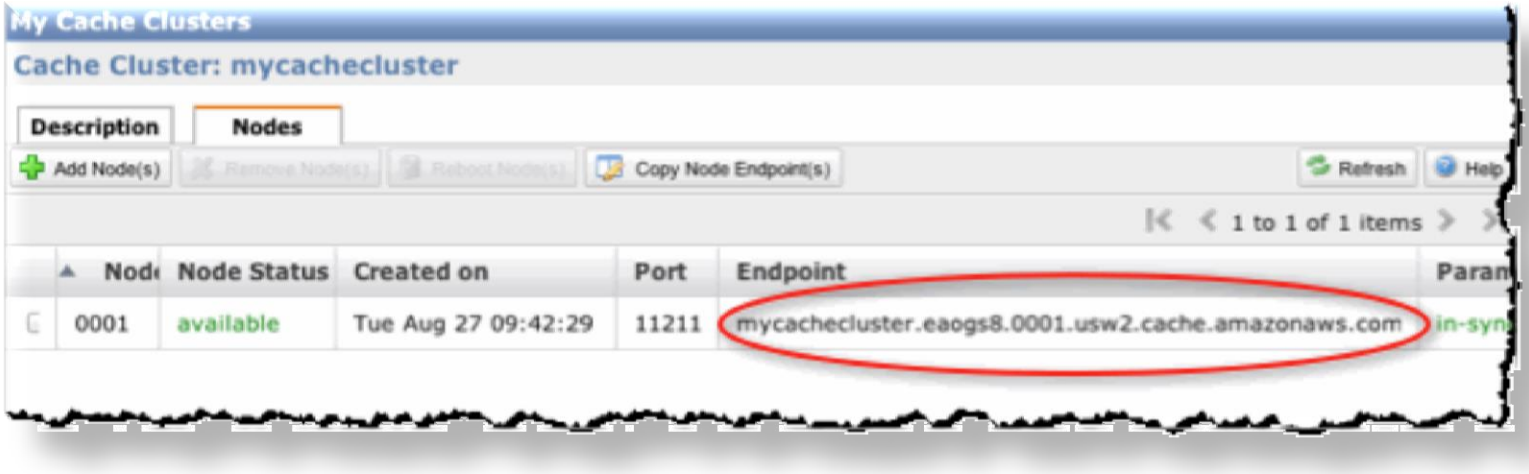

Connect to Config Endpoint for Auto discovery

- Requires special client.
- May require re-code of app.

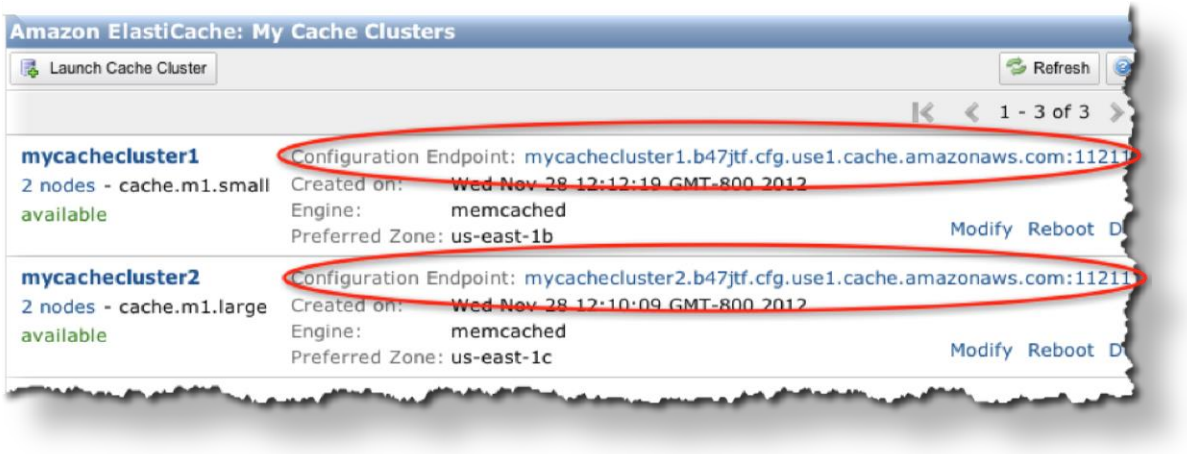

A graphical view of auto-discovery

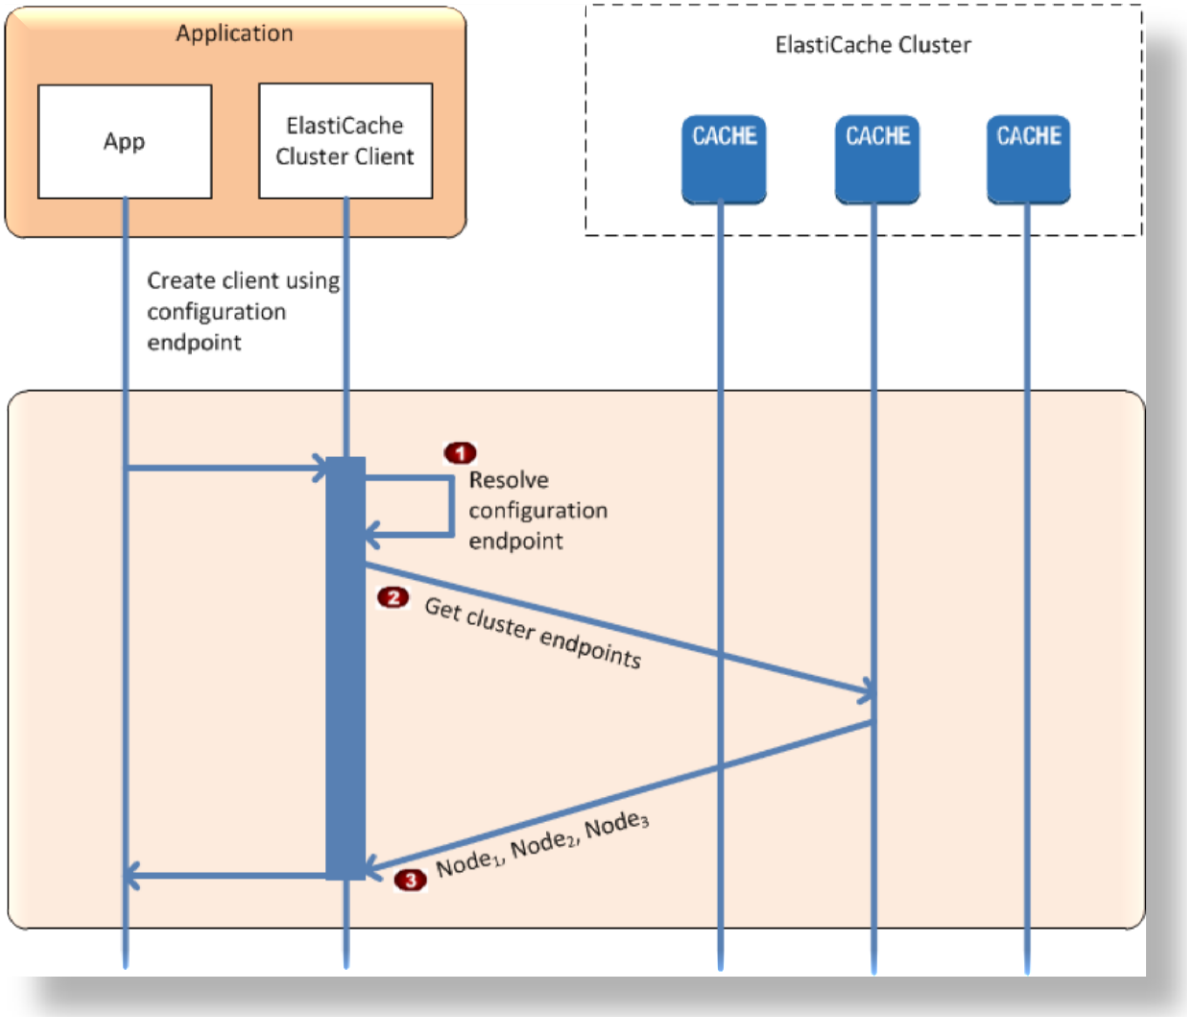

Play with Data

```
telnet cachecluster.mm.0001.usw2.cache.amazonaws.com 11211
```

```
set a 0 0 5       // Set key "a" with no expiration and 5 byte value
hello             // Set value as "hello"
STORED
get a             // Get value for key "a"
VALUE a 0 5
hello
END
get b             // Get value for key "b" results in cache miss
END
```

Knowing about Cache can be helpful…

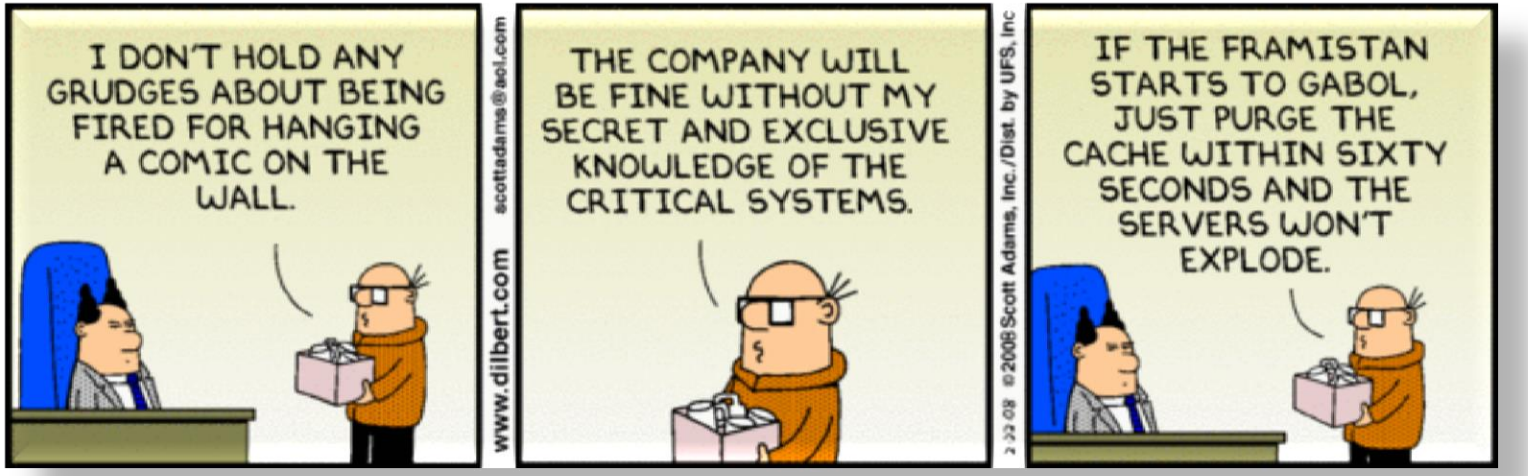

The Dark Side

- Beta
- Limits on Replication and Commands
- Only one AZ for Memcache?

- Permissions
  - Seem convoluted. Need to map to VPC use EC2 security groups on top of cache security groups

What about the cost

- Based on EC2 resources consumed
  - Size of instance
  - Instance Type
  - Data in / out

What I don’t like

- Not cross regional
- Memcache cannot be seeded
- Does not scale automatically
  - Monitor and Script
- May need rewrite application or use amazon client

Client Libraries

- Java
- PHP
- Python
- Ruby
- .Net

Let's Play

Experiments!

Summary

- You can easily deploy two different cache engines in the cloud
- Cluster them
- Only pay for EC2 no other surcharges

On that note…


HI, MY NAME IS. . .
DON'T BOTHER.
CLICK
Dilbert.com DilbertCartoonist@gmail.com
MY APP DOES FACIAL RECOGNITION AND SEARCHES ALL SOCIAL MEDIA TO GIVE ME YOUR FULL BIOGRAPHY.
12-17-10 ©2010 Scott Adams, Inc./Dist. by UFS, Inc.
HOW'S THAT WORKING OUT?
YOU'RE EITHER BART SIMPSON OR A HUGE DRY-ERASE MARKER.

Next Meeting

- When: Nov 21st
- Where: Packard Place ?
- Who: Mike McKeown
- What: High Availability and Disaster Recovery in Azure Applications

THANK YOU

@BmanClt

http://BonCode.blogspot.com