



CS 686: Special Topics in Big Data

Dynamo

Lecture 10

Tweets

- Dynamo – tune your storage system to 99.9% availability if you are okay with a little inconsistency. The magic behind the giant and always-on Amazon!
- Techniques have strong vitality when they're tailored to serve market's need!
- Dynamo: "We have your data... all your data. We just need to figure out which is the right version."
- Dynamo only supports primary-key access to data, which is useful for services such as shopping carts and session management.

Review Challenges

- Finding any glaring issues with Dynamo can be a tall order
- The system uses established techniques, targets a narrow set of use cases
- We know that if Amazon is doing something they're probably going to do it well... Money is on the line!

Jeff Bezos' Mandate (~2002)

- Accidental post by Google employee that previously worked at Amazon
- 6 years: From online bookstore to \$1bn IaaS provider
- Bezos *reportedly* issued a mandate to his company:
 1. All teams will henceforth expose their data and functionality through service interfaces.
 2. Teams must communicate with each other through these interfaces.

Bezos' Mandate (2/3)

3. There will be no other form of inter-process communication allowed: no direct linking, no direct reads of another team's data store, no shared-memory model, no back-doors whatsoever. The only communication allowed is via service interface calls over the network.
4. It doesn't matter what technology they use.
5. All service interfaces, without exception, must be designed from the ground up to be externalizable. That is to say, the team must plan and design to be able to expose the interface to developers in the outside world. No exceptions.

Bezos' Mandate (3/3)

6. *Anyone who doesn't do this will be fired. Thank you; have a nice day!*

Use Cases

- Would you use Dynamo if you were:
 - A bank?
 - Running your own version of Dropbox?
 - (Designed specifically for cat images)
 - Analyzing data from SpaceX launches?
 - Building a social network for goats?
- ...**What** use case would Dynamo cover?

Discussion Groups

- Group 1: 4.4 – 4.5
- Group 3: 6.0 – 6.1
- Group 5: 6.3 – 6.6
- Group 2: 4.6 – 4.9
- Group 4: 6.2