

ElasticLINQ

Bringing the Power of LINQ
to NoSQL for .NET Developers

<http://bit.ly/BradsTalks>

<https://github.com/bradwilson>



Comparing SQL and NoSQL Document Stores

- o Flat tables
 - o *Normalization*
 - o *Strong relationships*
- o Centralized
 - o *Scale vertically*
 - o *ACID*
- o Known quantity
- o (Usually) expensive
- o Complex documents
 - o *De-normalized*
 - o *Weak relationships*
- o Distributed
 - o *Scale horizontally*
 - o *Usually not ACID*
- o New to most people
- o (Usually) open source

Use the best tools for the job

Considerations When Moving to a Document Store


- o Documents \sim domains
 - o *Typically the concurrency granularity*
 - o *Extra round trip for relationships*
 - o *Be careful of payload size*
- o Ad-hoc reporting can be challenging
 - o *Some document stores offer map/reduce*
 - o *3rd party data aggregation & reporting*
 - o *Replicate counted/reported data into SQL*

Why Elasticsearch?

- o Search: full text & facets (uses Lucene)
- o DB integration, or use it stand-alone
- o Built for the web
 - o *Simple HTTP-based API*
 - o *JSON is a first class citizen*
 - o *Dynamic schema support*
 - o *Low-friction sharding and replication*
- o Lots of great documentation
- o Available support contracts

[http://www.elasticsearch.\[org|com\]/](http://www.elasticsearch.[org|com]/)

Execution Pipeline



DEMO

Why ElasticLINQ?

- o JSON query language is challenging
- o Automatic mapping of CLR \leftrightarrow JSON
- o Because it's LINQ:
 - o *Composition & inspection*
 - o *Deferred execution*
 - o *Easier to adapt when moving from SQL*
 - o *Use the data models you already have*

<https://github.com/CenturyLinkCloud/ElasticLINQ>

DEMO

Why Couchbase?

- o Capable of massive scale
 - o *Tens of thousands of users*
 - o *Billions of documents*
 - o *Cross-data center replication*
 - o *Automatic sharding and replication*
- o Views: Javascript-based map/reduce
- o Integrates well with Elasticsearch
- o Open source & commercial options

<http://www.couchbase.com/>

Why Use Couchbase With Elasticsearch?

o **Couchbase:**

Document storage & retrieval

- o *No searching, just retrieval by index (and range)*
- o *Plugin support is limited*

o **Elasticsearch:**

Industrial-strength indexing

- o *Does not offer map/reduce transformation*
- o *Does not do cross-data center replication*

Extremely complementary technologies

Strategy: Type Key

```
{  
  "type": "beer",  
  "name": "Feral Saison",  
  "abv": 6.7,  
  "ibu": 30,  
  ...  
}
```

```
{  
  "type": "brewery",  
  "name": "Black Raven",  
  "city": "Redmond",  
  "state": "WA",  
  ...  
}
```

Can run afoul of Elasticsearch schema

Strategy: Type Wrapper

```
{
  "beer": {
    "name": "Feral Saison",
    "abv": 8.9,
    "ibu": 30,
    ...
  }
}

{
  "brewery": {
    "name": "Black Raven",
    "city": "Redmond",
    "state": "WA",
    ...
  }
}
```

Requires wrapping and unwrapping documents

DEMO

What's Next?

- o Explore Elasticsearch and Couchbase
 - o *CenturyLink Cloud is good for this 😊*
- o Try ElasticLINQ on your data sets
 - o *Lots of holes to be filled*
 - o *Custom search extensions*
- o Contribute
 - o *Docs, bugs, and code are always welcome!*

Thank You!

<http://bit.ly/BradsTalks>

<https://github.com/bradwilson>

