

Eventual Consistency

In the real world

or

Why You Already Know
Eventual Consistency

or

**Eventual Consistency
is better* than
Eventual Availability**

*depending on the use case

@roder

Matt Heitzenroder



basho



We <3 Distributed Systems




Basho has a Distributed Systems experts on the board – Eric Brewer

Brewer's Conjecture

Eric Brewer, 2000
Symposium on Principles of Distributed Computing

Eric Brewer, UC Berkley 2000



impossible for a distributed system to all 3 guarantees simultaneously

CAP Theorem

2002, Seth Gilbert and Nancy Lynch, MIT

formal proof in 2002

Life is full of
Tradeoffs

Amazon's Dynamo Paper

2007, Werner Vogels
Symposium on Operating Systems

addresses need for incrementally scalable, highly-available key-value storage system


Eventual Consistency


2007, Werner Vogels

Shopping cart is the defacto example of eventual consistency

Eventual Availability

We don't talk about this enough – what does it look like?





Eventual Consistency

In the real world

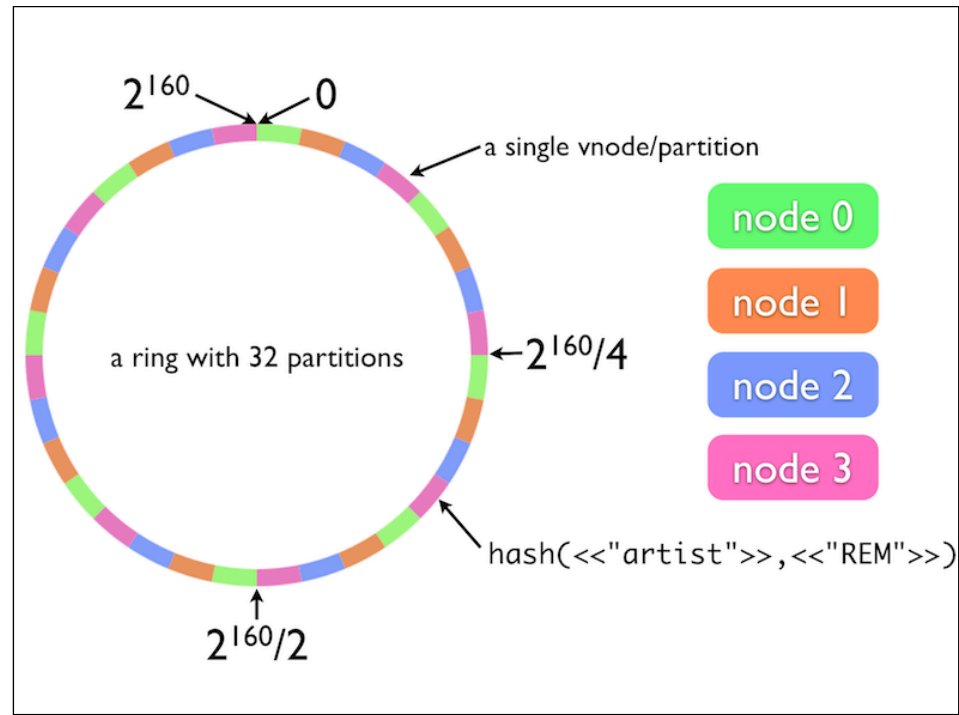







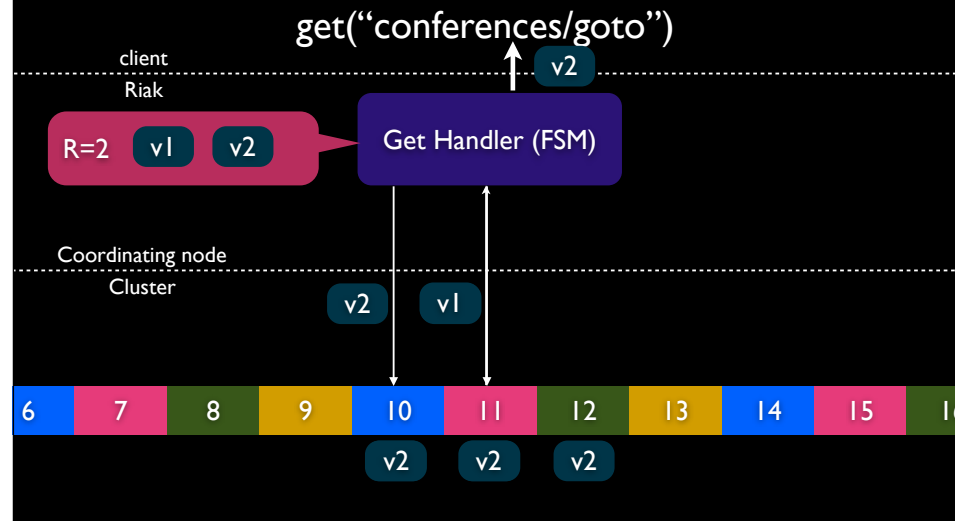
Eventual Consistency

In Riak



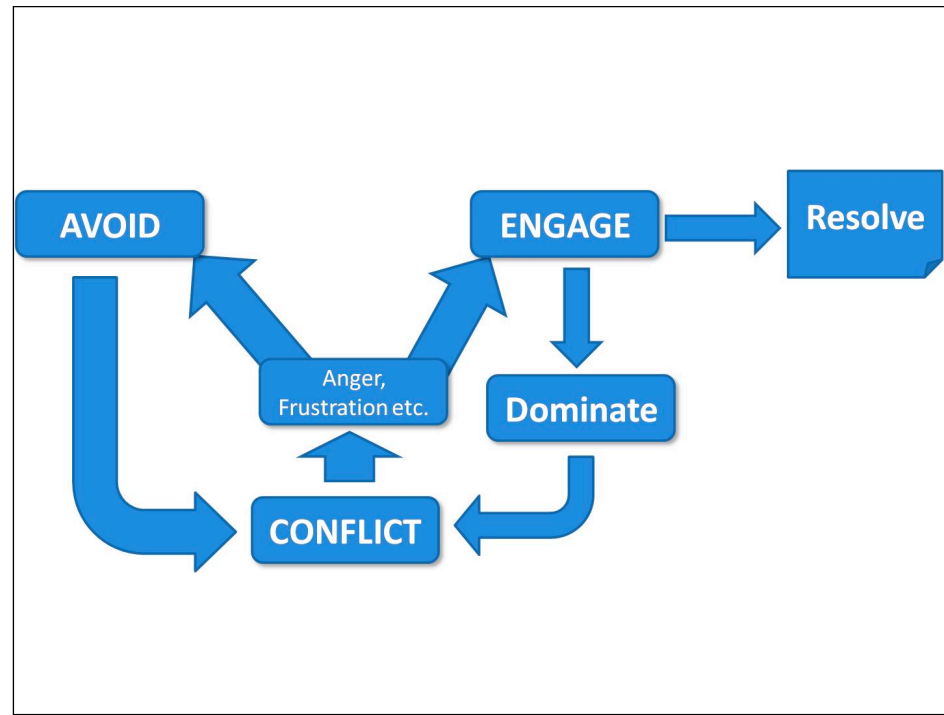


Read Repair



Divergent Object Versions

aka “Siblings”



*"We don't ever do conflict
resolution by picking a
particular sibling."*

Myron Marston, SEOMoz


```

1 class EventList
2   include Ripple::Document
3
4   property :date, Date, presence: true
5   property :hour, Integer, presence: true, numericality: { greater_than_or_equal_to: 0, less_than: 24 }
6   timestamps!
7
8   many :events
9
10  after_validation :assign_key, on: :create
11
12  on_conflict(:events) do |siblings, conflicted_attributes|
13    self.events = siblings.map(&:events).inject(:[])
14  end
15
16  def assign_key
17    self.key = [date.iso8601, hour].join('.')
18  end
19 end
20
21 class Event
22   include Ripple::EmbeddedDocument
23
24   property :timestamp, DateTime
25   property :id, String
26   property :metadata, Hash, default: {}
27 end

```

courtesy of Myron Marston, SEOMoz

“For an array property, we often take the union of all values in all siblings. This works great for array properties that we only ever add to.”

Myron Marston, SEOMoz

set union

“For a timestamp property, we often take the maximum sibling value or the minimum sibling value, depending on the semantics of that attribute.”

Myron Marston, SEOMoz

“For properties that don't have semantics that support these sorts of automatic resolution, we will often take the value from the sibling with the latest `updated_at` value.”

Myron Marston, SEOMoz

timestamp in the body of the data object



“Storing a communication between two users[...]will be written once[...]but it can be updated multiple times. The updates are resolved as a time sorted list.”

Will Moss, Bump

set union sorted by timestamp that is part of the metadata

“For every photo (or other large data item) sent via Bump we back it up to S3, but keep a little metadata about the item.[...] Resolutions are simply a matter of doing a set union between these two values.”

Will Moss, Bump

set union based on the metadata



in the real world, events happen concurrently.
We have ways of dealing with it and we must encode them.

<http://pbs.cs.berkeley.edu/>

*quantitatively demonstrate why eventual consistency is
"good enough" for many users*

Matt Heitzenroder

@roder