

Cassandra

A short Intro



Vitaly Kushner
astrails.com

- Distributed
- Non-relational
- Linearly scalable
- Multi data-center aware

BigTable

<http://labs.google.com/papers/bigtable.html>

Terminology Sucks

big time

```
my_cluster: {
  demo_app: {
    accounts: {
      "1": {
        "attributes": {
          "name": "astrails"
        },
        "users": {
          "1" => "",
          "3" => "",
        }
      }
    },
    users: {
      "1": {
        "name": "vitaly",
        "email": "vitaly@astrails.com"
      }
    }
  }
}
```

```
my_cluster: {  
  demo_app: {  
    accounts: {  
      "1": {  
        "attributes": {  
          "name": "astrails"  
        },  
        "users": {  
          "1" => "",  
          "3" => "",  
        }  
      }  
    },  
    users: {  
      "1": {  
        "name": "vitaly",  
        "email": "vitaly@astrails.com"  
      }  
    }  
  }  
}
```

<<= Cluster Name

```
my_cluster: {
  demo_app: {                                     <== Keyspace
    accounts: {
      "1": {
        "attributes": {
          "name": "astrails"
        },
        "users": {
          "1" => "",
          "3" => "",
        }
      }
    },
    users: {
      "1": {
        "name": "vitaly",
        "email": "vitaly@astrails.com"
      }
    }
  }
}
```

```
my_cluster: {
  demo_app: {
    accounts: {
      "1": {
        "attributes": {
          "name": "astrails"
        },
        "users": {
          "1" => "",
          "3" => "",
        }
      }
    },
    users: {
      "1": {
        "name": "vitaly",
        "email": "vitaly@astrails.com"
      }
    }
  }
}
```

<<= super column family


```
my_cluster: {
  demo_app: {
    accounts: {
      "1": {
        "attributes": {
          "name": "astrails"
        },
        "users": {
          "1" => "",
          "3" => "",
        }
      }
    },
    users: {
      "1": {
        "name": "vitaly",
        "email": "vitaly@astrails.com"
      }
    }
  }
}
```

<<= column family

```
my_cluster: {
  demo_app: {
    accounts: {
      "1": {                                     <<= key
        "attributes": {
          "name": "astrails"
        },
        "users": {
          "1" => "",
          "3" => "",
        }
      }
    },
    users: {
      "1": {
        "name": "vitaly",
        "email": "vitaly@astrails.com"
      }
    }
  }
}
```

```
my_cluster: {
  demo_app: {
    accounts: {
      "1": {
        "attributes": {          << super column
          "name": "astrails"
        },
        "users": {
          "1" => "",
          "3" => "",
        }
      }
    },
    users: {
      "1": {
        "name": "vitaly",
        "email": "vitaly@astrails.com"
      }
    }
  }
}
```

```
my_cluster: {
  demo_app: {
    accounts: {
      "1": {
        "attributes": {
          "name": "astrails" <<= column
        },
        "users": {
          "1" => "",
          "3" => "",
        }
      }
    },
    users: {
      "1": {
        "name": "vitaly",
        "email": "vitaly@astrails.com"
      }
    }
  }
}
```

```
my_cluster: {
  demo_app: {
    accounts: {
      "1": {
        "attributes": {
          "name": "astrails" <<= value
        },
        "users": {
          "1" => "",
          "3" => "",
        }
      }
    },
    users: {
      "1": {
        "name": "vitaly",
        "email": "vitaly@astrails.com"
      }
    }
  }
}
```

```
my_cluster: {
  demo_app: {
    accounts: {
      "1": {
        "attributes": {
          "name": "astrails"
        },
        "users": {
          "1" => "",
          "3" => "",
        }
      }
    },
    users: {
      "1": {
        "name": "vitaly",
        "email": "vitaly@astrails.com"
      }
    }
  }
}
```

<<= column

```
my_cluster: {                                     <<= cluster
  demo_app: {                                     <<= keyspace
    accounts: {                                  <<= super column family
      "1": {                                     <<= key
        "attributes": {                         <<= supercolumn
          "name": "astrails"                    <<= column
        },
        "users": {
          "1" => "",
          "3" => "",
        }
      }
    },
    users: {                                     <<= column family
      "1": {                                     <<= key
        "name": "vitaly",                       <<= column
        "email": "vitaly@astrails.com"
      }
    }
  }
}
```

```

my_cluster: {                                     <<= cluster
  demo_app: {                                     <<= database
    accounts: {                                  <<= super column family
      "1": {                                     <<= key
        "attributes": {                         <<= supercolumn
          "name": "astrails"                    <<= column
        },
        "users": {
          "1" => "",
          "3" => "",
        }
      }
    },
    users: {                                     <<= column family
      "1": {                                     <<= key
        "name": "vitaly",                       <<= column
        "email": "vitaly@astrails.com"
      }
    }
  }
}

```



```

my_cluster: {                                     <<= cluster
  demo_app: {                                     <<= database
    accounts: {                                   <<= table
      "1": {                                      <<= key
        "attributes": {                          <<= supercolumn
          "name": "astrails"                     <<= column
        },
        "users": {
          "1" => "",
          "3" => "",
        }
      }
    },
    users: {                                     <<= column family
      "1": {                                      <<= key
        "name": "vitaly",                        <<= column
        "email": "vitaly@astrails.com"
      }
    }
  }
}

```

```
my_cluster: {                                     <<= cluster
  demo_app: {                                     <<= database
    accounts: {                                   <<= table
      "1": {                                     <<= key
        "attributes": {                         <<= supercolumn
          "name": "astrails"                    <<= column
        },
        "users": {
          "1" => "",
          "3" => "",
        }
      }
    },
    users: {                                     <<= table
      "1": {                                     <<= key
        "name": "vitaly",                       <<= column
        "email": "vitaly@astrails.com"
      }
    }
  }
}
```

```
my_cluster: {                                     <<= cluster
  demo_app: {                                     <<= database
    accounts: {                                   <<= table
      "1": {                                       <<= key
        "attributes": {                           <<= collection
          "name": "astrails"                       <<= column
        },
        "users": {
          "1" => "",
          "3" => "",
        }
      }
    },
    users: {                                       <<= table
      "1": {                                       <<= key
        "name": "vitaly",                          <<= column
        "email": "vitaly@astrails.com"
      }
    }
  }
}
```

```
my_cluster: {                                     <<= cluster
  demo_app: {                                     <<= database
    accounts: {                                  <<= table
      "1": {                                     <<= key
        "attributes": {                         <<= collection
          "name": "astrails"                    <<= attribute
        },
        "users": {
          "1" => "",
          "3" => "",
        }
      }
    },
    users: {                                     <<= table
      "1": {                                     <<= key
        "name": "vitaly",                       <<= attribute
        "email": "vitaly@astrails.com"
      }
    }
  }
}
```

Indexing

```
"my_cluster": {
  "demo_app": {
    "users": {
      "b2c75e99-e079-449d-b562-293dfdbf006e": {
        "attributes": {
          "name": "Vitaly",
          "email": "vitaly@astrails.com",
          "remember_me": "74c4f3a2c85b66f216211edd5245d25ba251d83e"
        }
      }
    },
    "emails": {
      "vitaly@astrails.com": "b2c75e99-e079-449d-b562-293dfdbf006e",
    }
    "remember_me": {
      "74c4f3a2c85b66f216211edd5245d25ba251d83e": "b2c75e99-e079-449d-b562-293dfdbf006e",
    }
  }
}
```

```
"my_cluster": {
  "demo_app": {
    "users": {
      "b2c75e99-e079-449d-b562-293dfdbf006e": {
        "attributes": {
          "name": "Vitaly",
          "email": "vitaly@astrails.com",
          "remember_me": "74c4f3a2c85b66f216211edd5245d25ba251d83e"
        }
      }
    },
    "emails": {
      "vitaly@astrails.com": "b2c75e99-e079-449d-b562-293dfdbf006e",
    }
    "remember_me": {
      "74c4f3a2c85b66f216211edd5245d25ba251d83e": "b2c75e99-e079-449d-b562-293dfdbf006e",
    }
  }
}
```

```
"my_cluster": {
  "demo_app": {
    "users": {
      "b2c75e99-e079-449d-b562-293dfdbf006e": {
        "attributes": {
          "name": "Vitaly",
          "email": "vitaly@astrails.com",
          "remember_me": "74c4f3a2c85b66f216211edd5245d25ba251d83e"
        }
      }
    },
    "emails": {
      "vitaly@astrails.com": "b2c75e99-e079-449d-b562-293dfdbf006e",
    }
  }
}
"remember_me": {
  "74c4f3a2c85b66f216211edd5245d25ba251d83e": "b2c75e99-e079-449d-b562-293dfdbf006e",
}
}
```



```
"my_cluster": {
  "demo_app": {
    "users": {
      "b2c75e99-e079-449d-b562-293dfdbf006e": {
        "attributes": {
          "name": "Vitaly",
          "email": "vitaly@astrails.com",
          "remember_me": "74c4f3a2c85b66f216211edd5245d25ba251d83e"
        }
      }
    },
    "emails": {
      "vitaly@astrails.com": "b2c75e99-e079-449d-b562-293dfdbf006e",
    }
  }
}
```

```
"my_cluster": {
  "demo_app": {
    "users": {
      "b2c75e99-e079-449d-b562-293dfdbf006e": {
        "attributes": {
          "name": "Vitaly",
          "email": "vitaly@astrails.com",
          "remember_me": "74c4f3a2c85b66f216211edd5245d25ba251d83e"
        }
      }
    },
    "emails": {
      "vitaly@astrails.com": "b2c75e99-e079-449d-b562-293dfdbf006e",
    }
    "remember_me": {
      "74c4f3a2c85b66f216211edd5245d25ba251d83e": "b2c75e99-e079-449d-b562-293dfdbf006e",
    }
  }
}
```

```
"my_cluster": {
  "demo_app": {
    "users": {
      "b2c75e99-e079-449d-b562-293dfdbf006e": {
        "attributes": {
          "name": "Vitaly",
          "email": "vitaly@astrails.com",
          "remember_me": "74c4f3a2c85b66f216211edd5245d25ba251d83e"
        }
      }
    },
    "emails": {
      "vitaly@astrails.com": "b2c75e99-e079-449d-b562-293dfdbf006e",
    }
    "remember_me": {
      "74c4f3a2c85b66f216211edd5245d25ba251d83e": "b2c75e99-e079-449d-b562-293dfdbf006e",
    }
  }
}
```

Playtime

```
> gem install cassandra
```

```
...
```

```
> cassandra_helper cassandra
```

```
...
```

```
<ClusterName>my_cluster</ClusterName>
```

```
<Keyspaces>
```

```
  <Keyspace Name="demo_app">
```

```
    <ColumnFamily CompareWith="UTF8Type" Name="users" ColumnType="Super" CompareSubcolumnsWith="UTF8Type" />
```

```
    <ColumnFamily CompareWith="UTF8Type" Name="emails" ColumnType="Super" CompareSubcolumnsWith="UTF8Type" />
```

```
    <ColumnFamily CompareWith="UTF8Type" Name="tokens" />
```

```
  </Keyspace>
```

```
</Keyspaces>
```

```
require 'rubygems'
require 'cassandra'

include Cassandra::Constants

client = Cassandra.new('demo_app')

client.insert(:users, "1", {"name" => "vitaly", "email" => "vitaly@astrails.com"})
client.insert(:accounts, "1", {"attributes" => {"name" => "astrails"}, "users" => {"1" => ""}})

client.insert(:users, "2", {"name" => "boris"})
client.insert(:accounts, "1", {"users" => {"2" => ""}})

client.get(:users, "1")
```

Rails?

YES!

SmallRecord

WARNING: waporware :)

http://github.com/astrails/small_record

ActiveModel

```
class Account < ActiveRecord::Base
  attribute :name
  attribute :created_at, :type => :datetime

  has_many :users

  validates_presence_of :name
end
```

```
class User < ActiveRecord::Base
  attribute :name
  attribute :email
  attribute :created_at, :type => :datetime
  attribute :account_id

  def account
    @account ||= Account.find(:account_id)
  end

  validates_presence_of :name
  validates_presence_of :email
end
```

```
Account.find(params[:account]).users.first.name
```

production:

adapter: cassandra

host: 127.0.0.1

port: 9160

keyspace: chronos

development:

adapter: mock

test:

adapter: mock

production:

adapter: cassandra

host: 127.0.0.1

port: 9160

keyspace: chronos

development:

adapter: **mock**

test:

adapter: **mock**

Some Links

- <http://incubator.apache.org/cassandra>
- <http://labs.google.com/papers/bigtable.html>
- <http://blog.evanweaver.com/articles/2009/07/06/up-and-running-with-cassandra/>
- <http://arin.me/blog/wtf-is-a-supercolumn-cassandra-data-model>
- http://github.com/astrails/small_record