



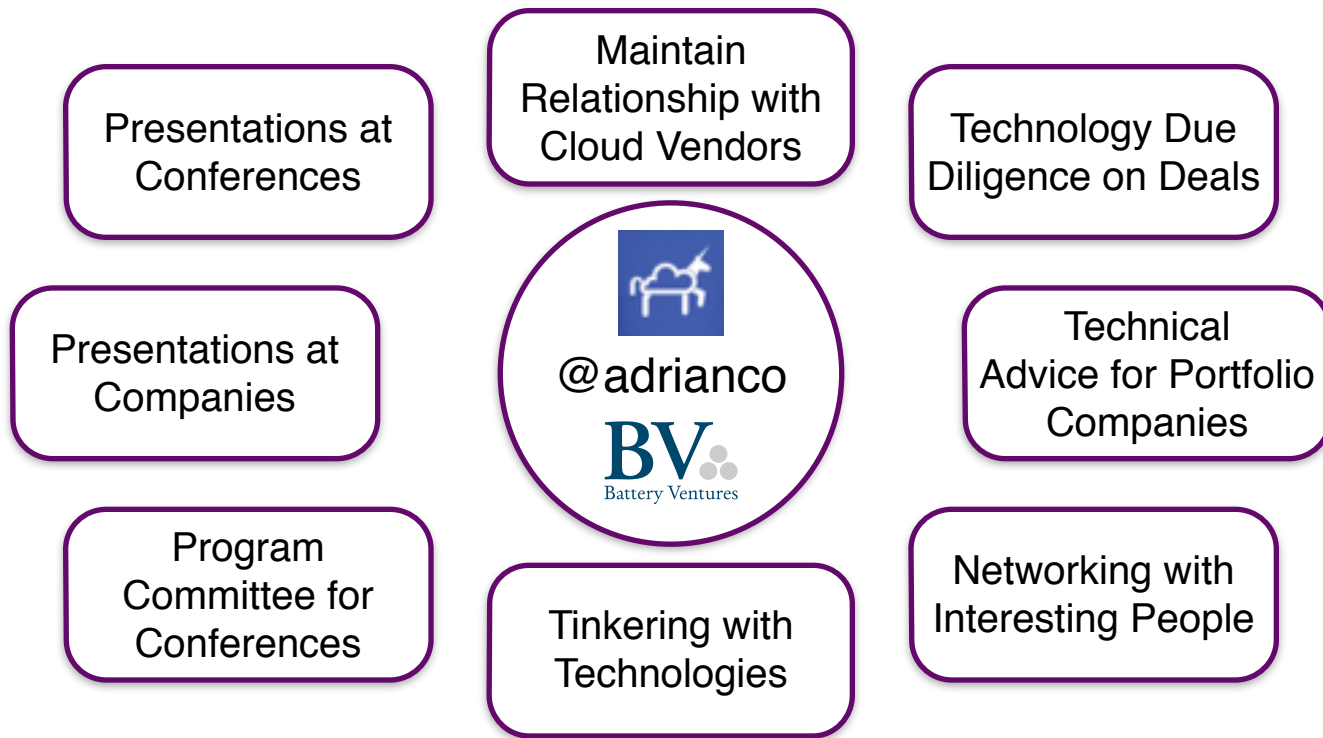
From Microservices to Teraservices

Adrian Cockcroft @adrianco
Technology Fellow - Battery Ventures
September 2015





What does @adrianco do now?







Topics for Today

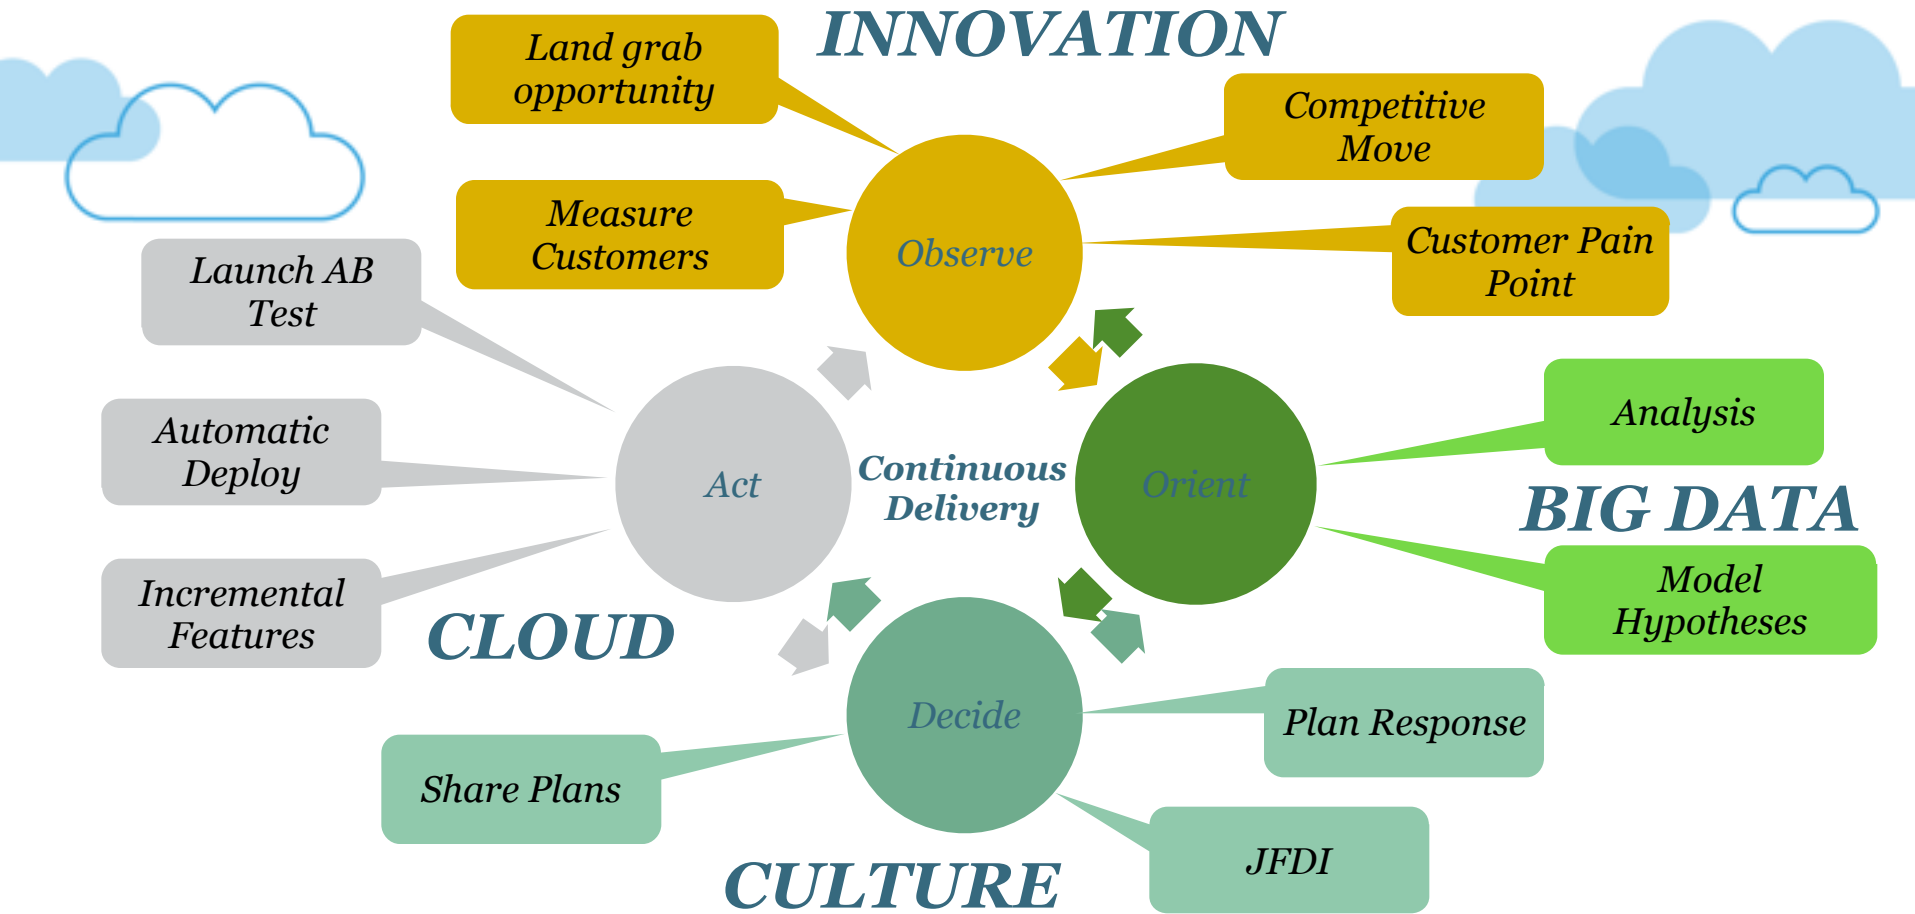
*Why Microservices?
Simulating Architectures
Terabytes of Memory*



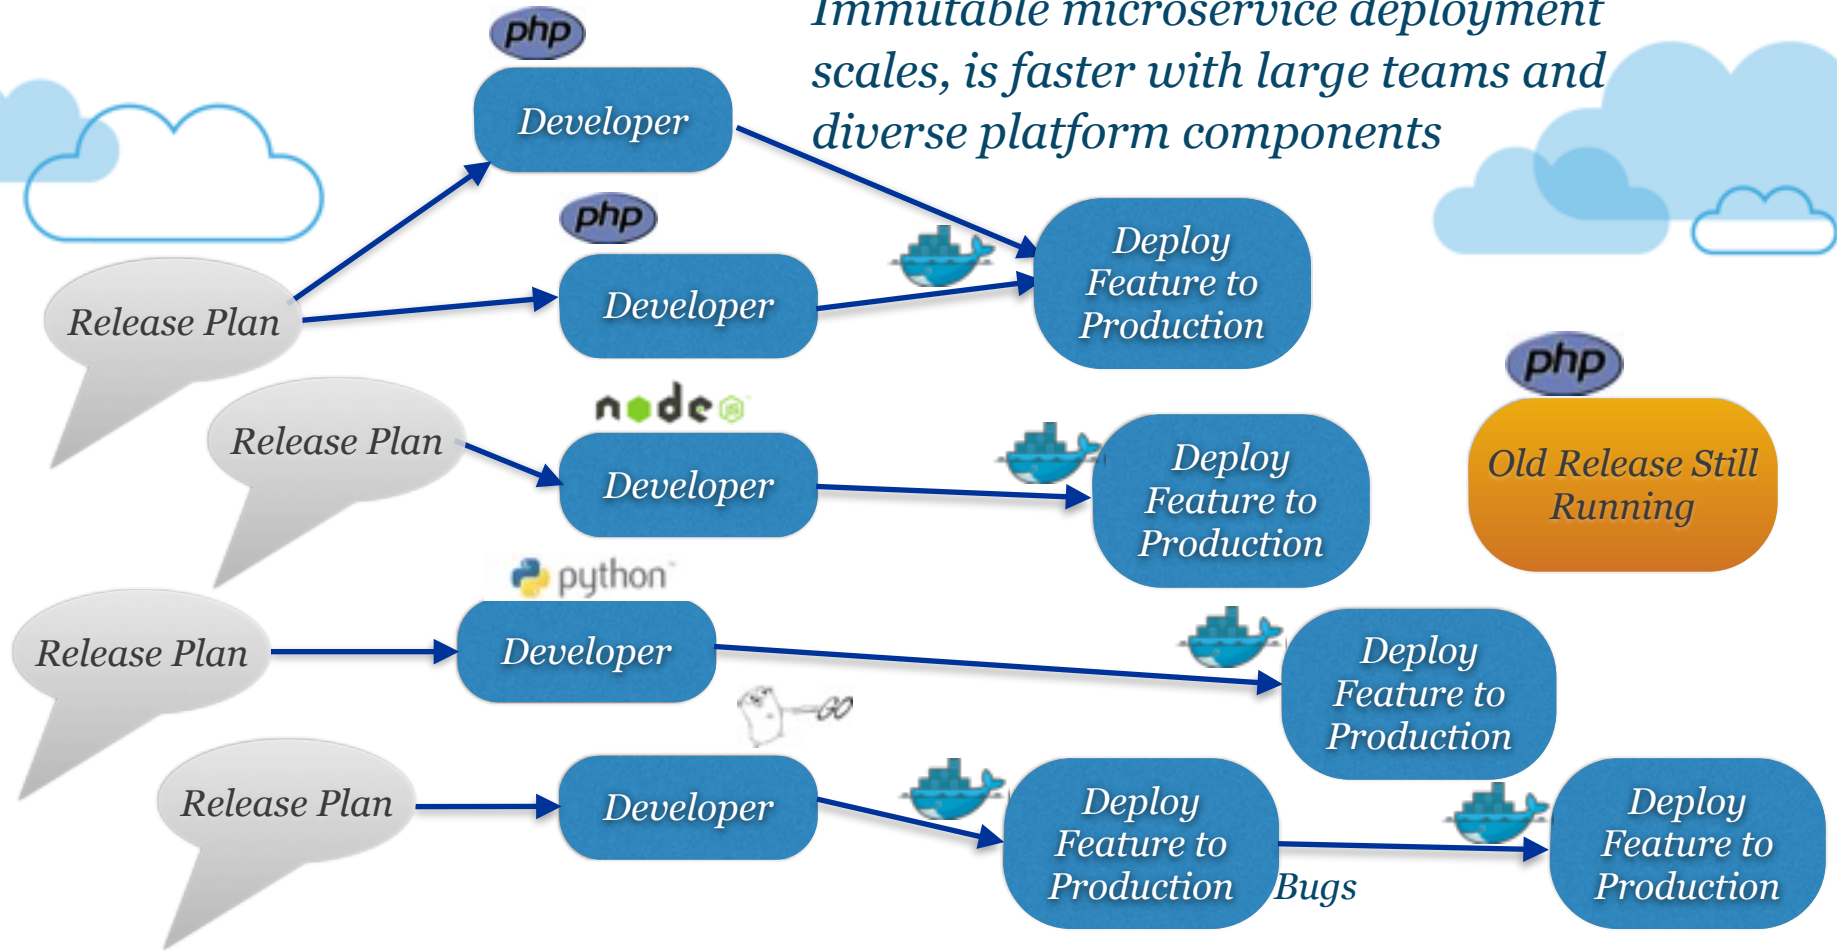


Key Goals of the CIO?
Align IT with the business
Develop products faster
Try not to get breached






Immutable microservice deployment scales, is faster with large teams and diverse platform components




What Happened?






Microservices

If every service has to be updated at the same time it's not loosely coupled



A Microservice Definition

Loosely coupled service oriented architecture with bounded contexts



If you have to know too much about surrounding services you don't have a bounded context. See the Domain Driven Design book by Eric Evans.



Cloud Native Monitoring and Microservices

Cloud Native Microservices



- *High rate of change*

Code pushes can cause floods of new instances and metrics

Short baseline for alert threshold analysis – everything looks unusual

- *Ephemeral Configurations*

Short lifetimes make it hard to aggregate historical views

Hand tweaked monitoring tools take too much work to keep running

- *Microservices with complex calling patterns*

End-to-end request flow measurements are very important

Request flow visualizations get overwhelmed




The background features decorative blue lines and cloud icons. A line starts at the top left, goes right, then down, then right again with several small arrows pointing right. Another line starts at the top right, goes left, then down, then left again with two small arrows pointing left. A cloud icon is positioned near the top right line.


Challenges for Microservice Platforms



Managing Scale



*It's much more challenging
than just a large number of
machines*



A Possible Hierarchy

Continents

Regions

Zones

Services

Versions

Containers

Instances

How Many?

3 to 5

2-4 per Continent

1-5 per Region

100's per Zone

Many per Service


1000's per Version

10,000's



The image features decorative blue lines and arrows. A line enters from the top left, turns right, and then down to join a horizontal line. This horizontal line contains several right-pointing arrows, followed by a white triangle with a blue outline, and then more right-pointing arrows. The line continues to the right, passing under a light blue cloud, and then loops back down and left to join another horizontal line. This second horizontal line has two left-pointing arrows and then turns down.

Flow

*Some tools can show
the request flow
across a few services*



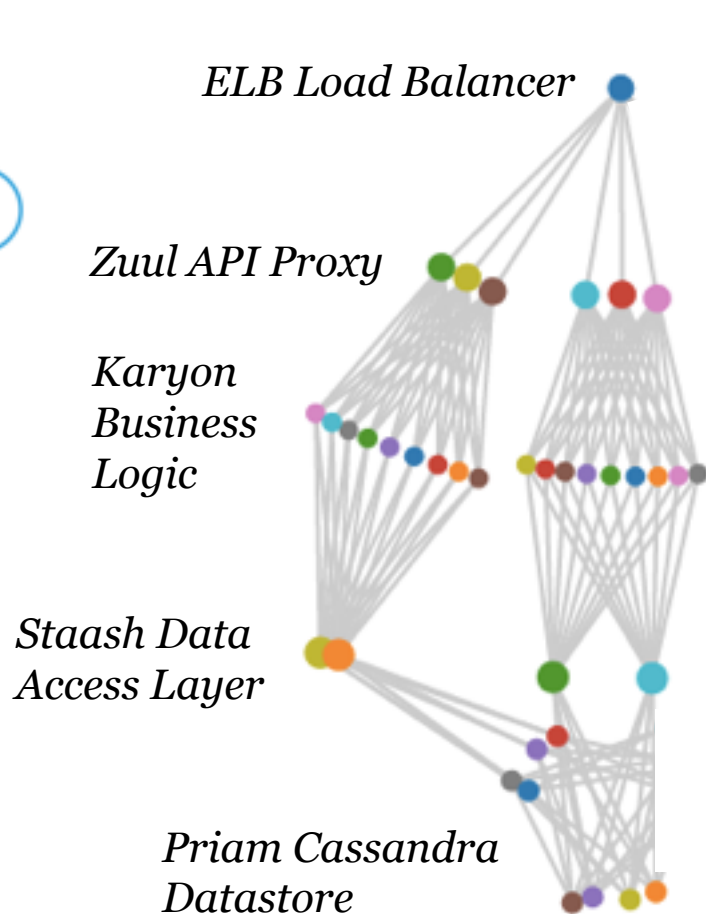
*But interesting
architectures have a
lot of microservices!
Flow visualization is
a big challenge.*





Failures





*Simple NetflixOSS
style microservices
architecture on three
AWS Availability Zones*

*Zone partition/failure
What should you do?
What should monitors show?*

*By design, everything works
with 2 of 3 zones running.
This is not an outage, inform
but don't touch anything!
Halt deployments perhaps?*

*Challenge: understand and
communicate common
microservice failure patterns.*



Testing





*Testing monitoring tools at scale
gets expensive quickly...*




The background features decorative blue lines and clouds. A line starts at the top left, goes right, then down, then right again with several small arrows pointing right. It then continues right with a larger blue arrow pointing right, followed by a cloud icon. Another line starts at the bottom left, goes right with two small blue arrows pointing right, then turns down.

Simulation



Simulated Microservices



Model and visualize microservices
Simulate interesting architectures
Generate large scale configurations
Eventually stress test real tools

See github.com/adrianco/spigo
Simulate Protocol Interactions in Go
Visualize with D3



Definition of an architecture

Header can include
chaos monkey victim

```
{
  "arch": "cassandra",
  "description": "Simple Cassandra model for the Cassandra Summit 2015 paper",
  "version": "arch-0.0",
  "victim": "",
  "services": [
    { "name": "cassandra", "package": "priamCassandra", "count": 6, "regions": 1, "dependencies": ["cassandra", "eureka"]},
    { "name": "restdata", "package": "staash", "count": 6, "regions": 1, "dependencies": ["cassandra"]},
    { "name": "app", "package": "karyon", "count": 12, "regions": 1, "dependencies": ["restdata"]},
    { "name": "proxy", "package": "zuul", "count": 6, "regions": 1, "dependencies": ["app"]},
    { "name": "www-elb", "package": "elb", "count": 0, "regions": 1, "dependencies": ["proxy"]},
    { "name": "www", "package": "denominator", "count": 0, "regions": 0, "dependencies": ["www-elb"]}
  ]
}
```

New tier
name

Tier
package

Node
count

Region
count: 1

List of tier
dependencies

Single Region Architectures




*Service only view
3 zones in 1 region*



*Instance level view
3 zones in 1 region*



*Instance level view
3 zones in 1 region
double scale*



*Instance level view
3 zones in 1 region
quadruple scale*

2 and 3 Region Architectures



*Instance level view
3 zones in 2 regions*




*Instance level view
3 zones in 3 regions*


4, 5 and 6 Region Cassandra



*Instance level view
3 zones in 4 regions*



*Instance level view
3 zones in 5 regions*



*Instance level view
3 zones in 6 regions*



Adding endpoints & clusters



*More realistic architecture for a simple
Netflix-like web service*

*Instance level view
3 zones in 1 region*

*Separate endpoints for API and WWW
Three Cassandra clusters*

Spigo Nanoservice Structure



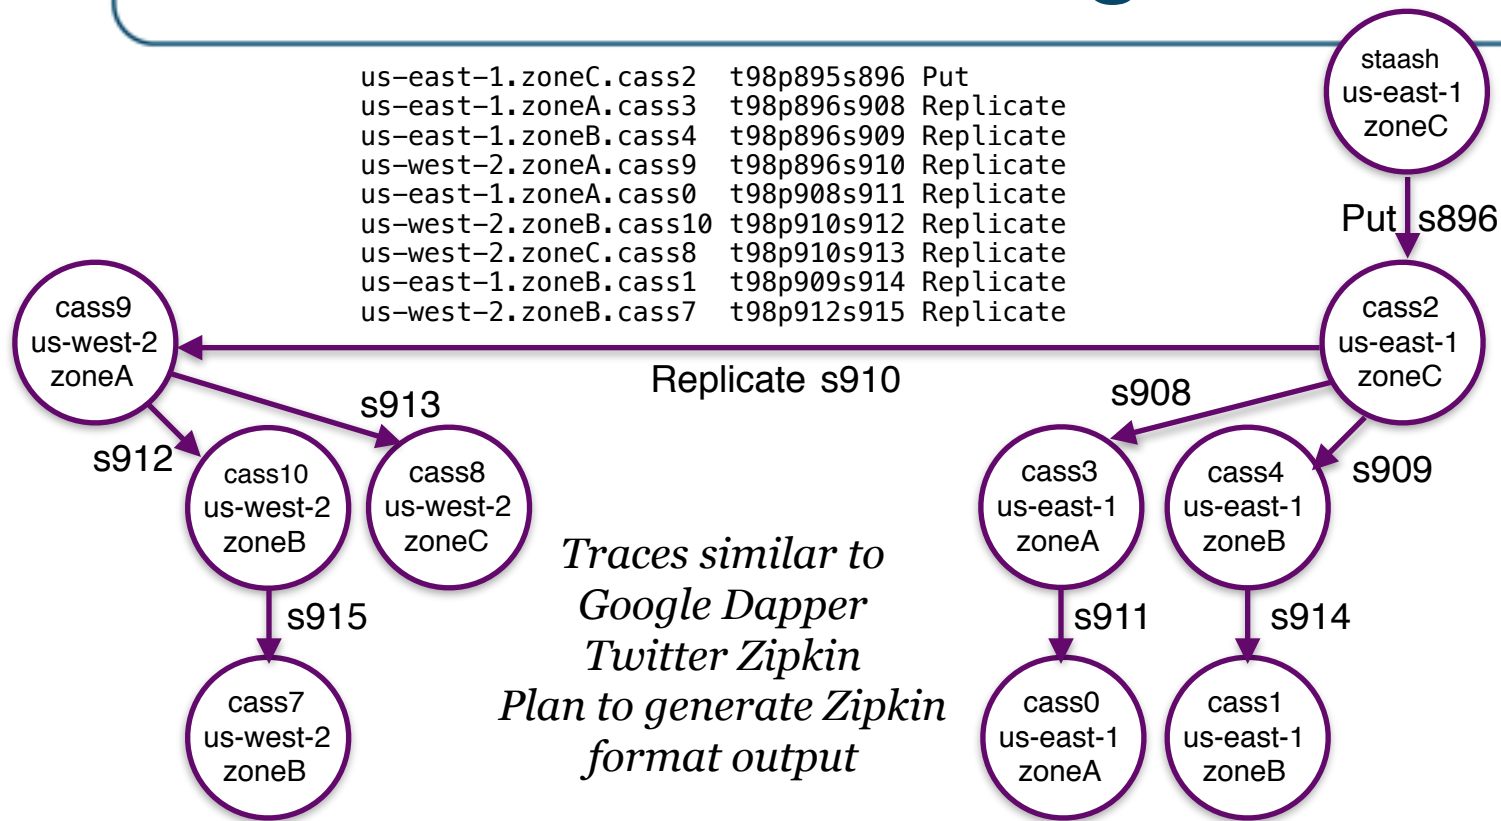
```
func Start(listener chan gotocol.Message) {
    ...
    for {
        select {
        case msg := <-listener:
            switch msg.Imposition {
            case gotocol.Hello:      // get named by parent
                ...
            case gotocol.NameDrop:   // someone new to talk to
                ...
            case gotocol.GetRequest: // upstream request handler
                ...
            case gotocol.GetResponse:// downstream response handler
                ...
            case gotocol.Goodbye:    // tell parent I'm going away now
                gotocol.Message{gotocol.Goodbye, nil, time.Now(), name}.GoSend(parent)
                return
            }
        case <-eurekaTicker.C:      // poll the service registry
            ...
        }
    }
}
```

priamCassandra.go package total about 200 lines of Go



Flow Trace Recording

us-east-1.zoneC.cass2	t98p895s896	Put
us-east-1.zoneA.cass3	t98p896s908	Replicate
us-east-1.zoneB.cass4	t98p896s909	Replicate
us-west-2.zoneA.cass9	t98p896s910	Replicate
us-east-1.zoneA.cass0	t98p908s911	Replicate
us-west-2.zoneB.cass10	t98p910s912	Replicate
us-west-2.zoneC.cass8	t98p910s913	Replicate
us-east-1.zoneB.cass1	t98p909s914	Replicate
us-west-2.zoneB.cass7	t98p912s915	Replicate





Why Build Spigo/Simianviz?



Generate test microservice configurations at scale
Stress monitoring tools display capabilities

Eventually (i.e. not implemented yet)

Dynamically vary configuration: autoscale, code push

Chaos monkey for microservice, zone, region failures

D3 websocket dynamic browser interface

github.com/adrianco/spigo





Teraservices





Terabyte Memory Directions



Engulf dataset in memory for analytics

Balanced config for memory intensive workloads

Replace high end systems at commodity cost point

Explore non-volatile memory implications





A Terabyte Memory Option



Diablo - a Battery Ventures portfolio company

DDR4 DIMM containing flash 64/128/256GB

Migrates pages to/from companion DRAM DIMM

Shipping now as volatile memory, future non-volatile



Memory1: 1st All-Flash System



TM



- **NO CHANGES** to CPU or Server
- **NO CHANGES** to Operating System
- **NO CHANGES** to Applications

- ✓ **UP TO 256GB DDR4 MEMORY PER MODULE**
- ✓ **UP TO 4TB MEMORY IN 2 SOCKET SYSTEM**




In-Memory Database: *Minimize Infrastructure Cost*

Problem: IMDB memory requirements force purchase of additional unneeded CPUs

Solution: Increased memory-per-socket eliminates need for additional costly CPUs


Quad-Processor Configuration (With DRAM Only)



DRAM capacity limitations necessitate additional DIMM slots




Dual-Processor Configuration (System Memory Expanded By Memory1)



With **more memory per socket**, IMDBs can be supported by **less expensive servers**

Takeaway



*Teraservices
moving to
mainstream*





See www.battery.com for a list of portfolio investments

Q&A

Adrian Cockcroft @adrianco
<http://slideshare.com/adriancockcroft>
Technology Fellow - Battery Ventures
July 2015

Enterprise IT

