

Beamly 

Why Devops Needs



Ali Asad Lotia | ali.asad.lotia@gmail.com | @aalotia

What's Beamly?

- TV Focused Social Network
- Smart TV planner
- Personalised TV Magazine
- <http://about.beamly.com>



Availability

- Multiple Countries
- Multiple Platforms



Behind the curtain

- AWS
- Multiple Regions
- *milli*-Services
 - Scala
 - NodeJS



Platform Team

- Build / Release Automation
- Persistence
- Platform performance/metrics/logs
- Core libs




Deploys - The Mutable Days

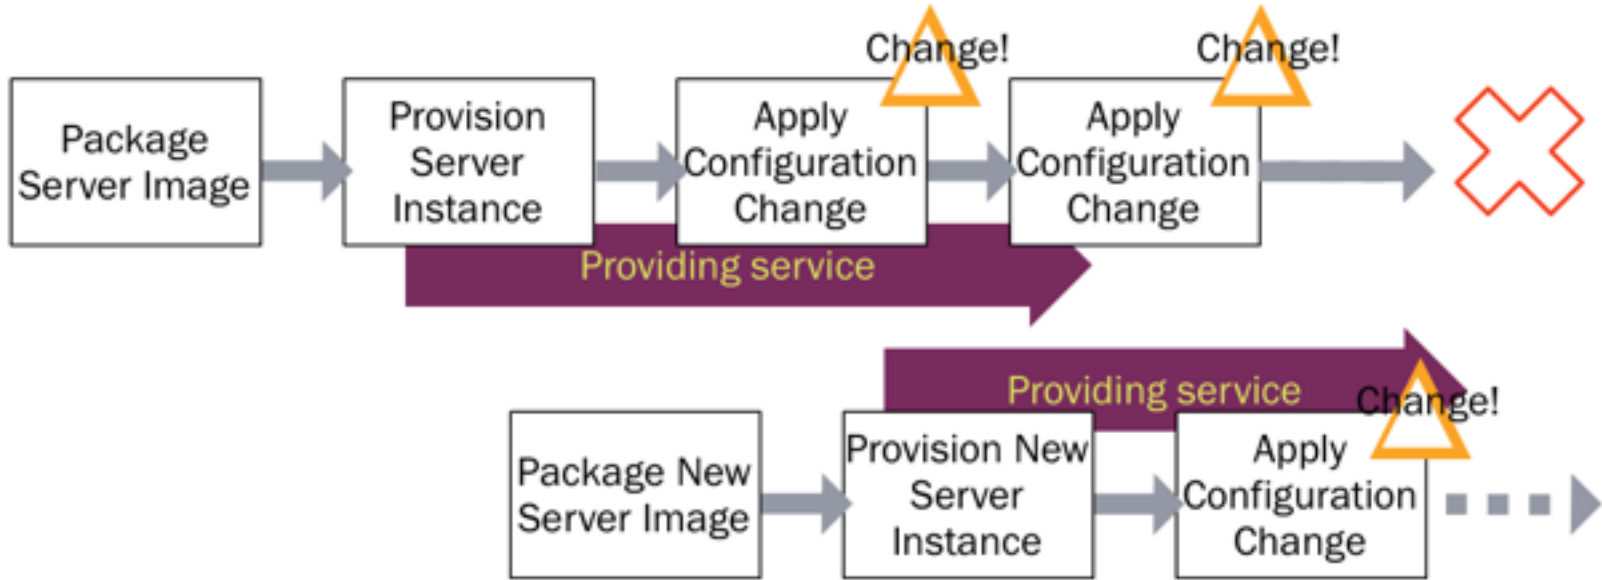
- Generate build artefacts
- Define config in Puppet
- Deploy build artefacts to server
- Deploy config to server



Build



Phoenix Server



Phoenix Server

- Short-lived compared to “pet” servers
- Reduced configuration drift
- But why have any?



Disconnected Dev and Ops

- Isolate state change in code
- Why not config?



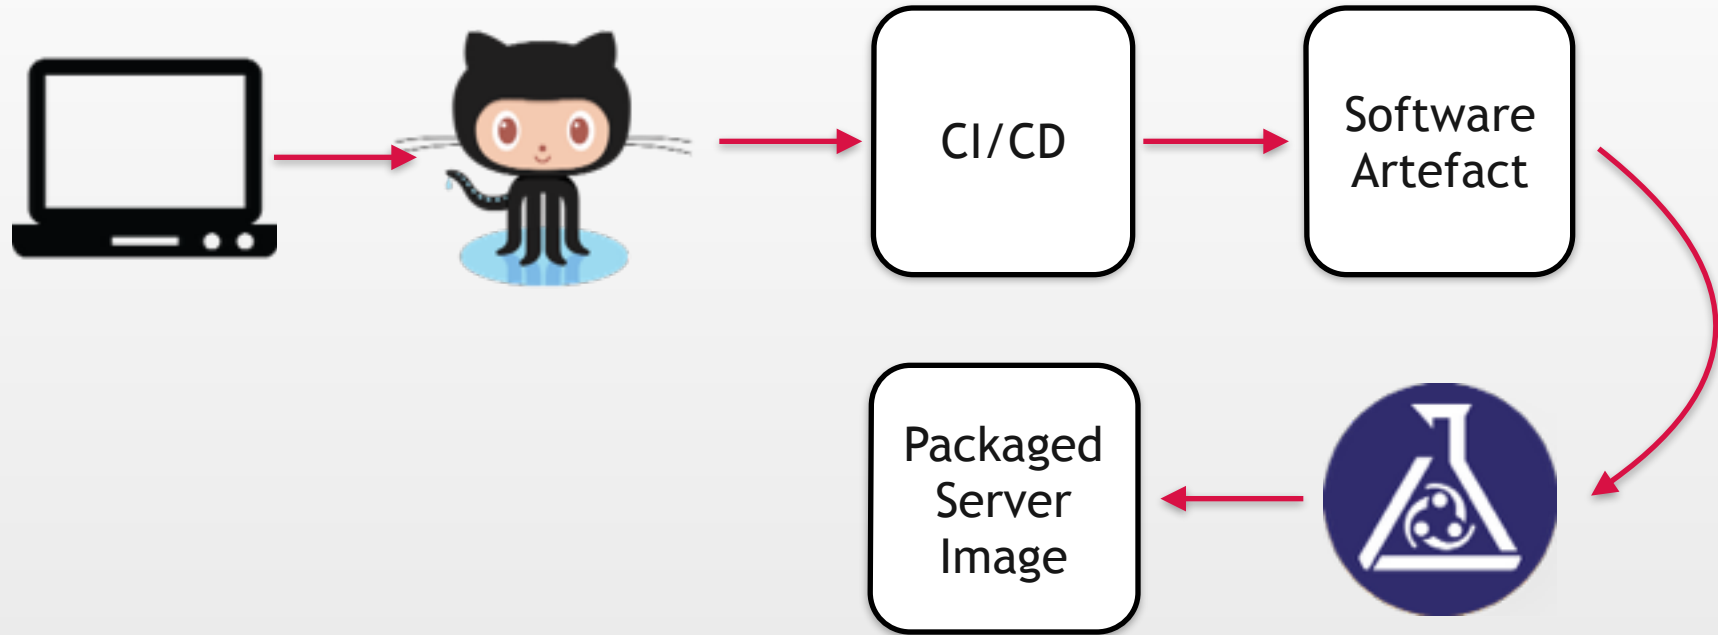
“maybe you use a language like Lisp that pretends the computer is some purely functional fantasy land with padded walls for little babies”

Zed Shaw

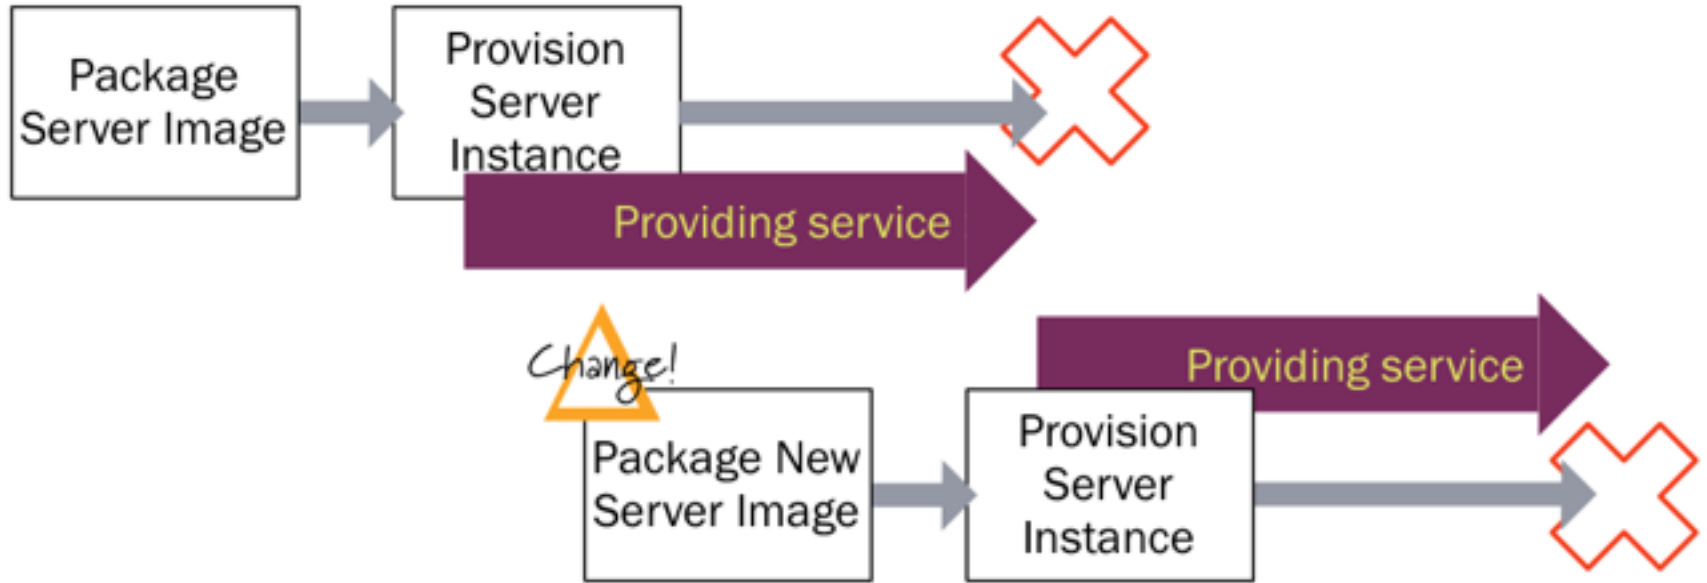
<http://c.learncodethehardway.org/book/introduction.html>



Artefacts: S/W + Configuration



Immutable Server



Requirement: Examine Server Images

- List Server Images for Component
- Identify root device
- Create volume from snapshot
- Attach volume to running instance



AWS Console: AMI details

The screenshot displays the AWS Management Console interface for an Amazon Machine Image (AMI). At the top, there is a search bar with the text 'cassandra' and a filter set to 'Owned by me'. Below the search bar, a table lists AMIs, with the selected row being 'cassandra-09-201406171220' with AMI ID 'ami-12af017a'. The details for this AMI are shown below, including its name, source, owner, visibility, status, and platform. The 'Details' tab is active, showing various attributes such as AMI ID, Owner, Status, Platform, Image Type, Description, Root Device Type, Kernel ID, Block Devices, AMI Name, Source, State Reason, Architecture, Virtualization type, Root Device Name, RAM disk ID, and Product Codes.

Name	AMI Name	AMI ID	Source	Owner	Visibility	Status	Platform
cassandra-09-201406171220	ami-12af017a	440-██████████-c...	440-██████████	Private	available	Other Linux	

Image: ami-12af017a

Details | Permissions | Tags | Edit

AMI ID	ami-1██████████	AMI Name	cassandra-09-201406171220
Owner	440-██████████	Source	440-██████████:cassandra-09-201406171220
Status	available	State Reason	-
Platform	Other Linux	Architecture	x86_64
Image Type	machine	Virtualization type	paravirtual
Description	cassandra v09	Root Device Name	/dev/sda1
Root Device Type	ebx	RAM disk ID	-
Kernel ID	dk-88a75e1	Product Codes	-
Block Devices	/dev/xda1:caep-██████████-tracsterband, /dev/xdb:uphemerd		



Identifying the Server Image

```
import boto.ec2
conn = boto.ec2.connect_to_region("us-east-1")
images = conn.get_all_images(owners="self",
                             filters={"tag:COMPONENT_NAME":
                                     "cassandra"})
```

```
images
[Image:ami-02f51d6a,
 Image:ami-10f11878,
 Image:ami-1cd92d74,
 Image:ami-26ad504e,
 ...]
```



A New Hope

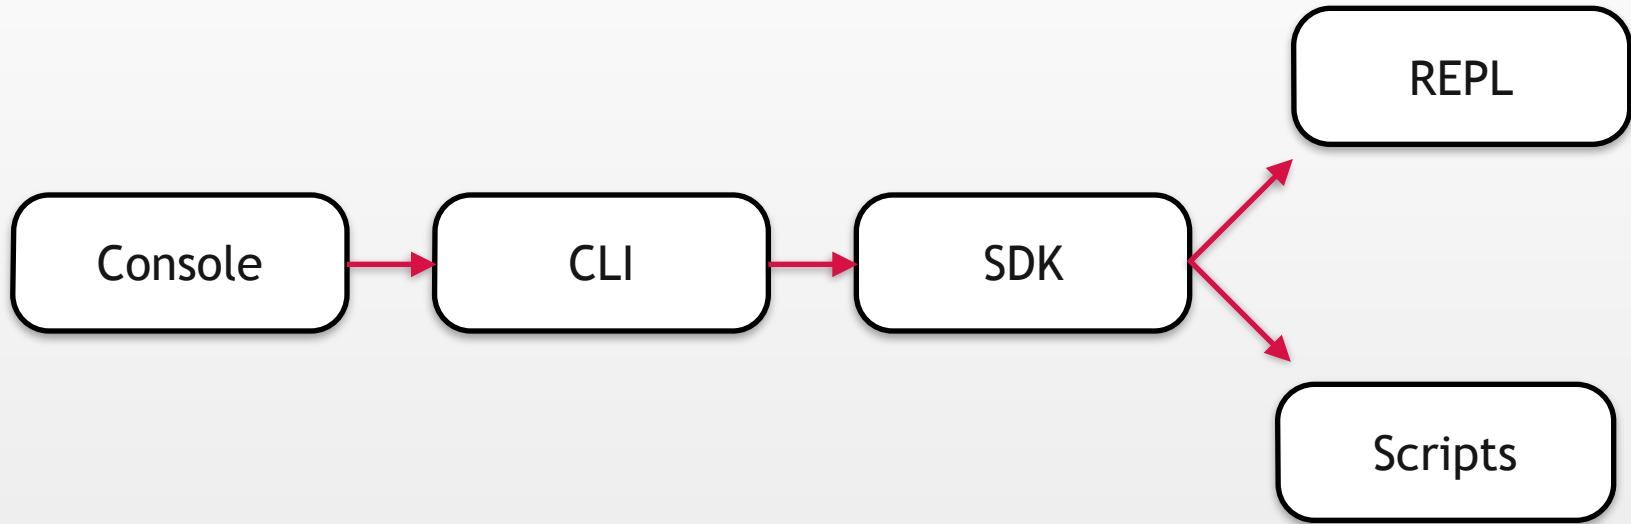
```
(require '[amazonica.aws.ec2 :as ec2])
(require '[clojure.pprint :as pp])
(def images
  (ec2/describe-images
   :owner "self"
   :filters [{:name "tag:COMPONENT_NAME"
              :values ["cassandra"]}]))
(pp/pprint images)
{:images
 [{:description "cassandra v32",
   :block-device-mappings
 [{:device-name "/dev/sda1",
   :ebs
   {:snapshot-id "snap-0*****1",
    :volume-type "standard"...}]}]}
```



(> data objects)



Infrastructure tooling: Evolution



Why REPLs

- Facilitate exploration of APIs
- Minimise context switching
- Instant feedback
- Data rich-er



Interfaces: Desired properties

- Data-rich representations
- Fast feedback
- Safe modification



Team Reactions

- “Soooo many brackets!!!!”
- “How do I iterate over this?”
- “I want to change this value”
- “Wow, this is really powerful”



Offloading State

“Often you can pass the buck to a service which someone else maintains”

Immutable Server

Kief Morris

13 June 2013

<http://martinfowler.com/bliki/ImmutableServer.html>



The “passed buck”

- Application Data
- Metrics
- Logs
- Facilitated autoscaling



Autoscaling: Supporting Infrastructure

- Auto scaling groups
- Launch configurations
- Load balancers
- DNS



Was the buck passed?

- Persisted by cloud provider
- Provider defined data model
- Varied query facilities
- Poor tooling for ancillary config



Herding the infra-cats

Requirement: Don't use us-east-1 c AZ

- Configurations had drifted
- Collections of entities
- Manipulate *safely*



The Clojure Strikes Back

```
(defn get-all-asgs-for-region
  [region]
  (loop [{:keys [next-token auto-scaling-groups]}
        (as/describe-auto-scaling-groups {:endpoint region})]
    (if (nil? next-token)
      auto-scaling-groups
      (let [asg-desc-result
            (as/describe-auto-scaling-groups
             {:endpoint region} :next-token next-token)
            new-next-token
            (:next-token asg-desc-result)
            asgs
            (:auto-scaling-groups asg-desc-result)]
        (recur {:next-token
                new-next-token
                :auto-scaling-groups
                (concat auto-scaling-groups asgs)}))))))
```



```
[{:tags
  [...],
  :health-check-type "EC2",
  ...
  :availability-zones ["us-east-1b" "us-east-1a" "us-east-1d"],
  :launch-configuration-name "***-live-c0us-20140519100859",
  :health-check-grace-period 120,
  :max-size 16,
  :desired-capacity 2,
  :min-size 2,
  :instances
  [...],
  :vpczone-identifier "subnet-f**,subnet-6**,subnet-4**",
  :load-balancer-names [],
  :enabled-metrics
  [{...}],
  :created-time #<DateTime 2014-04-24T11:54:42.921+02:00>,
  :termination-policies ["ClosestToNextInstanceHour"],
  :auto-scaling-group-name "**-live-c0us",
  :suspended-processes
  [{:process-name "AZRebalance",
    :suspension-reason "User suspended at 2014-04-24T09:54:43Z"}]]}
...]
```



Observing platform performance

- Knock-on/trickle down effects
- Needed to explore multiple streams
- Sensu handlers limited
- Riemann <http://riemann.io>



Metrics details

- Sensu gathering graphite metrics
- Key had much info
- Transform for Riemann



Return of the Lambda

```
(defn process-graphite-metric-name
  [{:keys [service metric time]}]
  (let [[component source environment region aws-host metric-details]
        (cs/split service #"\" 6)
        host-ip
        (extract-host-ip-from-aws-hostname aws-host)
        new-service
        (cs/join "." (vector component source metric-details))]
    {:service new-service
     :environment environment
     :region region
     :host host-ip
     :component component
     :metric (or metric 0)
     :time time
     :state "ok"
     :description service
     :tags ["from-sensu" component]}))
```



Riemann - so far

- Embedded REPL!
- Overridable
- Extensible
- Responsive primary author



Tracking our Services

- Zion - system knowledge base
- Component details
 - Tech
 - Ownership
 - Dependencies



Infrastructure as Data

- Config
- Metrics
- Logs



Future Work

- Analyze Logs + Metrics
- Catch and correct misconfigurations
- Scripts with upcoming fastload?
- Cyanite to replace graphite



“Lisp isn't a language, it's a building material.”

- Alan Kay



<http://www.paulgraham.com/quotes.html>



Clojure: Pros

- Core data structures
- Data manipulation
- Community
- Shared aesthetic





XKCD http://imgs.xkcd.com/comics/lisp_cycles.png CC 2.5 BY-NC



References

- Phoenix Server Image <http://martinfowler.com/bliki/images/immutableServer/PhoenixServerLifecycle.png>
- Immutable Server Image <http://martinfowler.com/bliki/images/immutableServer/ImmutableServerLifecycle.png>
- Amazonica - Clojure AWS library <https://github.com/mcohen01/amazonica>
- Cyanite <https://github.com/pyr/cyanite>



Thanks

- You
- Colleagues at Beamly for feedback
- Clojurians who've built awesome stuff





LONDON

Drury House
34-43 Russell Street
London
WC2B 5HA



NEW YORK CITY

84 Wooster Street
Suite 703
New York
NY 10012



SYDNEY

22-36 Mountain Street
Suite 1.10
Ultimo
NSW 2007