



Design Patterns for High Availability

Lessons learned building Amazon CloudFront

CTD303

Alex Smith
Head of M&E Architecture, APAC
Amazon Web Services

Harvo Jones
Sr. Software Development Engineer
Amazon CloudFront

What to Expect from the Session

- Learn about the design patterns for high availability of Amazon CloudFront
- Learn how you can implement the patterns in your own services or applications built on top of AWS

What is a Content Delivery Network (CDN)?



Amazon CloudFront locations worldwide

68 POPs

21 Countries

43 Cities

5 Continents



North America

South America

EMEA

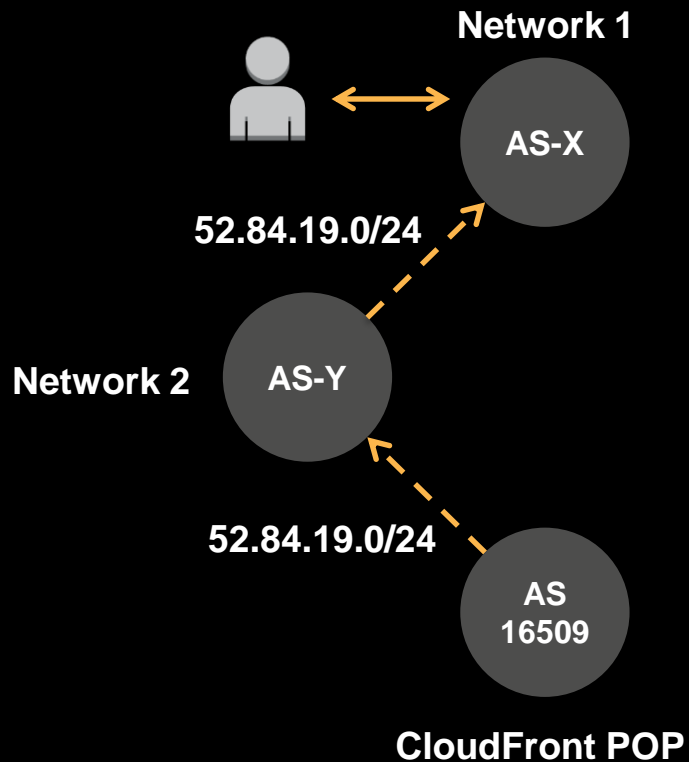
APAC




● AWS Region ● CloudFront Edge Location




How does Amazon CloudFront work?




How does Amazon CloudFront work?




Sequence of events



Sequence of events



Sequence of events




How to monitor availability




How CloudFront monitors availability

1. Analysis of server-side metrics




How CloudFront monitors availability

1. Analysis of server-side metrics
2. **Canaries**




How CloudFront monitors availability


1. Analysis of server-side metrics
2. Canaries
3. **Third-party global HTTP tests**



How CloudFront monitors availability



1. Analysis of server-side metrics
2. Canaries
3. Third party global HTTP tests



Availability interruption




Segfault in CloudFront DNS server


```
27: int dns_get_domain_length(const char *domain) { /* <== domain is NULL */
28:     const char *pos;
29:     unsigned char ch;
30:
31:     pos = domain;
32:     while ((ch = *pos++) != 0)          /* <== SEGFAULT */
33:         pos += (unsigned int) ch;
34:
35:     return pos - domain;
36: }
```




Potential causes of application failure




Potential causes of application failure




Potential causes of application failure




Potential causes of application failure




Potential causes of application failure




Potential causes of application failure



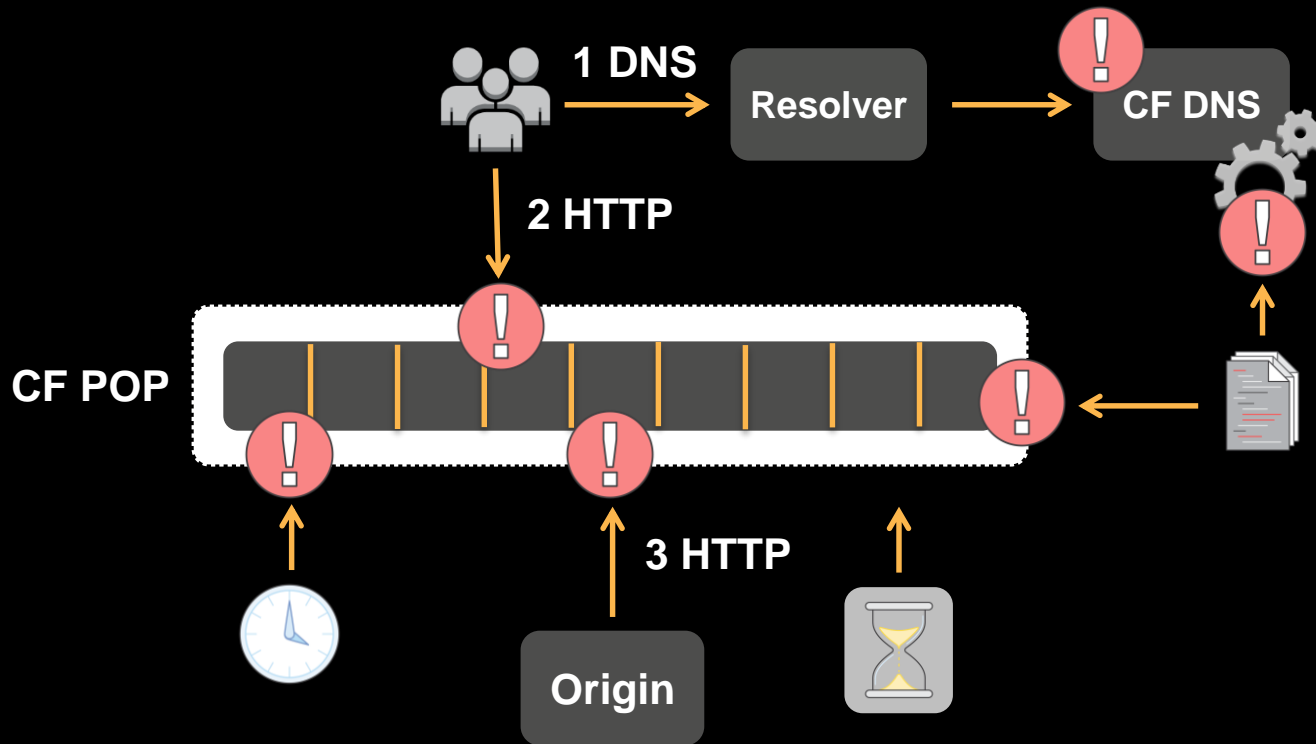
Potential causes of application failure



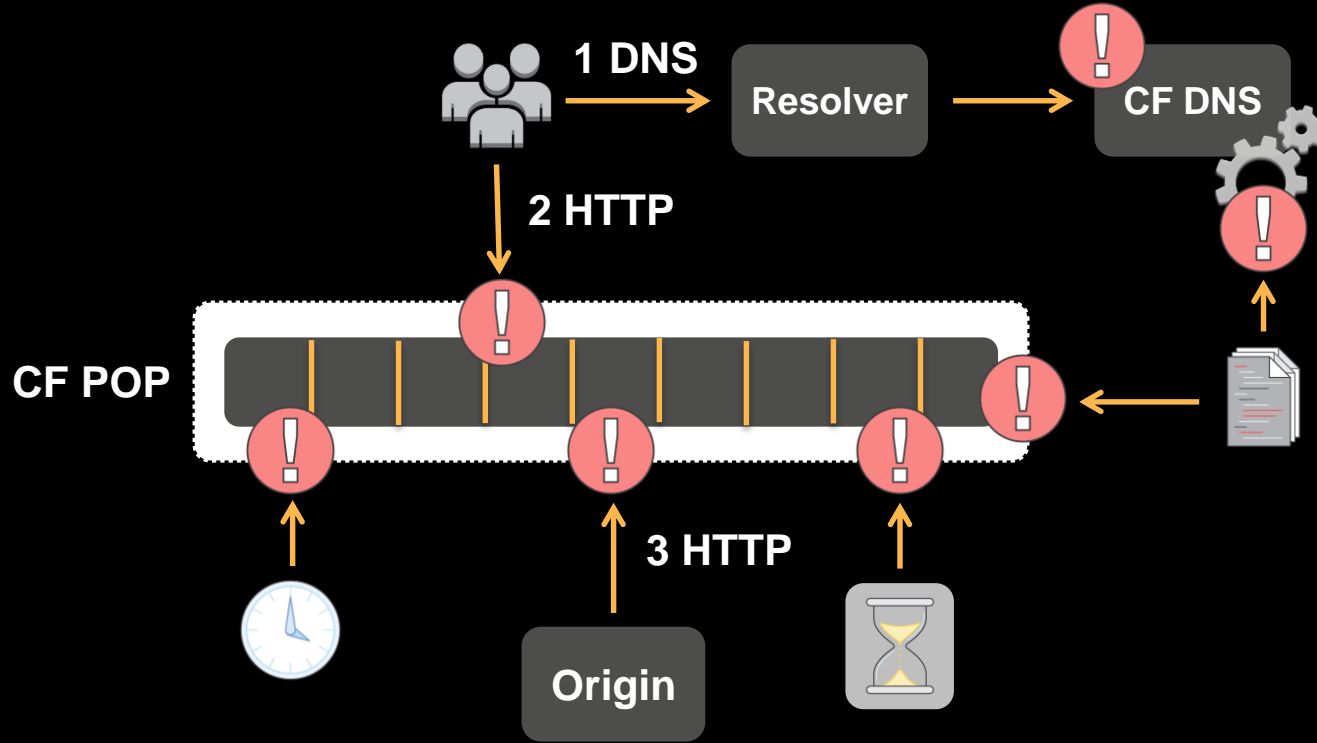
Potential causes of application failure




Potential causes of application failure



Potential causes of application failure



Risks to availability




Rapid Iteration on Capabilities




Design patterns for availability



Design patterns for availability



Maximizing
availability with
Food Tasting



Flash Crowds
without scaling
the peak




Defense in Depth
Strategies




Time Bomb Jitter
Protection

Design pattern 1: FoodTaster



**Maximizing
availability with
Food Tasting**



Flash Crowds
without scaling for
the peak



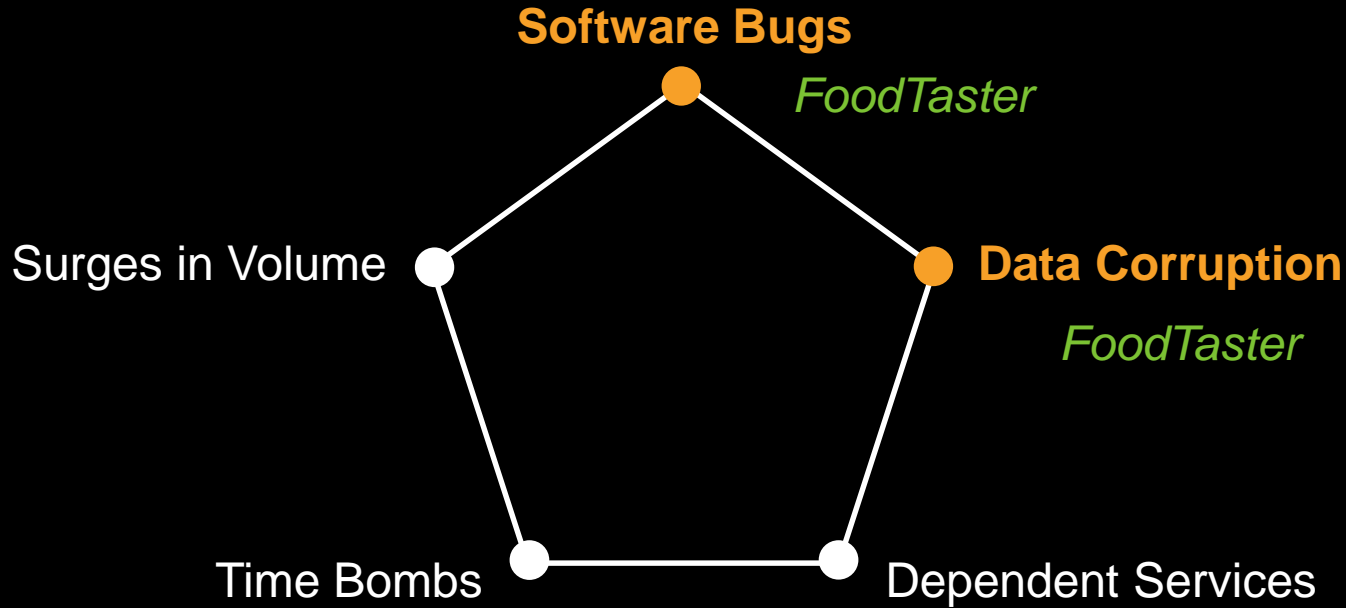
Defense in Depth
Strategies



Time Bomb Jitter
Protection

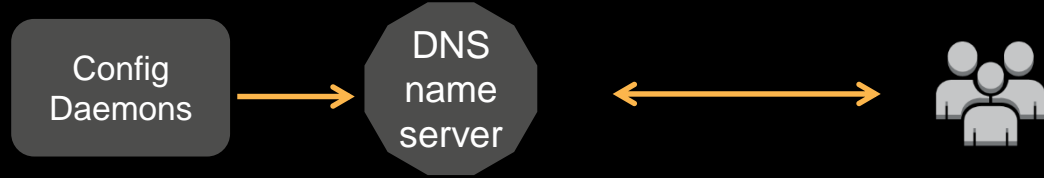


Risks to availability



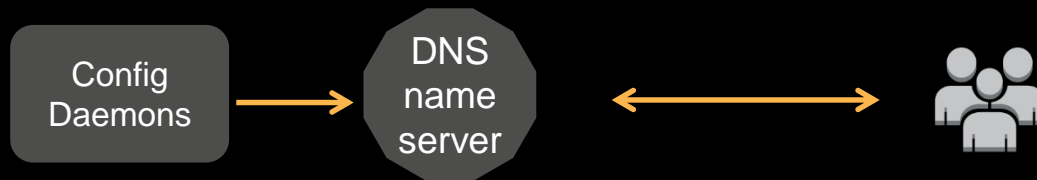
Praegustator – DNS FoodTaster

Before




Praegustator – DNS FoodTaster

Before



After



Praegustator – DNS FoodTaster

```
$ ls
```

```
dns/ foodtaster/ landing/
```

```
$ ls foodtaster/
```

```
customer.data@
```

```
pop.data@
```

```
resolvers.data@
```

```
routing.data
```

```
$ ls dns/
```

```
customer.data
```

```
pop.data
```


```
resolvers.data
```

```
routing.data
```



Food Tasting in AWS

- Straightforward approach
- Deployments are inexpensive, redeployments more too
- Approaches like CodeDeploy make life even easier



Food Tasting in AWS



- Example - **GeoIP Database**
- Automatically pushed out to every host
- Likely already have checks
- More to consider than just invalid user configuration

Food Tasting in AWS

- Simple is better (Works for Amazon CloudFront!)
- Assume we already use a deployment system (e.g., AWS CodeDeploy)
- Complete in-situ tests before returning complete



CodeDeploy AppSec example

hooks:

BeforeInstall:

- location: Scripts/UnzipResourceBundle.sh
- location: Scripts/UnzipDataBundle.sh

AfterInstall:

- location: Scripts/RunResourceTests.sh
- timeout: 180

ApplicationStart:

- location: Scripts/RunFunctionalTests.sh
- timeout: 3600

validateService:

- location: Scripts/MonitorService.sh
- timeout: 3600
- runas: codedeployuse

CodeDeploy AppSec example

hooks:

BeforeInstall:

- location: Scripts/UnzipResourceBundle.sh
- location: Scripts/UnzipDataBundle.sh

AfterInstall:

- location: Scripts/RunResourceTests.sh
timeout: 180


ApplicationStart:

- location: Scripts/RunFunctionalTests.sh
timeout: 3600

validateService:


- location: Scripts/MonitorService.sh
timeout: 3600
runas: codedeployuse

Food Tasting in AWS




- Acts as a quality gate
- Roll-back can be automatic
- Verification never affects user facing traffic

Design pattern 2: Flash Crowds



Maximizing
availability with
Food Tasting



Flash Crowds
without scaling for
the peak



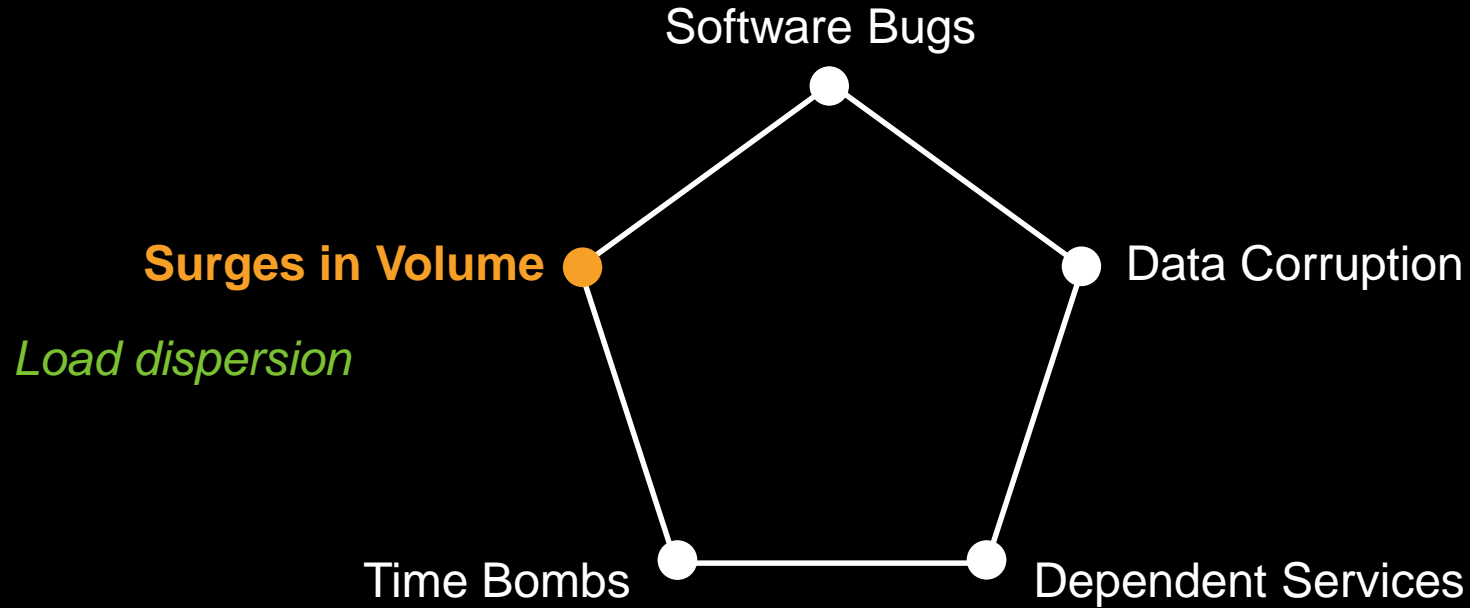
Defense in Depth
Strategies




Time Bomb Jitter
Protection




Risks to availability




HTTP Flash Crowds



Static & Dynamic Strategies




Packet-flow config



Feedback loop




Load Dispersal




Flash Crowds within AWS

Auto Scaling works really well


But.



TPS (Prod Cluster)



TPS (Prod Cluster)




Problem: Flash Crowds

Auto Scaling works really well

But.

Sometimes, 60s is too long.



AWS solutions


- **Plan Ahead**



Planning ahead

Typical examples:

- TV programs
- Live events (sports)
- Game releases
- Established traffic patterns



Planning ahead

Typical examples:

- TV programs
- Live events (sports)
- Game releases
- Established traffic patterns

AWS Solutions:



Infrastructure Event Management




Scheduled Auto Scaling Group



Auto Scaling Integration

Schedules and the Big Red Button


- Integration of **program scheduling** and Auto Scaling
- Program scheduling includes expected online audience parameters
- Big red button



Schedules and the Big Red Button


- Integration of program scheduling and Auto Scaling
- Program scheduling includes expected online audience parameters
- Big red button

I CAN'T
PLAN
AHEAD!



AWS solutions

- Plan Ahead
- **Cache Things**



Cache things



Cache things

I can't!

My website is...



Cache things

I can't!

My website is...

DYNAMIC!



Cache things

- What's really dynamic?
 - Newspapers
 - Voting sites
 - Forum sites
- Updates don't need to be more than once a second




Cache things (even for a second)

- 10000s req/s
- Assuming every Amazon CloudFront POP
- 10000s req/s -> ~ 10s req/s



AWS solutions

- Plan Ahead
- Cache Things
- **Serve only what you have to**



Serve only what you have to

Menu



AWS re:Invent

Products ▾

Solutions

Pricing

More ▾

English ▾

My Account ▾

Sign In to the Console

Introducing AWS Server Migration Service
Migrate on-premise servers to AWS easier and faster

[Learn more »](#)



Manage Your Resources

[Sign In to the Console](#)

AWS Console Mobile App

View your resources on iOS
and Android devices

[Download the Mobile App »](#)

Serve only what you have to

Menu



AWS re:Invent

Products ▾

Solutions

Pricing

More ▾

English ▾

My Account ▾

Sign In to the Console

Introducing AWS Server Migration Service
Migrate on-premise servers to AWS easier and faster

[Learn more »](#)



Manage Your Resources

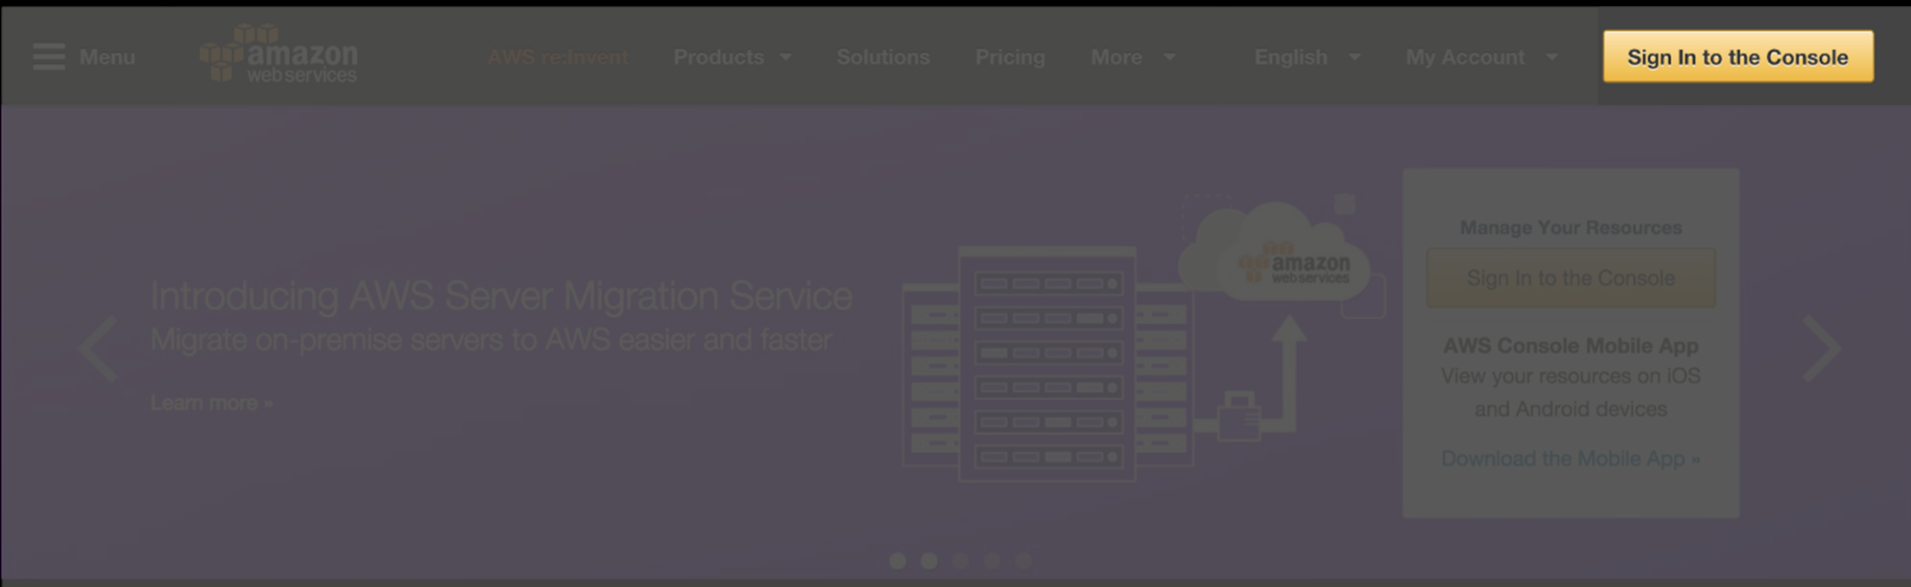
Sign In to the Console

AWS Console Mobile App

View your resources on iOS
and Android devices


[Download the Mobile App »](#)

Serve only what you have to




Serve only what you have to

- AJAX Includes




Serve only what you have to


- AJAX Includes
- CloudFront Cache Keys



Design pattern 3: Defense in Depth



Maximizing
availability with
Food Tasting



Flash Crowds
without scaling for
the peak



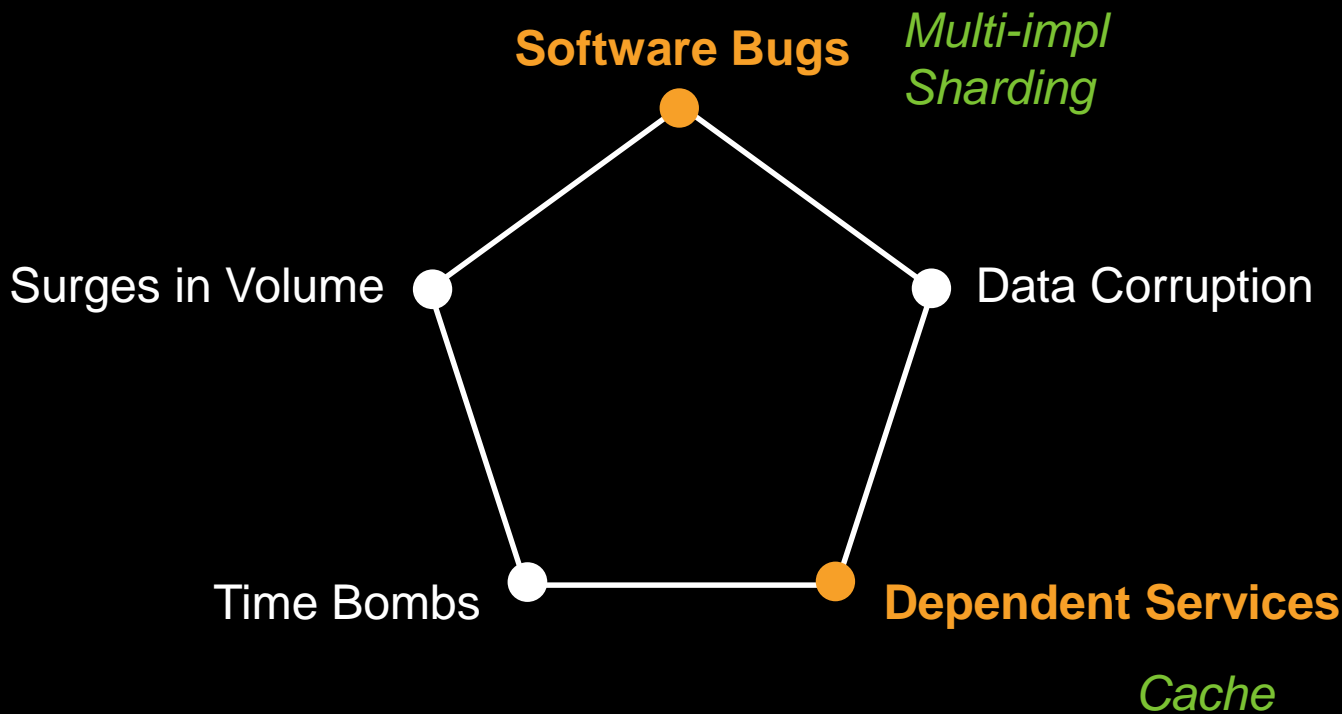
**Defense in Depth
Strategies**




Time Bomb Jitter
Protection



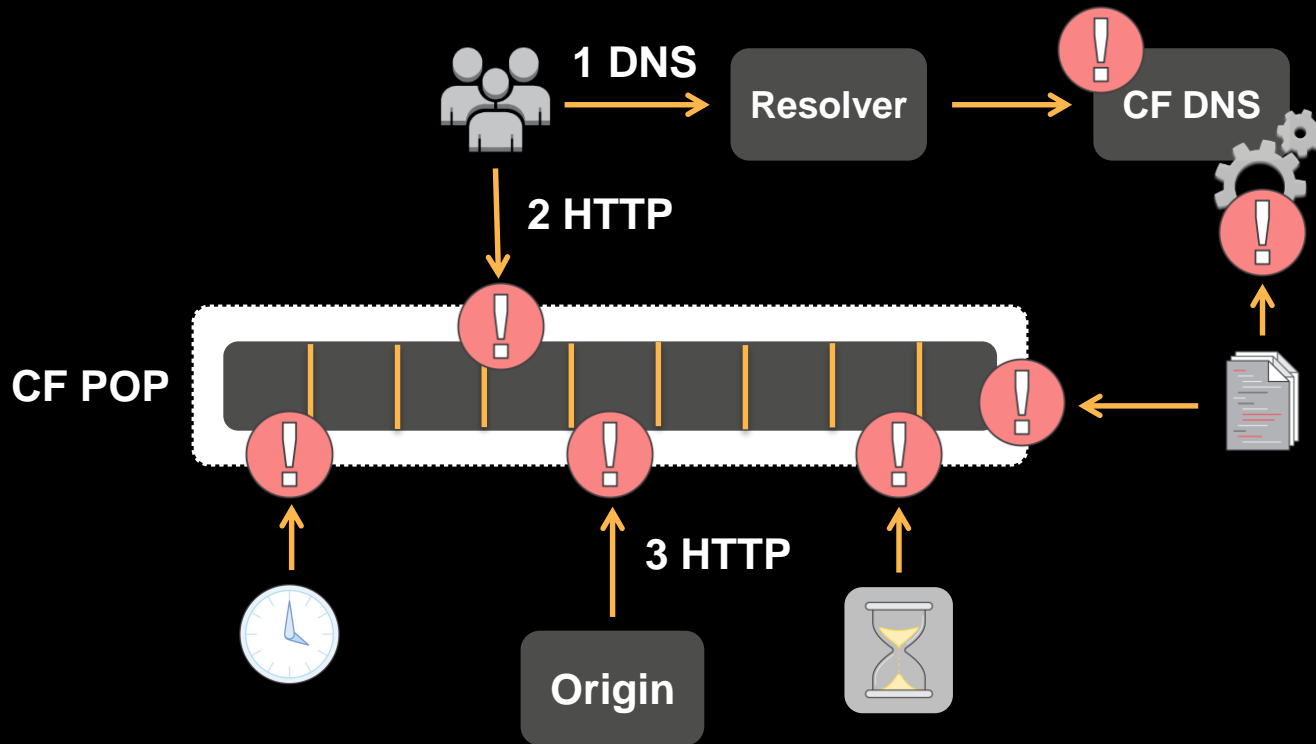
Risks to availability



Potential causes of application failure



Potential causes of application failure



Ideas to avoid crashing



- Comprehensive test coverage
- Simplify systems

Failures are going to happen.

How can we survive them when they do?

Ideas to survive crashing



- Reduce blast radius
- Reject input that previously made you crash



Ideas to survive crashing



- Reduce blast radius
 - shard customers to separate processes
- Reject input that previously made you crash



Ideas to survive crashing



- Reduce blast radius
 - shard customers to separate processes
 - recover quickly
- Reject input that previously made you crash



Ideas to survive crashing



- Reduce blast radius
 - shard customers to separate processes
 - recover quickly
 - multiple implementations
- Reject input that previously made you crash



DNS Multi-Imp

Should this be a fallback system?

Place it in front of every DNS name server?



DNS Multi-Imp



Should this be a fallback system?

- We want to know it always works

Place it in front of every DNS name server?

DNS Multi-Imp



Should this be a fallback system?


- We want to know it always works



Place it in front of every DNS name server?


- Would trade point of failure one for another

DNS Stripes





DNS Stripes

Stripe 1
(34 POPs)



Stripe 2
(34 POPs)



CloudFront case study

`Math.abs()` on the lowest 32-bit signed integer, -2^{31} , yields a negative number.
Leading to invalid DNS configuration.

Two's complement: flip bits and add 1.

`Math.abs(Integer.MIN_VALUE):`

$-2^{31} = 10000000\ 00000000\ 00000000\ 00000000$

bits flipped = $01111111\ 11111111\ 11111111\ 11111111$

+1 = $10000000\ 00000000\ 00000000\ 00000000 \leq -2^{31}!$



CloudFront case study

Protections in place:


1. Config process crashing in a loop
2. DNS name server defensively ignores an invalid index
3. FoodTaster crashes protect King DNS name server
4. DNS stripes reduce common failure paths



AWS solution

We can use a Load Balancer with Layer-7 awareness

So we can have multiple backend types across a fleet




AWS solution

Enabled through microservices

Reimplementing the entire stack multiple times is a difficult cost / efficiency question

Specific high-risk services are easier



AWS solution

Example: **2FA Platform**

API (HTTP) front-end to a
back-end authentication
platform

Reasonably few SLOC

If it fails – no access to
platform

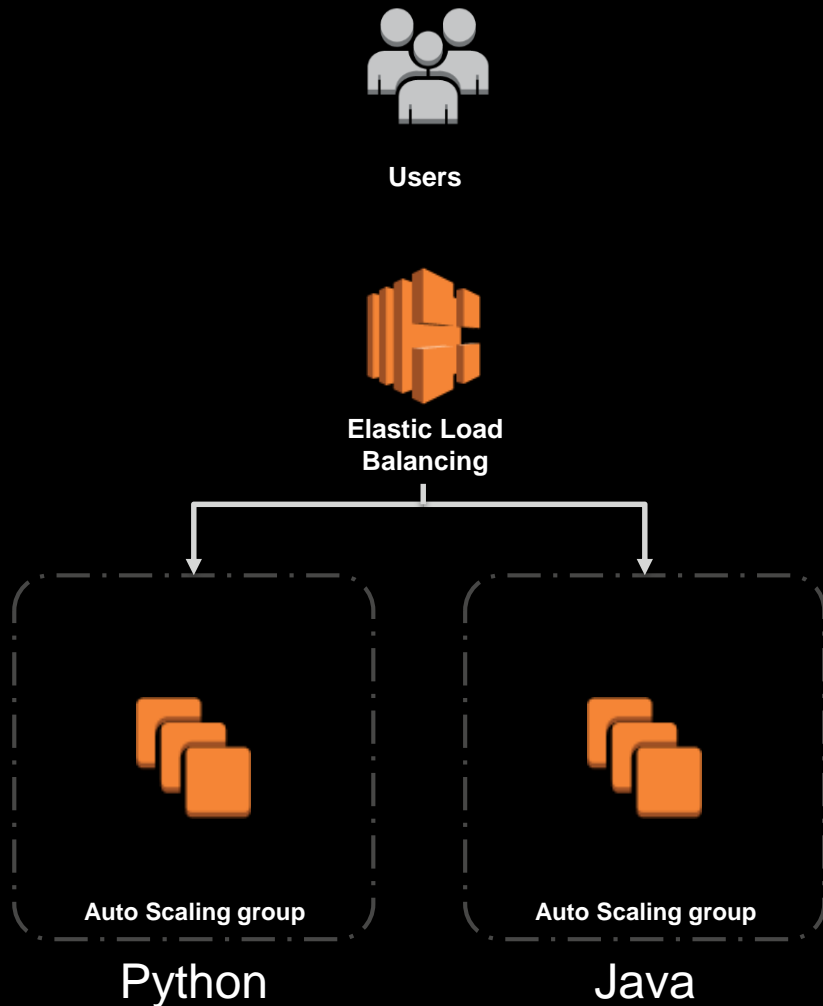
AWS solution

Example: **2FA Platform**


API (HTTP) front-end to a
back-end authentication
platform

Reasonably few SLOC


If it fails – no access to
platform



Design Pattern 4: Time Bomb Jitter Protection



Maximizing
availability with
Food Tasting



Flash Crowds
without scaling for
the peak

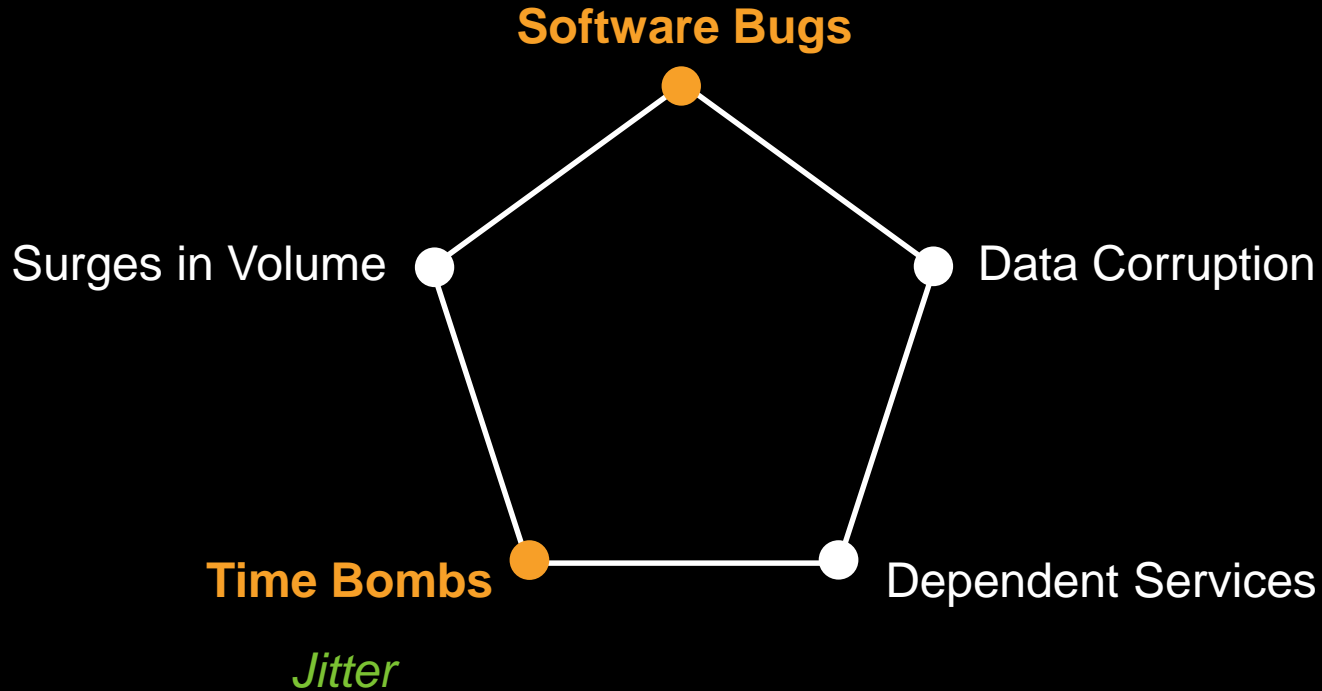


Defense in Depth
Strategies




**Time Bomb Jitter
Protection**


Risks to availability




Problem: homogeneous platforms




The problem




The problem




The problem



The problem




The problem




Traditionally..

- Instance level monitoring system with alerts
- Human response
- Monitor both percentage fill, and the fill rate



High availability matters




High availability matters


Recycling old instances is
unnecessary cost

Automated cleanups may be
too slow


Human operators **almost**
certainly too slow



The solution



The solution




AWS solution – within an Auto Scaling group


Amazon Linux AMI has
`/var/run` (and `/var/tmp`) on root

Ubuntu AMI uses tmpfs (10%
of RAM)

Jittering **root volume size** is
tricky within an ASG



AWS solution – within Auto Scaling



Consciously segregate your files to a specified directory – on a **separate volume** – with jittering

Code will follow (check SlideShare!)

Adds <10s to system startup

What else?

Time

- SSL Certificates
- Domain name registrations
- Deployment Schedules
- ???




Wrap up:

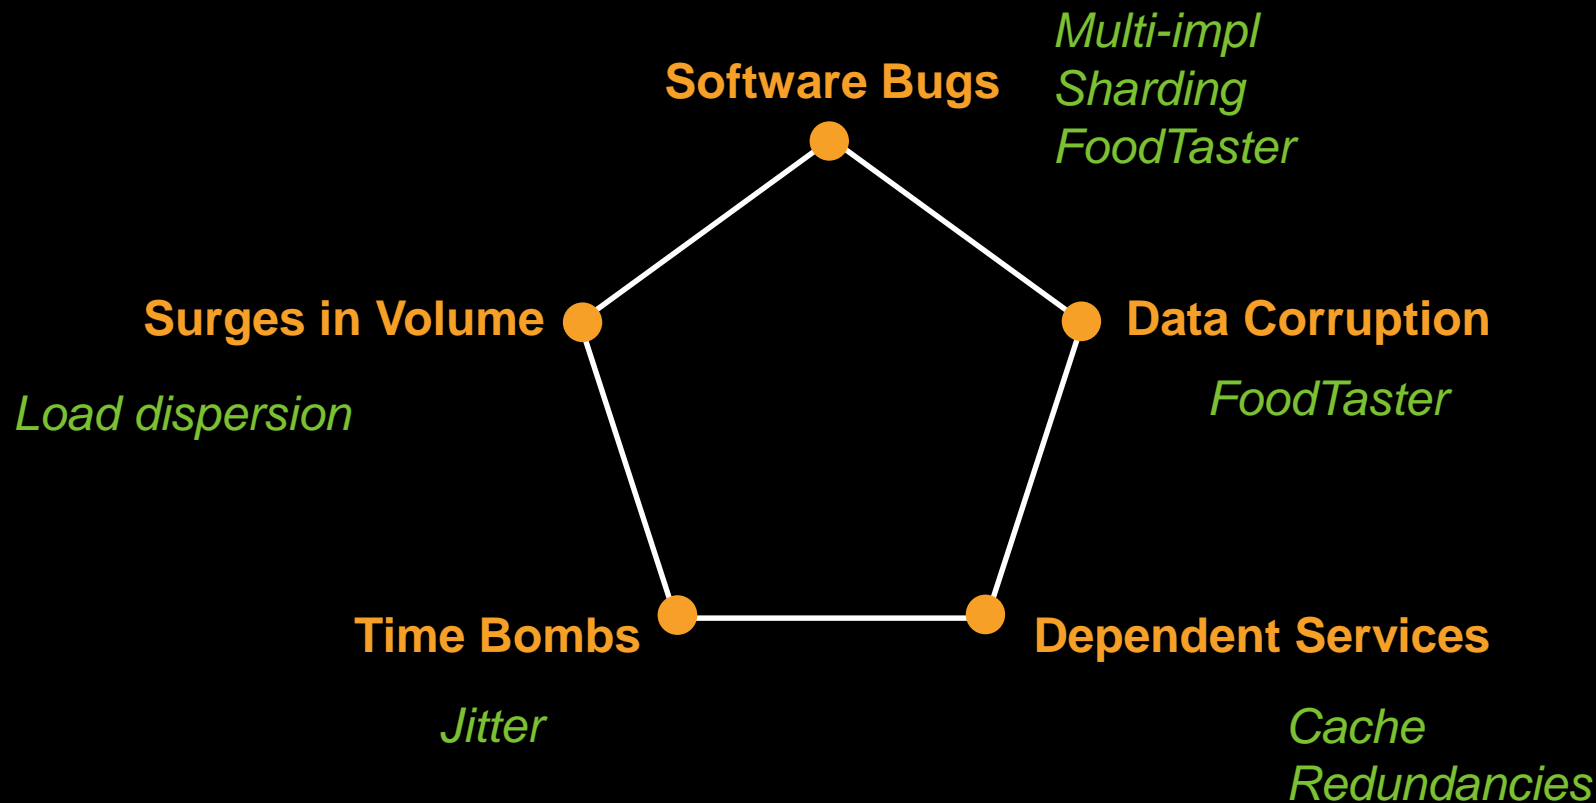
Lessons learned




Risks to availability



Risks & Common Patterns




Lessons learned



Maximizing
availability with
Food Tasting

**Don't rely on
validation in your
main application**



Flash Crowds
without scaling for
the peak

Auto Scaling
Integration;
Caching; Selective
Serving



Defense in Depth
Strategies


Implementation
striping can save
you



Time Bomb Jitter
Protection


Homogeneity can
hurt

Lessons learned



Maximizing
availability with
Food Tasting

Don't rely on
validation in your
main application



Flash Crowds
without scaling for
the peak

Auto Scaling
Integration;
Caching;
Selective Serving



Defense in Depth
Strategies


Implementation
striping can save
you



Time Bomb Jitter
Protection


Homogeneity can
hurt

Lessons learned



Maximizing
availability with
Food Tasting

Don't rely on
validation in your
main application



Flash Crowds
without scaling for
the peak

Auto Scaling
Integration;
Caching; Selective
Serving



Defense in Depth
Strategies


**Implementation
striping can save
you**



Time Bomb Jitter
Protection


Homogeneity can
hurt

Lessons learned



Maximizing
availability with
Food Tasting

Don't rely on
validation in your
main application



Flash Crowds
without scaling for
the peak

Auto Scaling
Integration;
Caching; Selective
Serving



Defense in Depth
Strategies

Implementation
striping can save
you



Time Bomb Jitter
Protection

**Homogeneity can
hurt**



AWS
re:Invent

Thank you!



**Remember to complete
your evaluations!**

Related Sessions

ARC309

Moving Mission Critical Apps from One Region to Multi-Region active/active

ARC204

From Resilience to Ubiquity - #NetflixEverywhere
Global Architecture