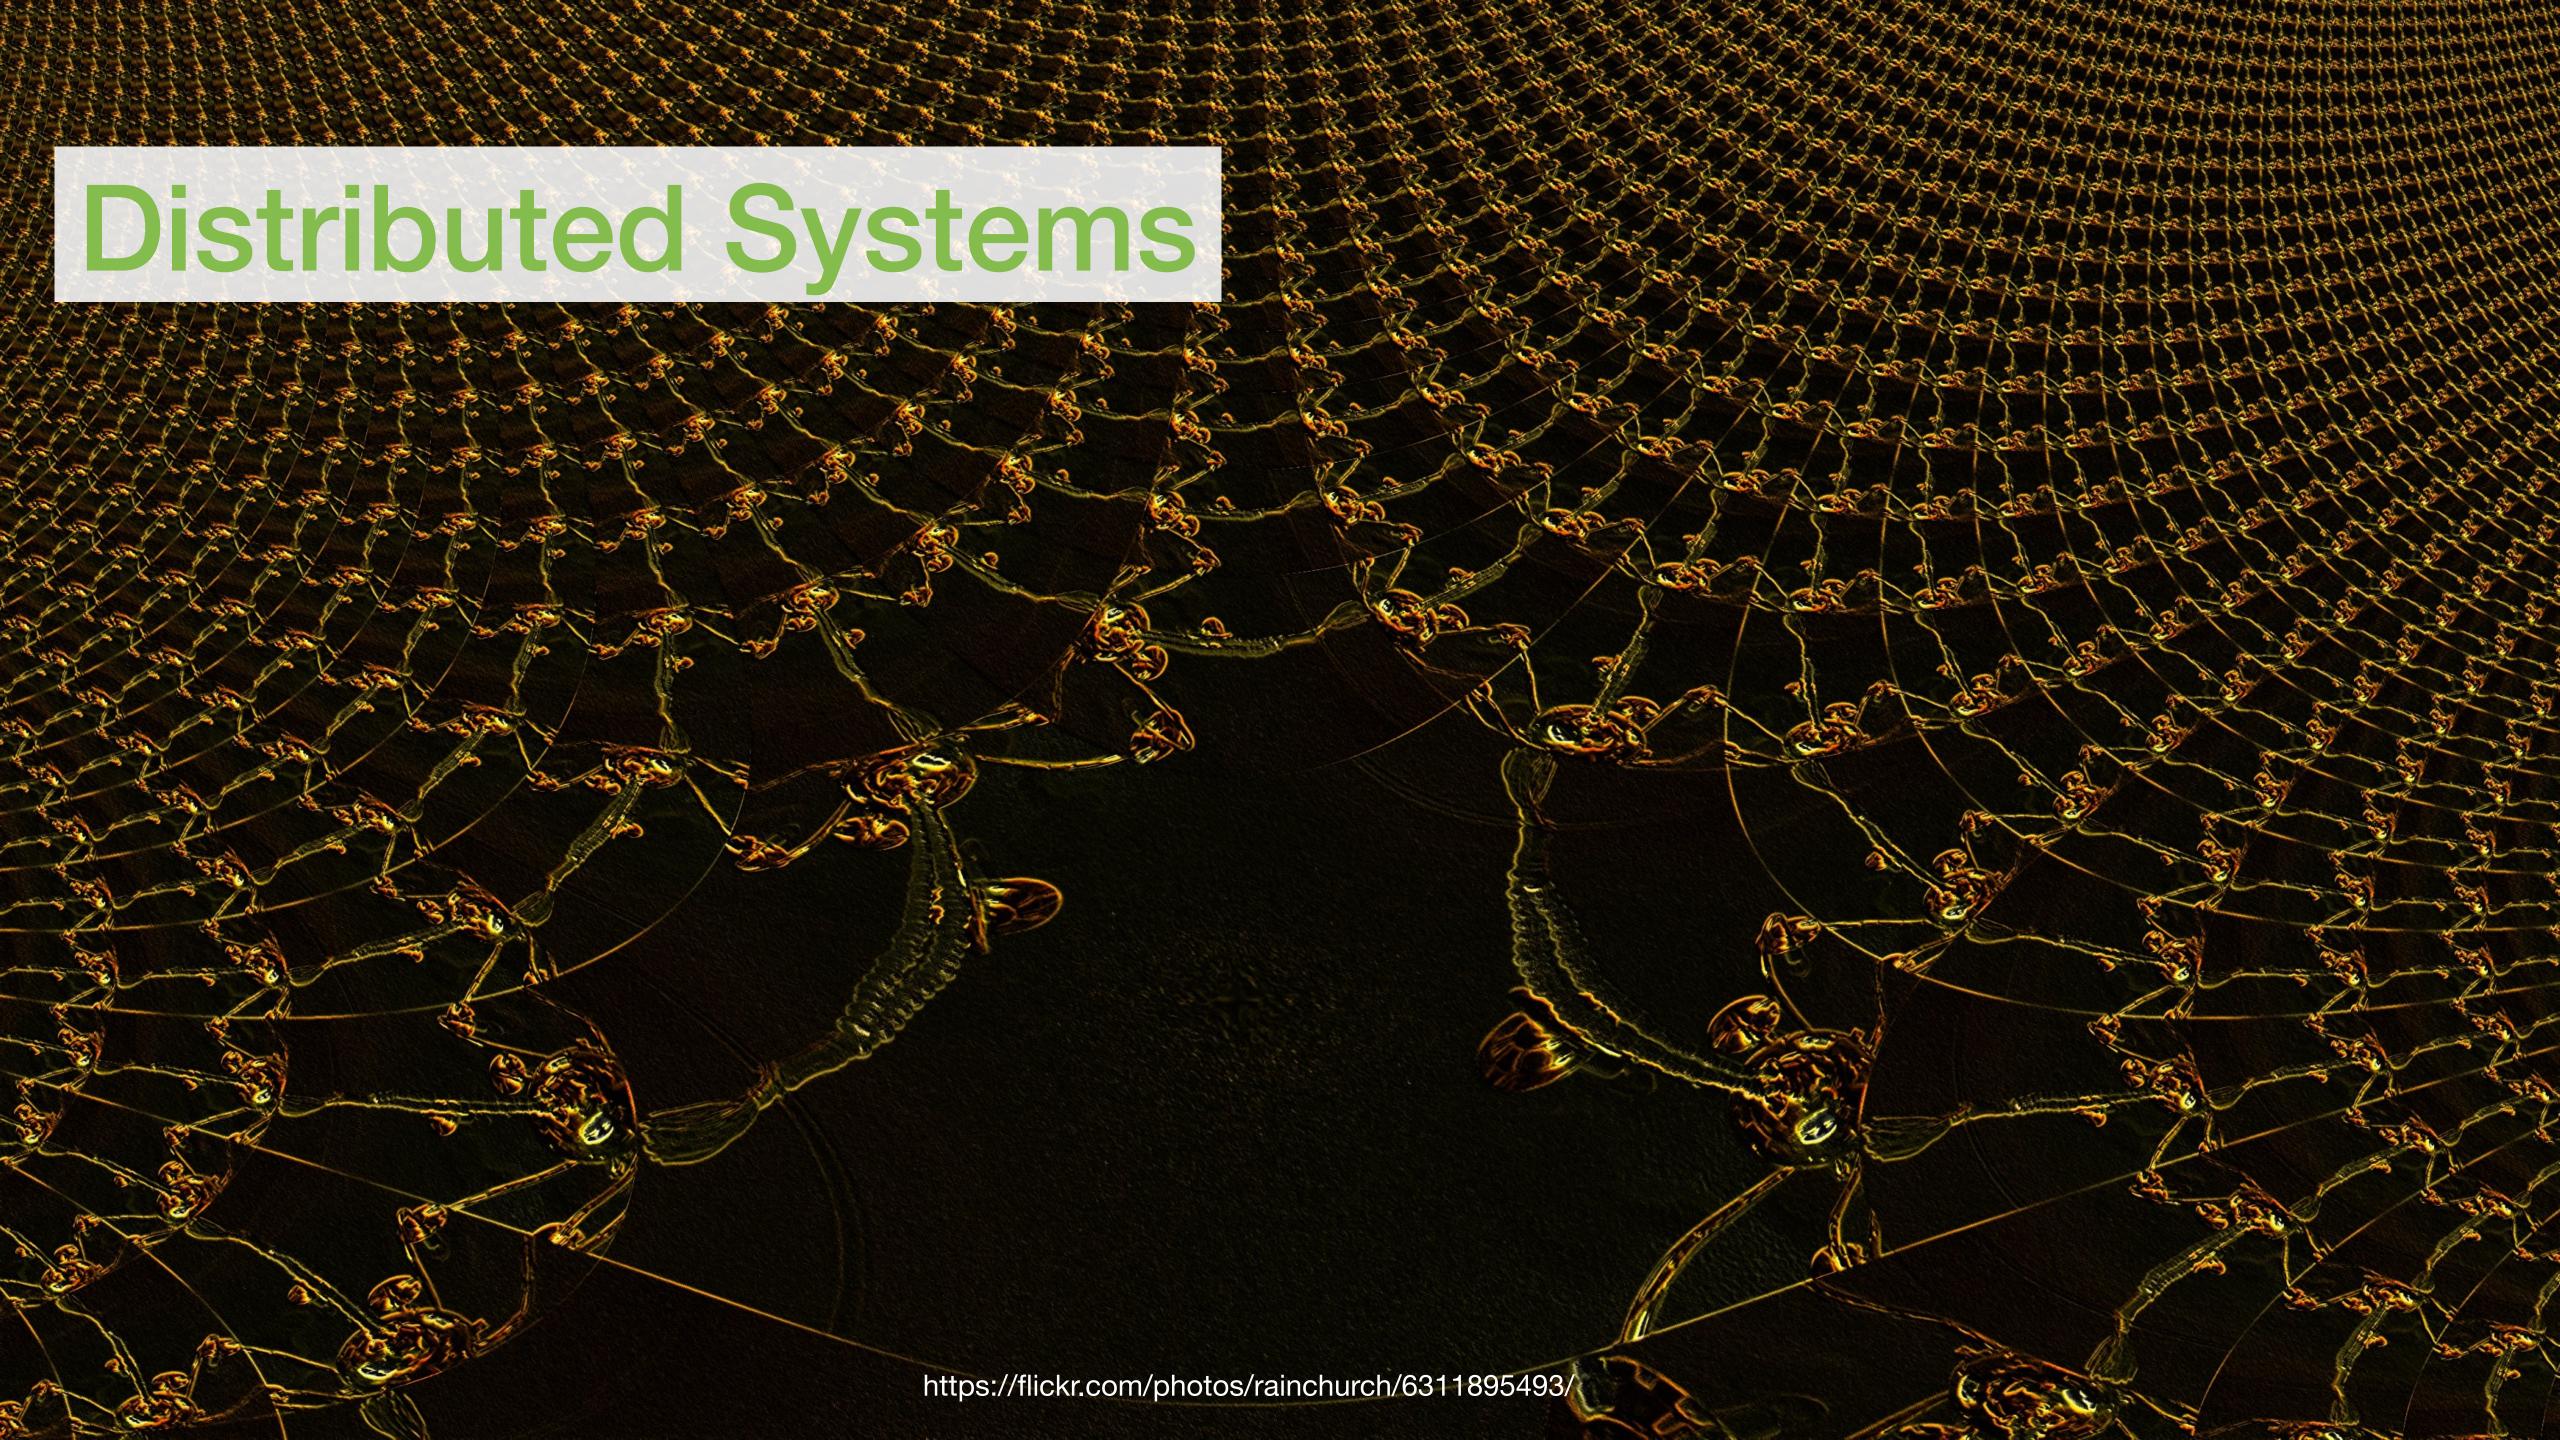


Retries

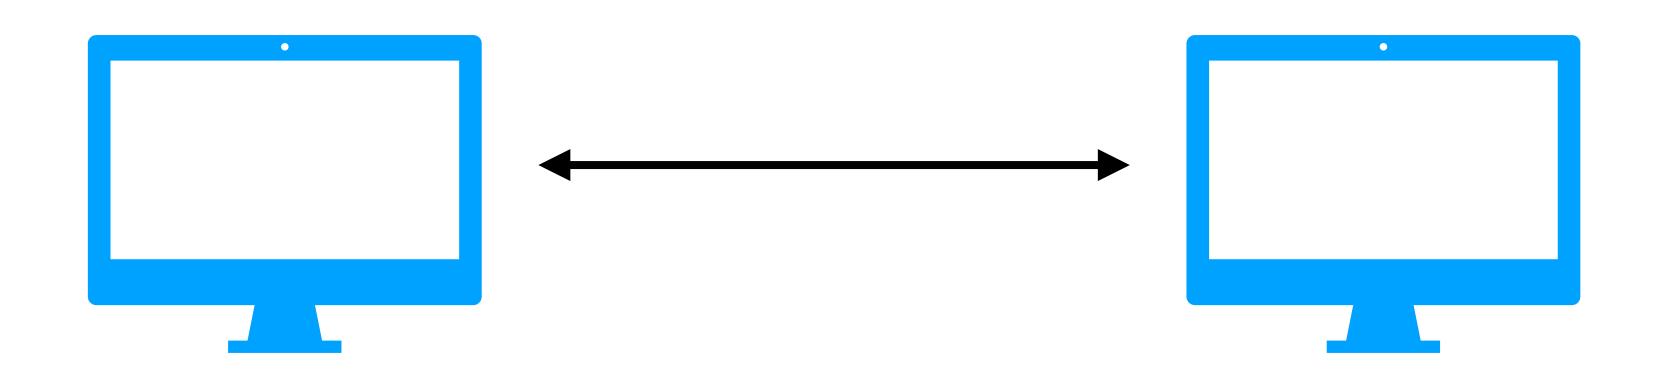






Distributed System

A system which consists of two or more computers connected by networks





"A distributed system is one in which the failure of a computer you didn't even know existed can render your own computer unusable."

- Leslie Lamport

Amazon Web Services' US-EAST-1 r trouble again, with EC2 and contain impacted

Major AWS outage took down Fortnite, impacted Alexa, Snapchat, and more Internal dependencies again prove problematic / Amazon says th Simon Sharwood Wed 29 Oct 2025 // 01:16 U7 AWS outage was в в с **UPDATED** Amazon Web Services' US-VITE disruption to online services, is having a amazon alexa Home News Sport Business Innovation Culture Arts Travel Earth Audio Video Live prove problematic. Amazon Alexa At 3:36 PM PDT on October 28 (10:36F Amazon apologises to customers that "Earlier today some EC2 launches impacted by huge AWS outage experienced increased latencies for EC aws 23 October 2025 Amazon throttled some requests for EC Share **<** Save □ resolve the issue. Liv McMahon Amazon Web Services Technology reporter Another impact of the incident created (by + Jess Weatherbed Service] ECS tasks for both EC2 and F Updated Oct 20, 2025, 6:41 PM GMT+2 US-EAST-1 Region." 130 Comments (All the outage. Image: Downdetector This site can't be reached d on creative industries, computing, and internet culture. Jess ws and hardware reviews.

+ NEWS + TECH + AMAZON

+ NEWS + TECH + AMAZON

outages. Image: Eight Sleep

Eight Sleep adds 'outage mode' to smart beds after AWS problems left them frozen



/ Thousands of I owners were una temperatures or while while Amaz were down.

by + Jess Weatherbed
Oct 22, 2025, 5:28 PM GMT+2







Comments (All

https://www.theverge.com/news/804289/eight-sleep-smart-bed-aws-outage-overheating-offline

"A distributed system is one in which the failure of a computer you didn't even know existed can render your own computer unusable."

mattress

- Leslie Lamport



THE THREE GOLDEN RULES OF DISTRIBUTED COMPUTING

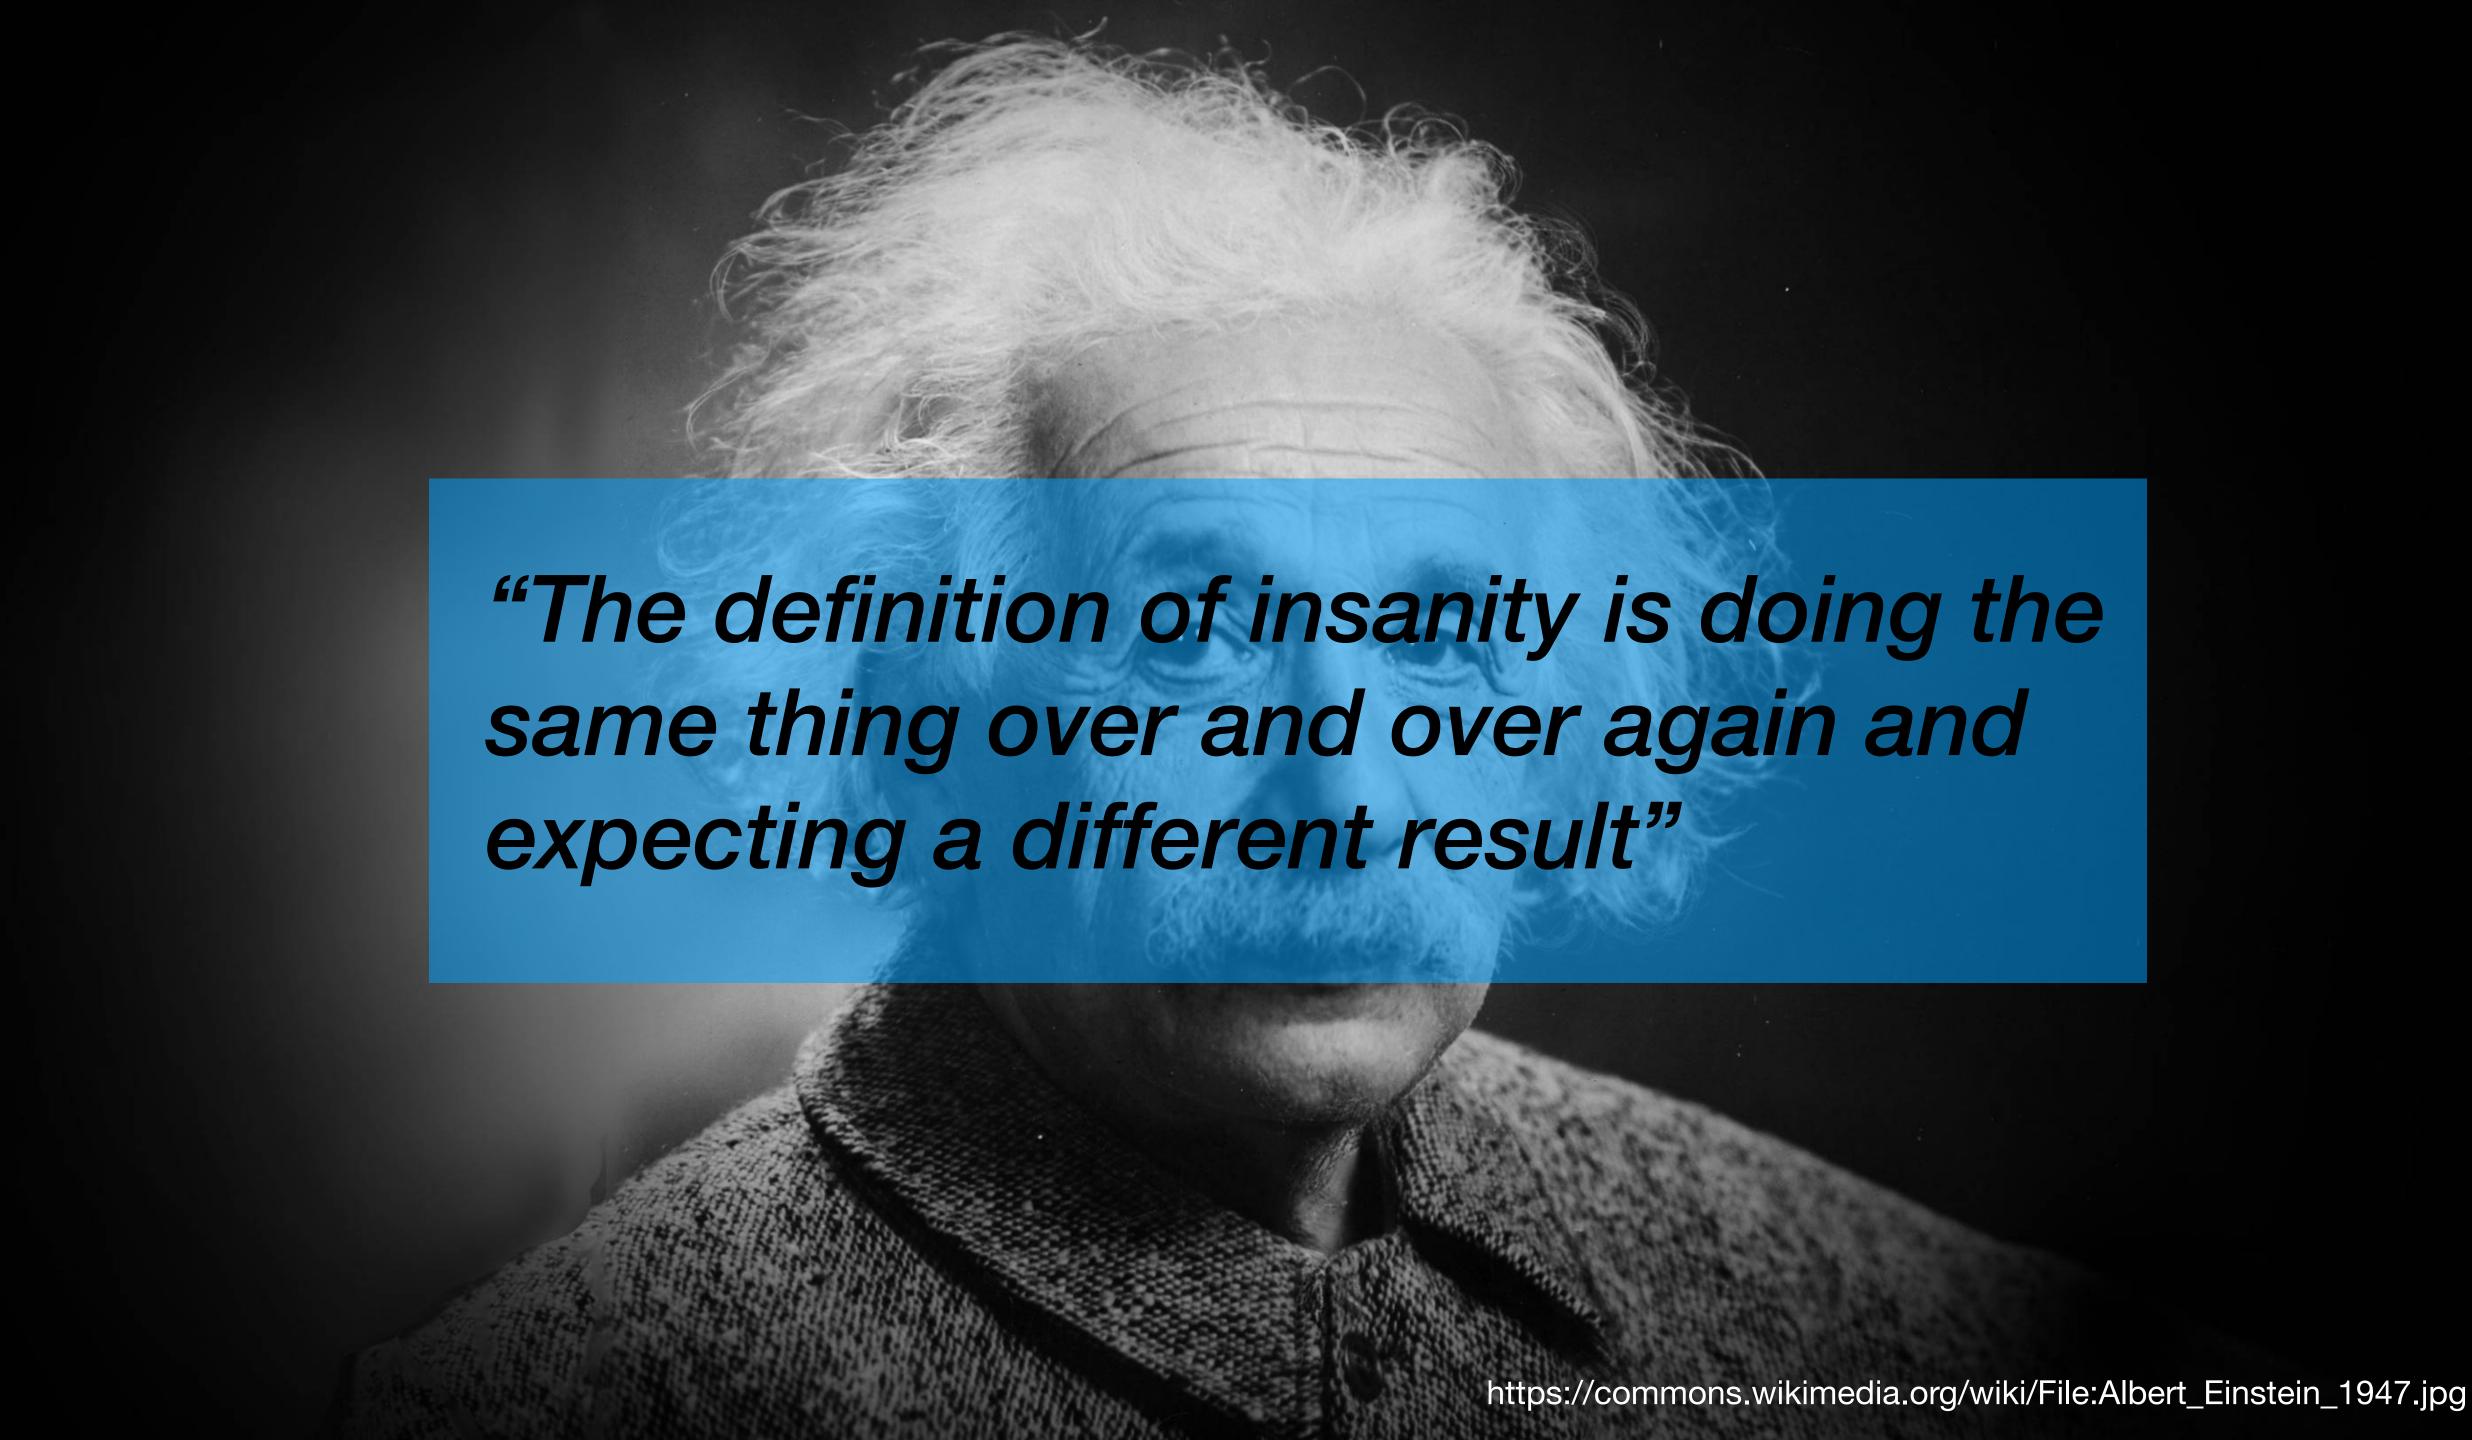
1. You can't beam information between two points instantly

2. Sometimes, you can't reach the thing you want to talk to

3. Resources are not infinite

Retries



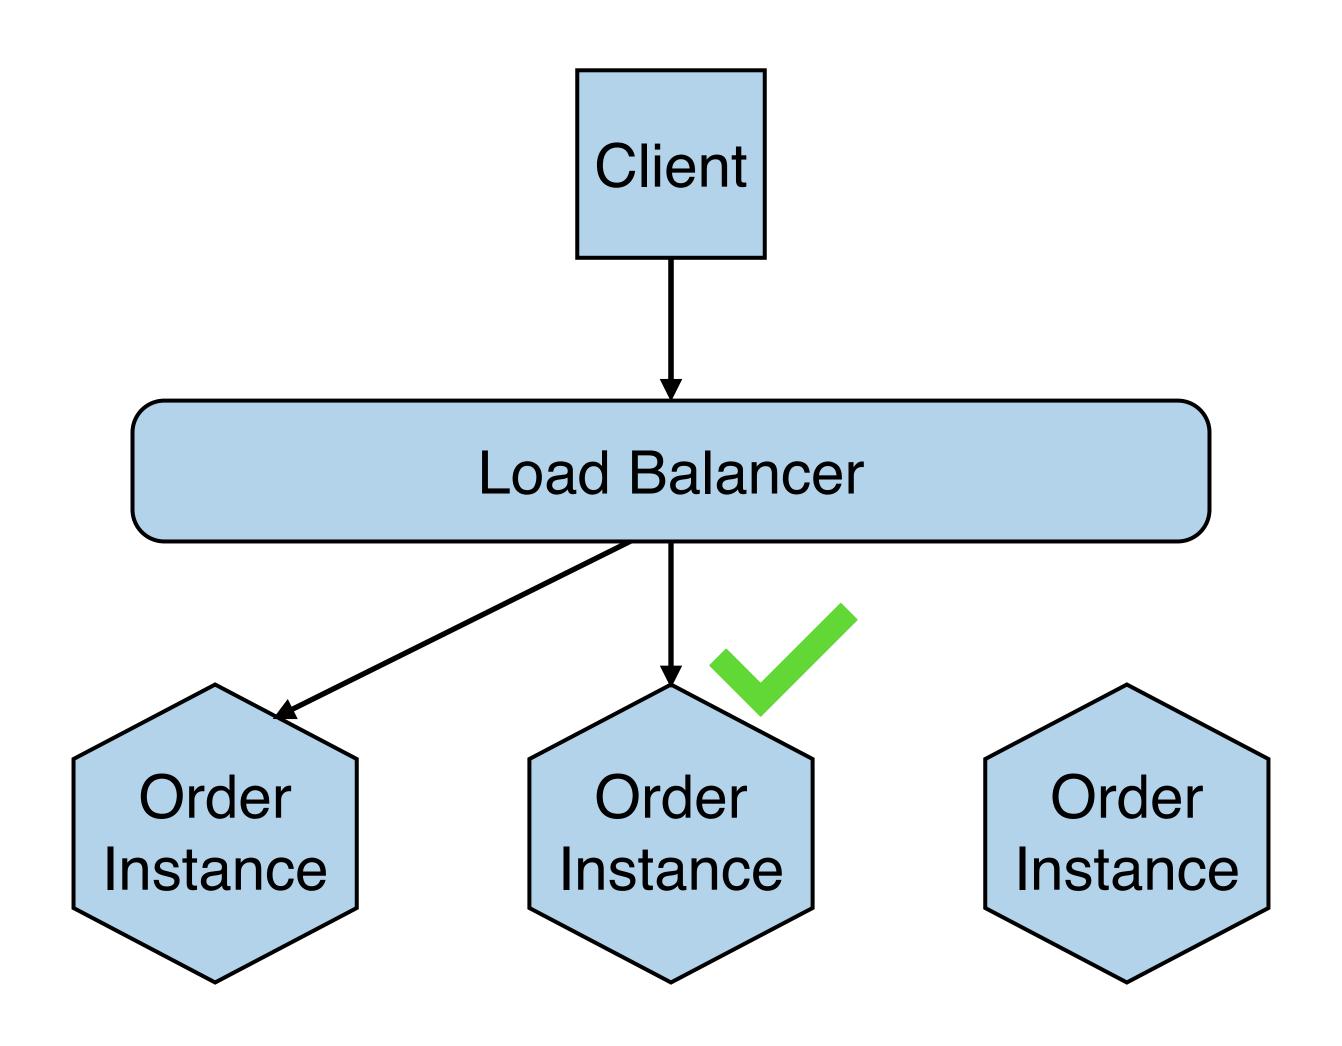


THE THREE GOLDEN RULES OF DISTRIBUTED COMPUTING

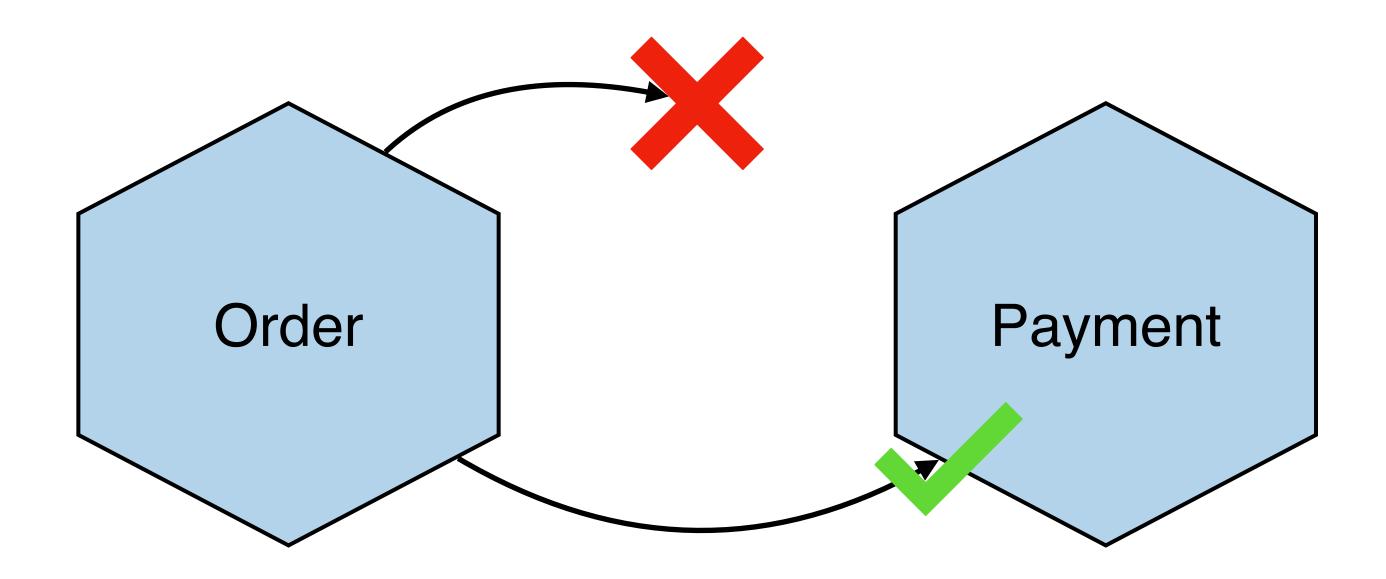
1. You can't beam information between two points instantly

2. Sometimes, you can't reach the thing you want to talk to

3. Resources are not infinite



So if the first attempt fails...



It might be worth trying again

How many retries?

THE THREE GOLDEN RULES OF DISTRIBUTED COMPUTING

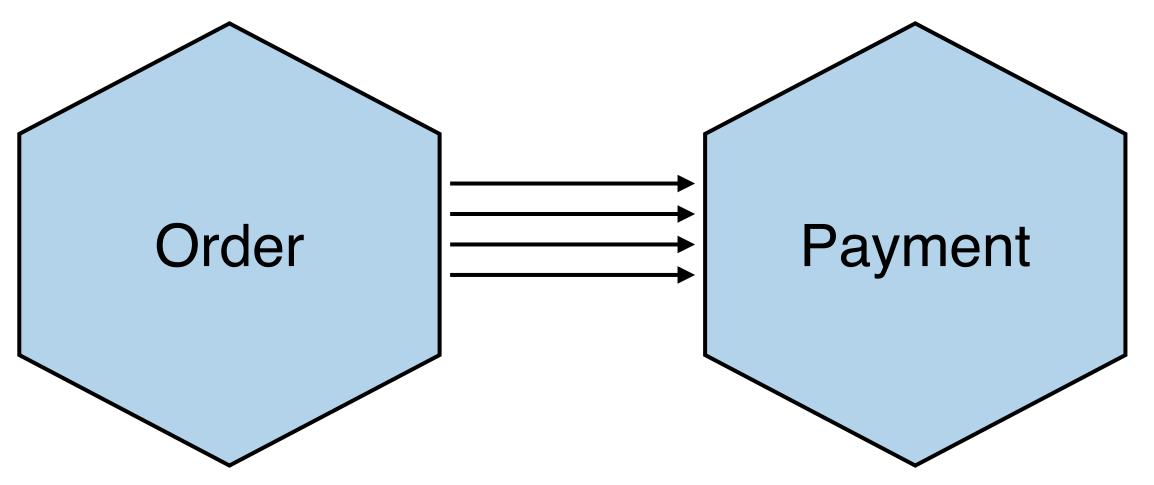
1. You can't beam information between two points instantly

2. Sometimes, you can't reach the thing you want to talk to

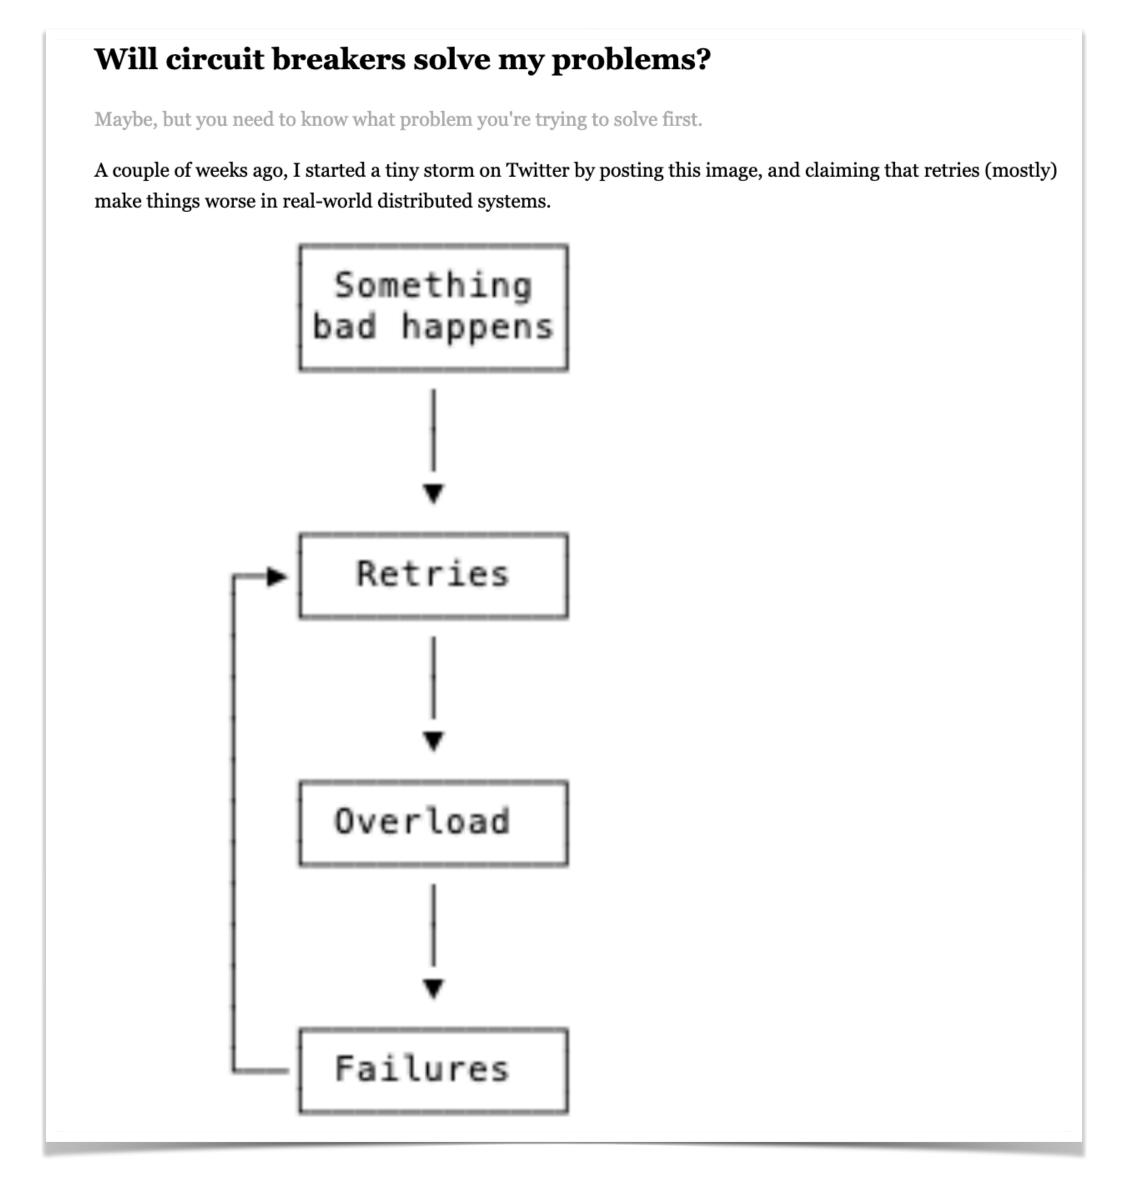
3. Resources are not infinite

Retrying too often can overload a server

Setting a max number of retries makes sense



And load shedding is a sensible defence here



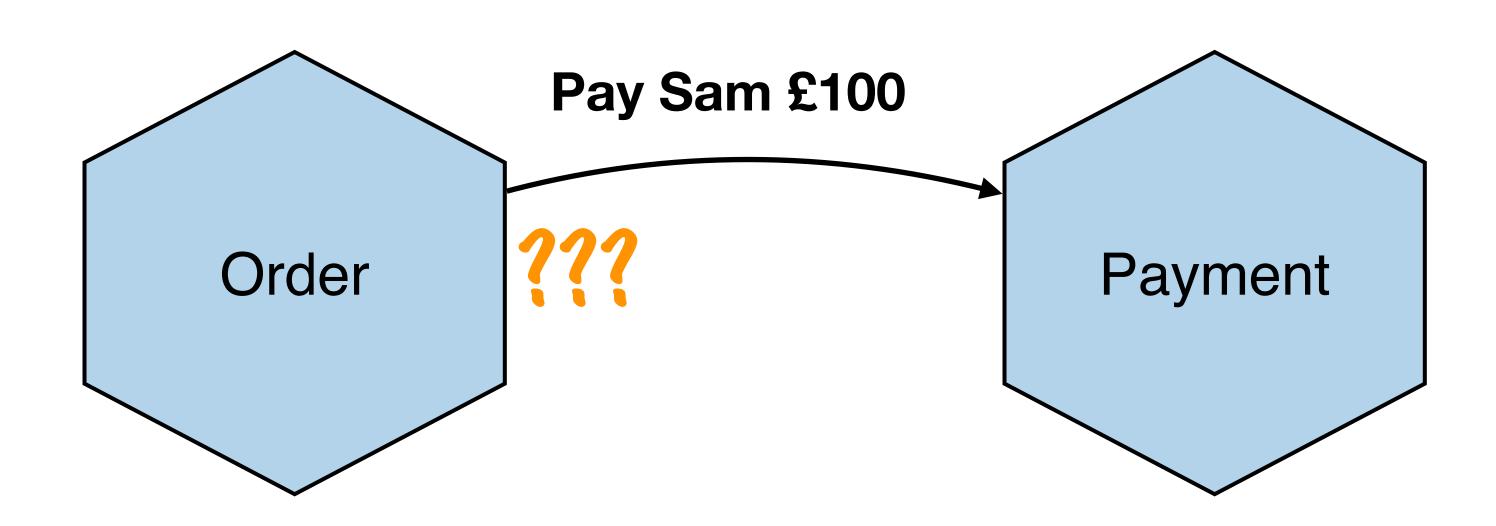
https://brooker.co.za/blog/2022/02/16/circuit-breakers.html



https://medium.com/square-corner-blog/incident-summary-2017-03-16-2f65be39297

Is it safe to retry?

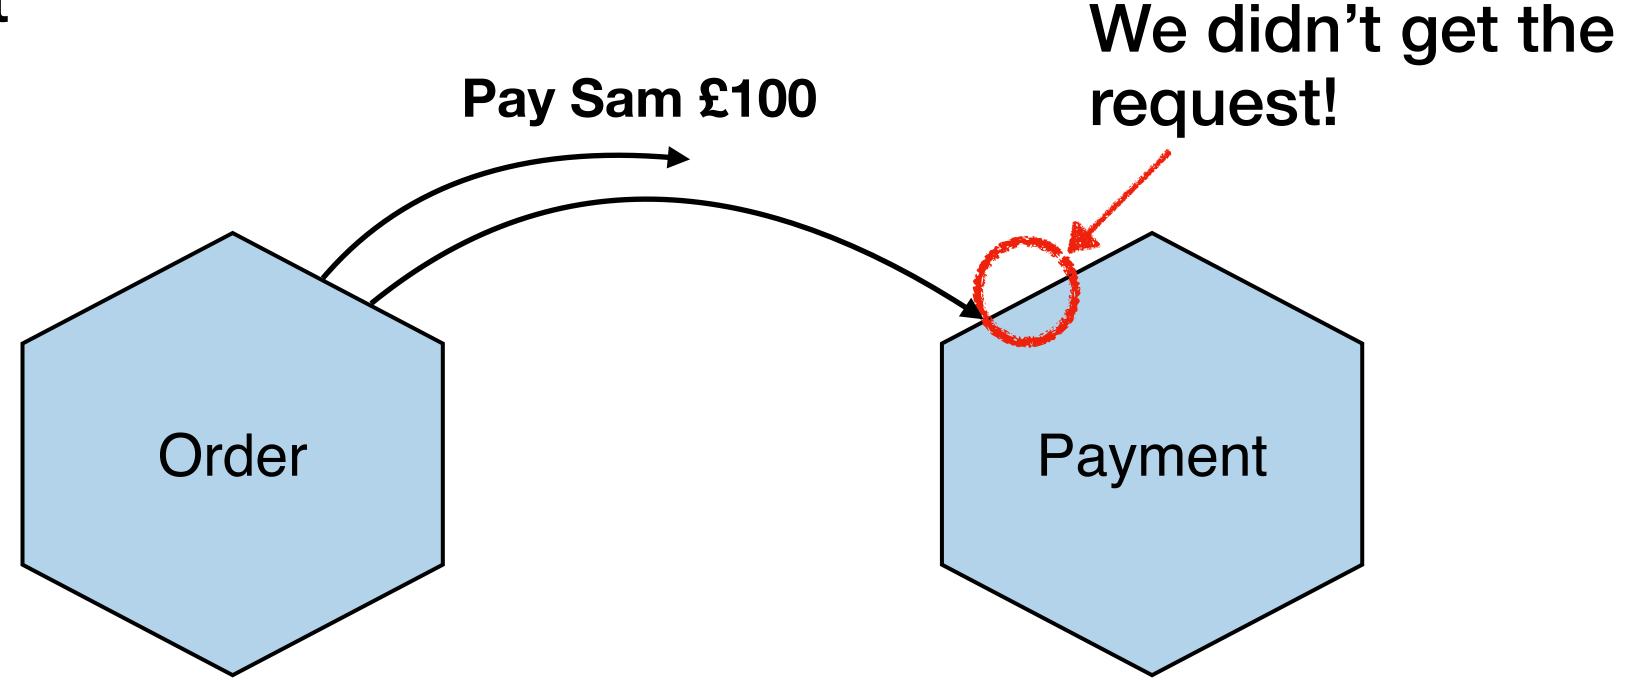
What happens if we don't get a response?



Two possibilities...

First Possibility

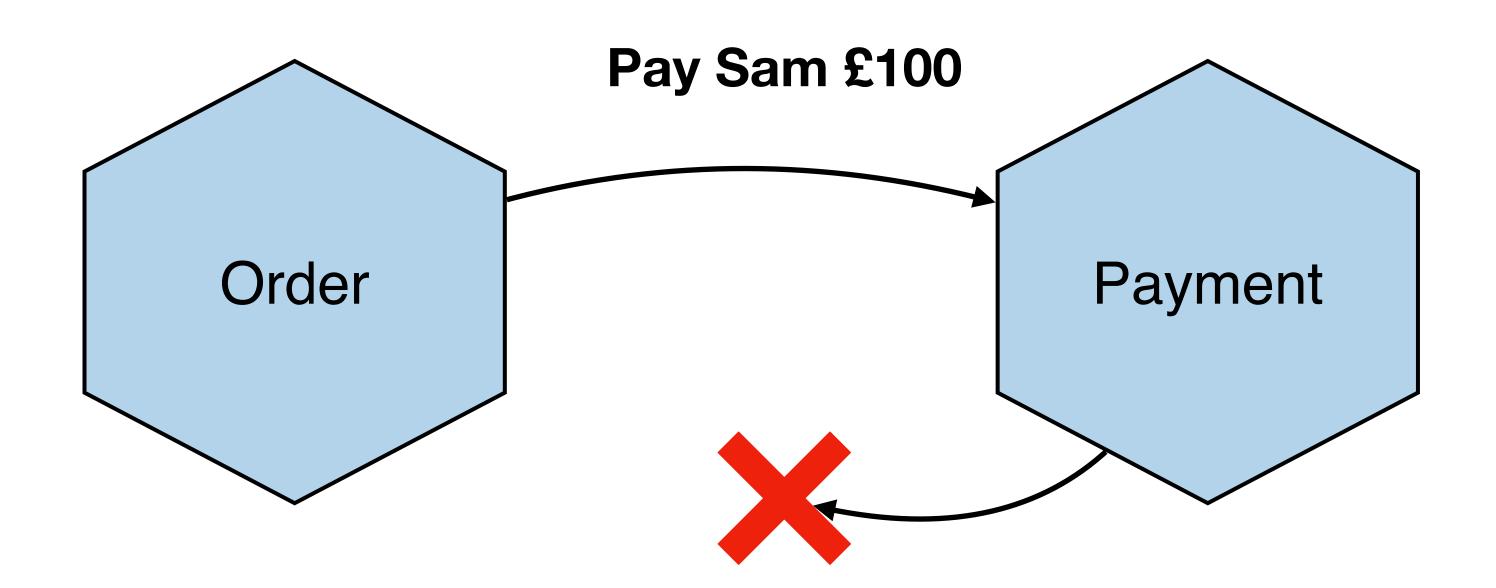
The payment wasn't carried out



Or we did, but crashed during processing...

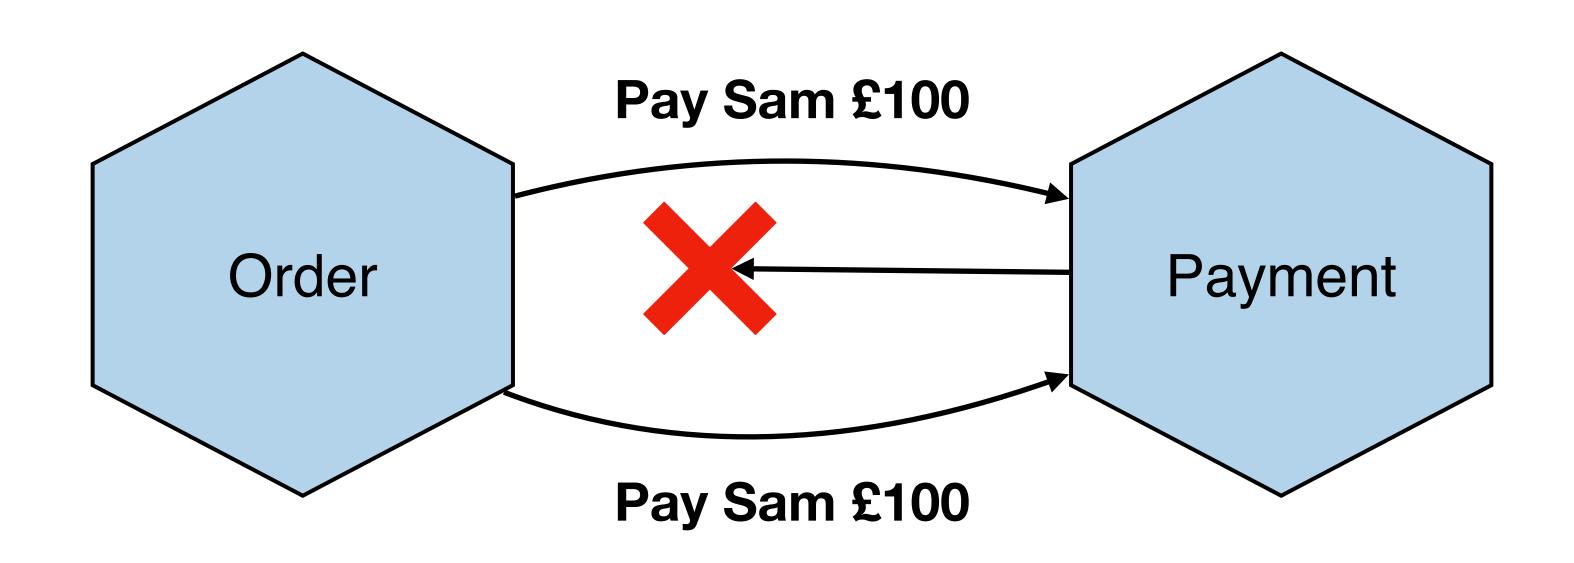
Second Possibility

The payment was performed, but the response wasn't processed





The new problem? What if that means Sam gets paid twice?

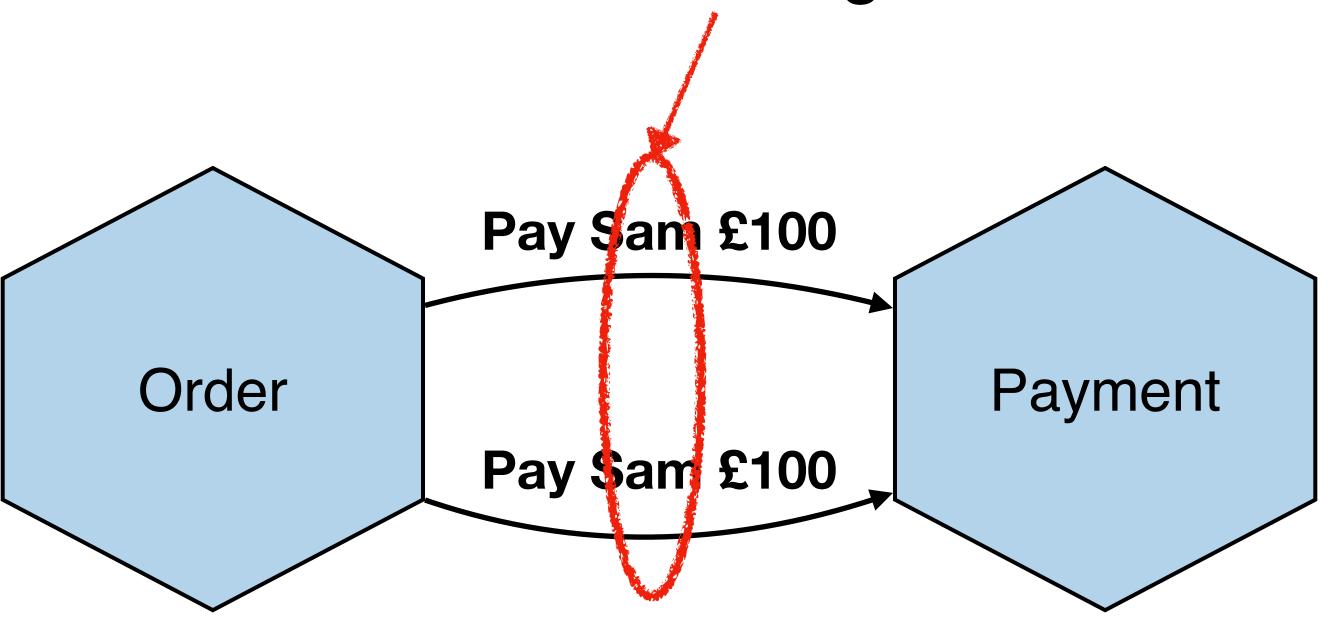




Idempotent Operation

An operation that can be applied multiple times without changing the result

This operation is NOT idempotent - performing the same operation twice does not give the same result



Idempotency Keys

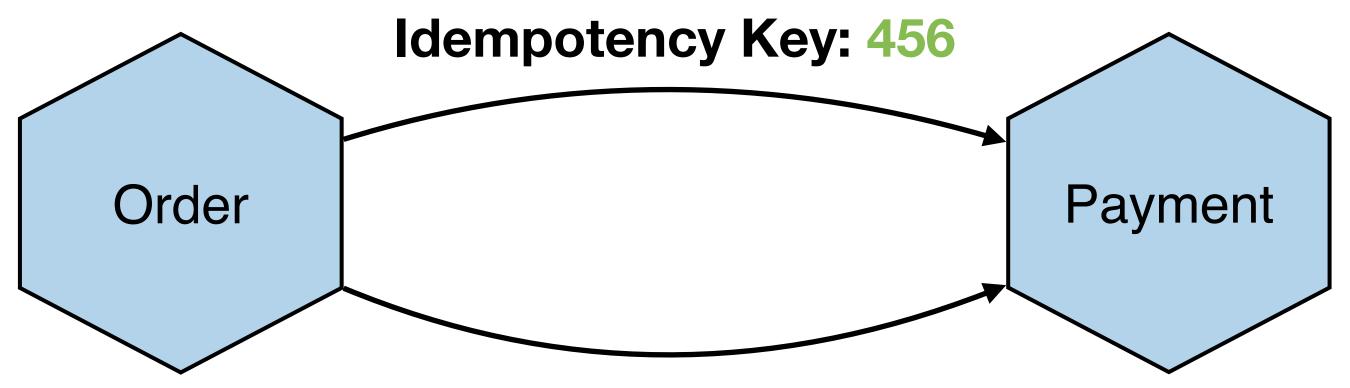
Server Side Request Fingerprinting

Idempotency Key

A unique ID used to identify a specific operation

Payment can now see this payment has already been made





Processed Payments

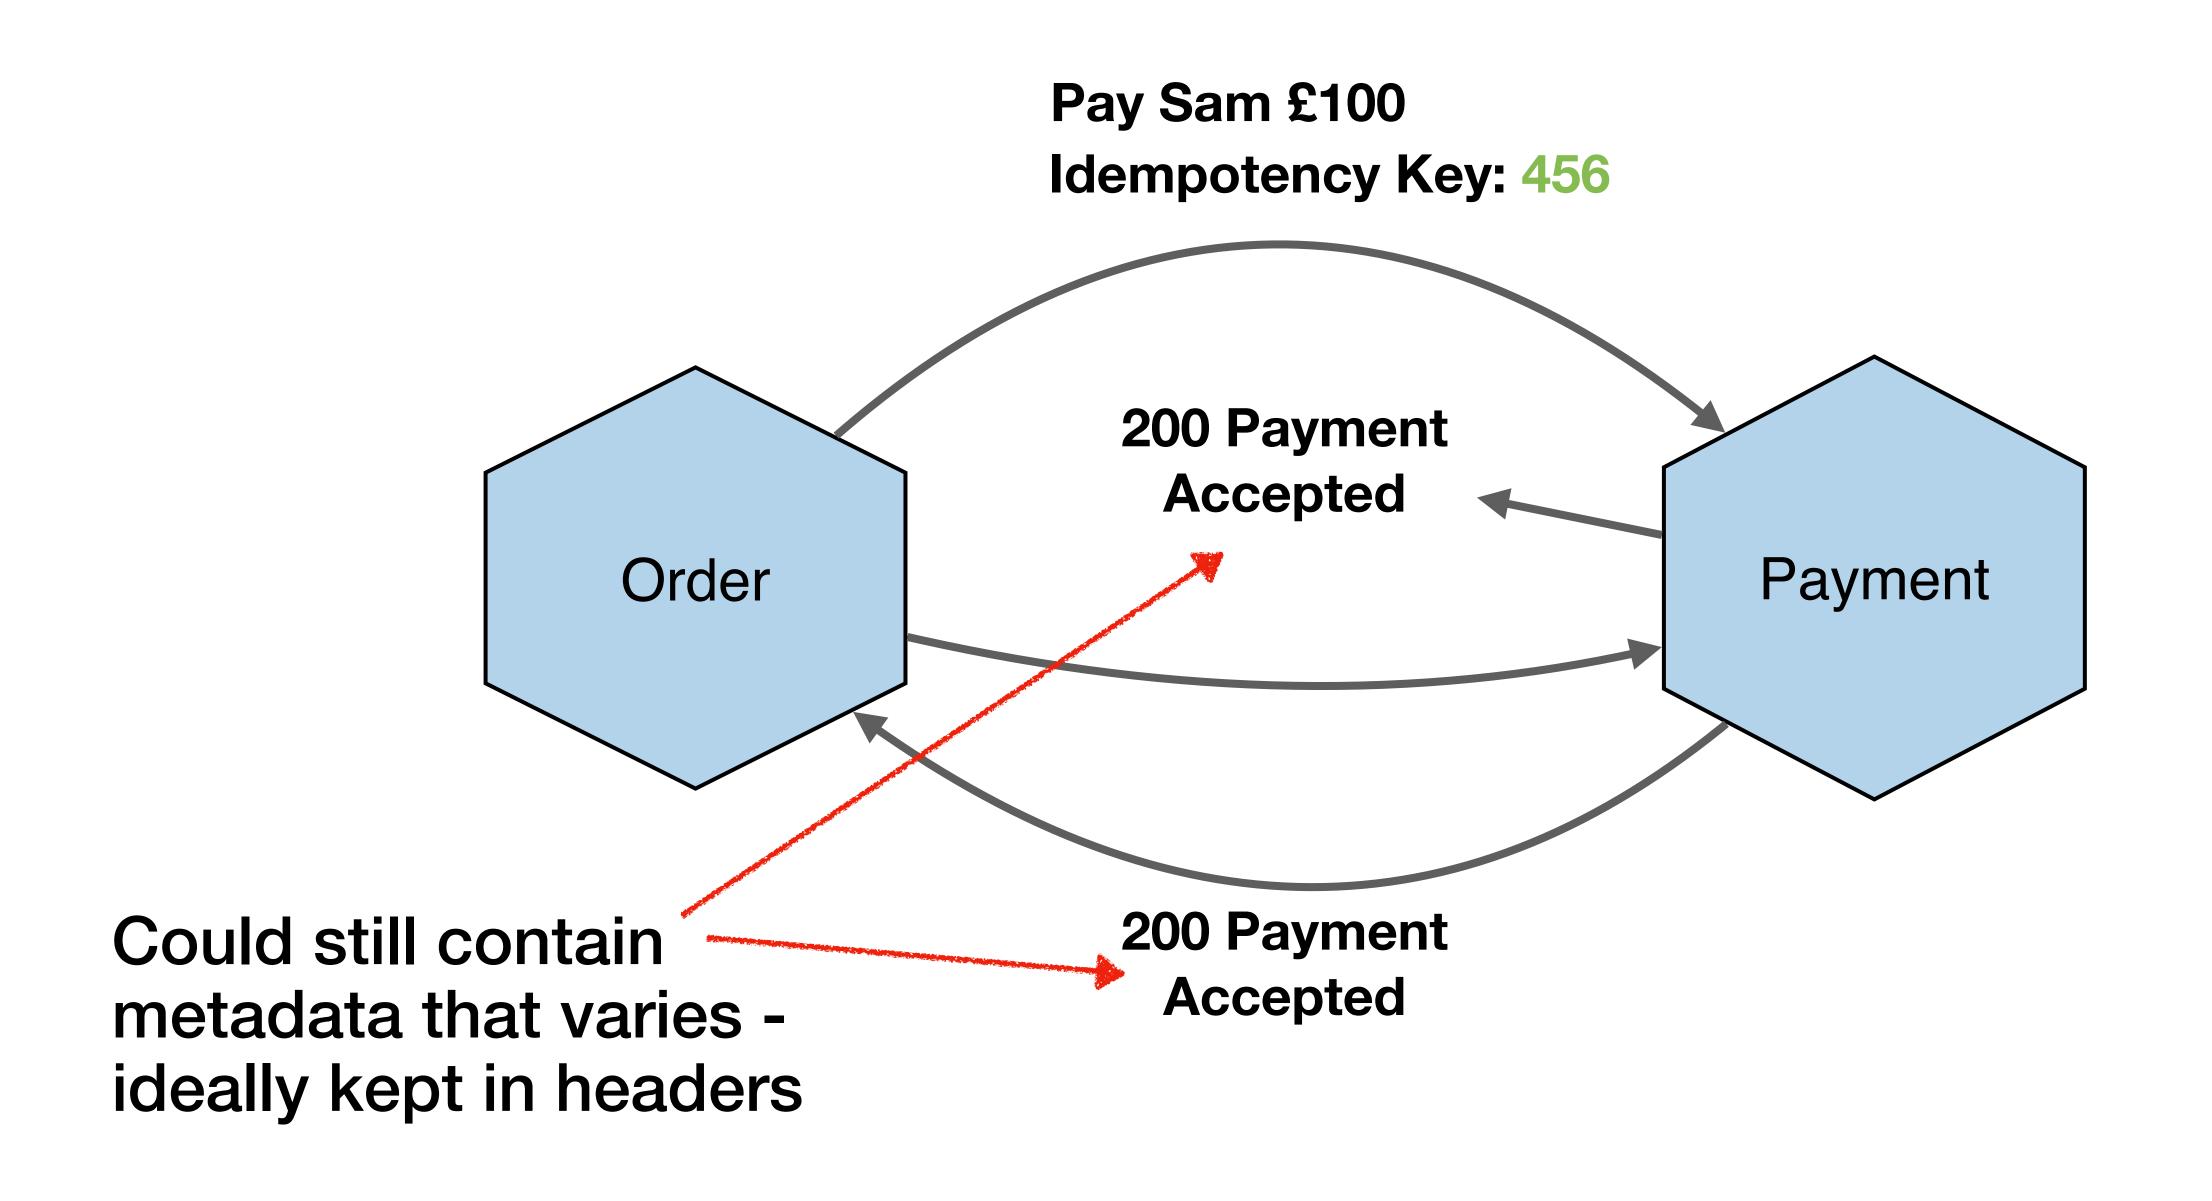
| Key | Result | |
|-----|-------------|--|
| 456 | 200 Success | |

Pay Sam £100

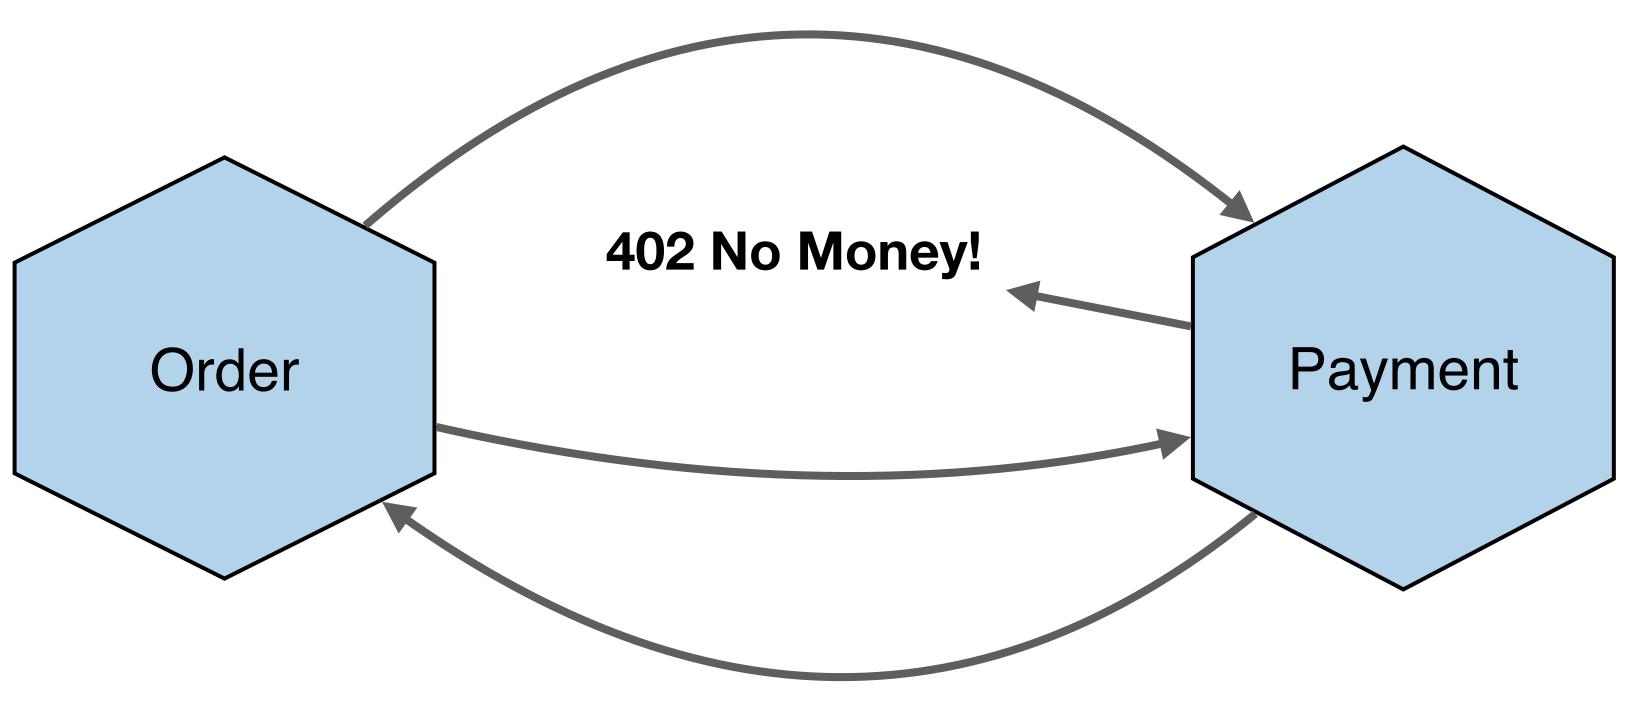
Idempotency Key: 456

What response should be sent when ignoring the retried operation?

Send the same (semantic) response as the original request would have been sent



Pay Sam £100 Idempotency Key: 456

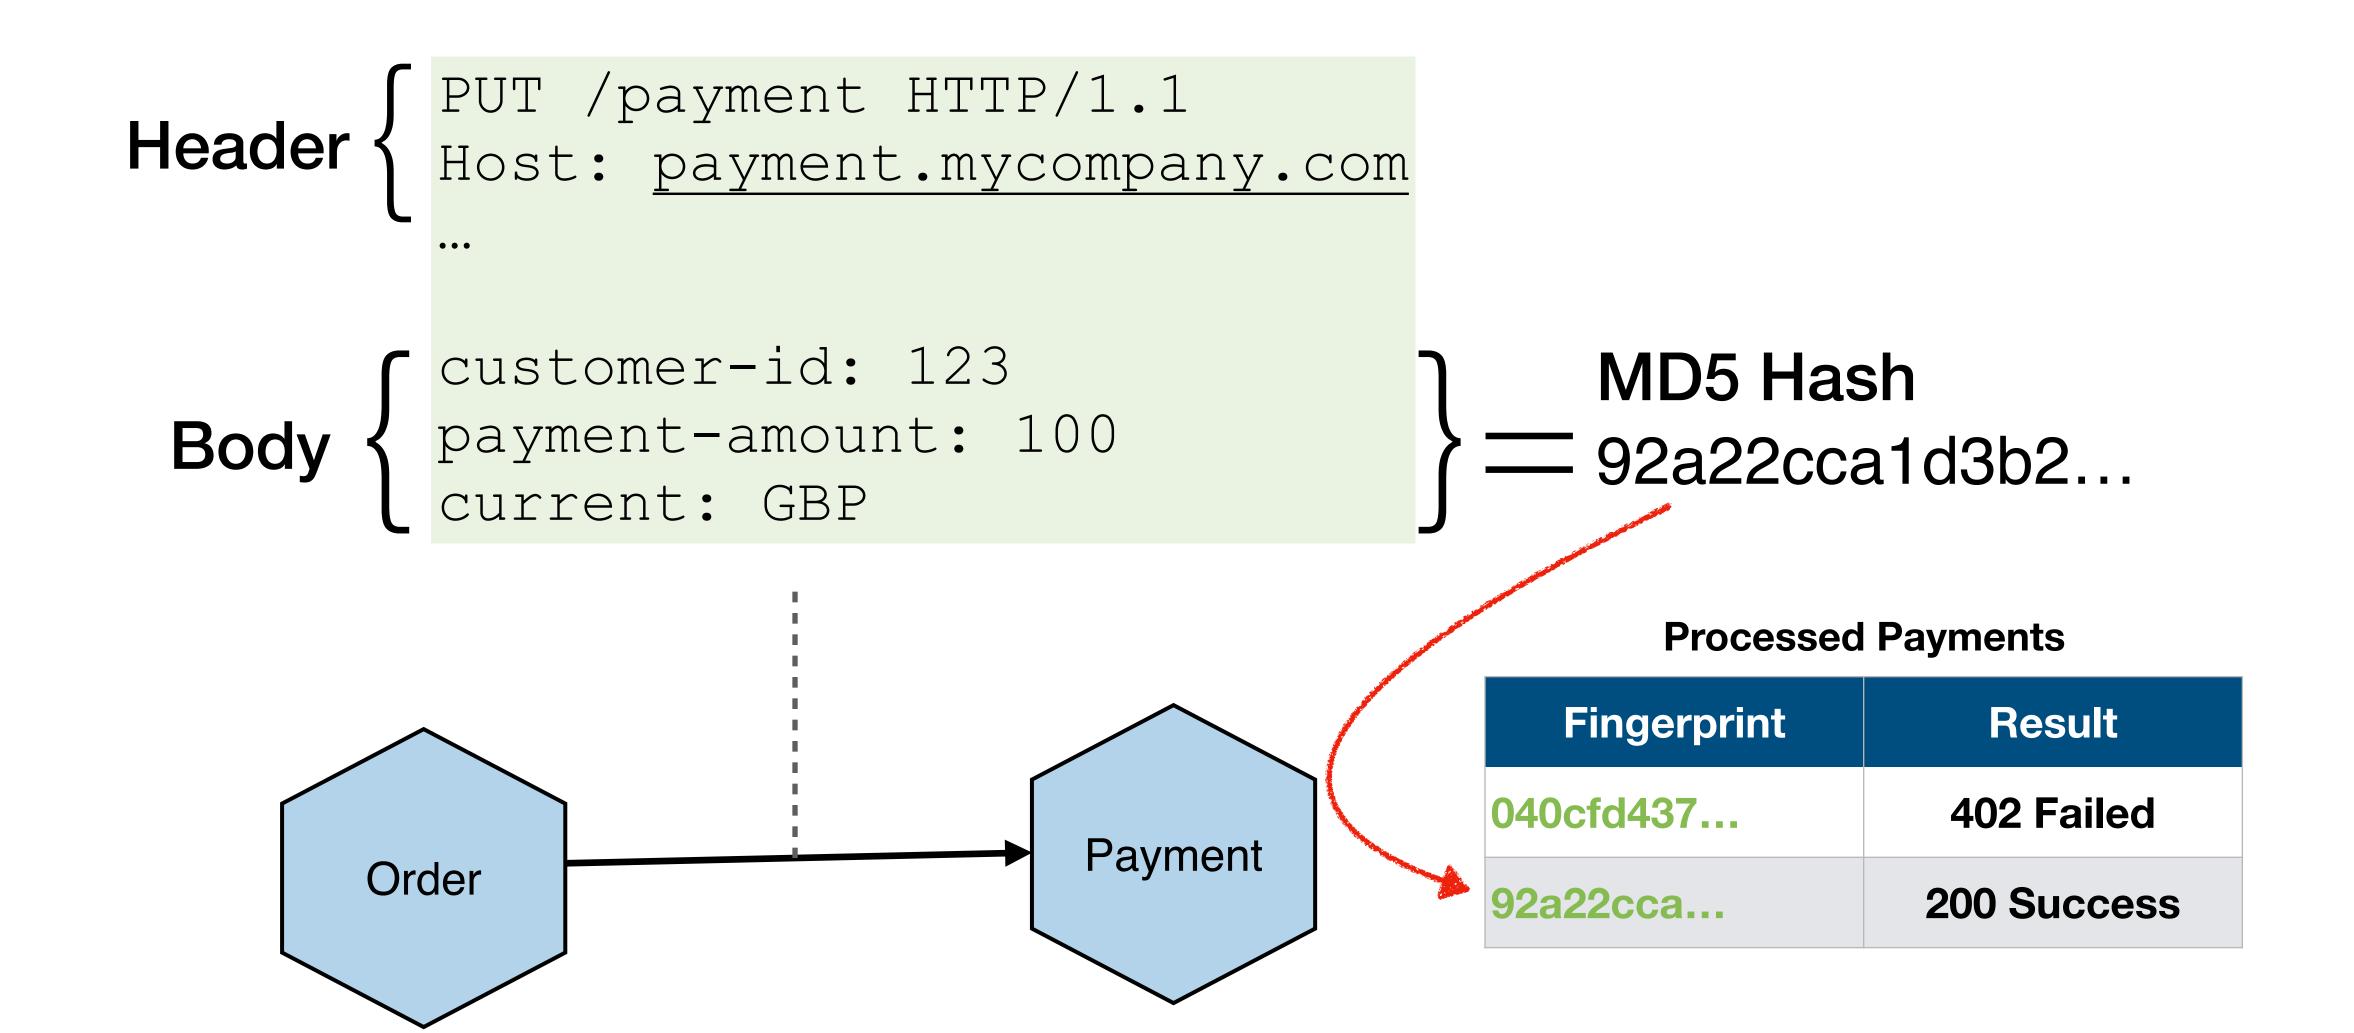


402 No Money!

Idempotency Keys IDs require a change in the client server protocol. What if that can't easily be changed?

Server Side Request Fingerprinting

Generating a fingerprint of the operation on the server side to spot duplicate requests



PUT /payment HTTP/1.1

Host:

payment.mycompany.com

customer-id: 123

payment-amount: 100

current: GBP

Timestamp: 1728468810

MD5 Hash 040cfd437dea56d... PUT /payment HTTP/1.1 Host:

payment.mycompany.com

customer-id: 123

payment-amount: 100

current: GBP

timestamp: 1728468853

MD5 Hash 997d0bb02e23f2a...

Option 1

Only use part of the request body for fingerprinting

```
PUT /payment HTTP/1.1
Host: payment.mycompany.com

customer-id: 123
payment-amount: 100
current: GBP

timestamp: 1728468810
```

Option 2

Move fields to the header

PUT /payment HTTP/1.1

Host: payment.mycompany.com

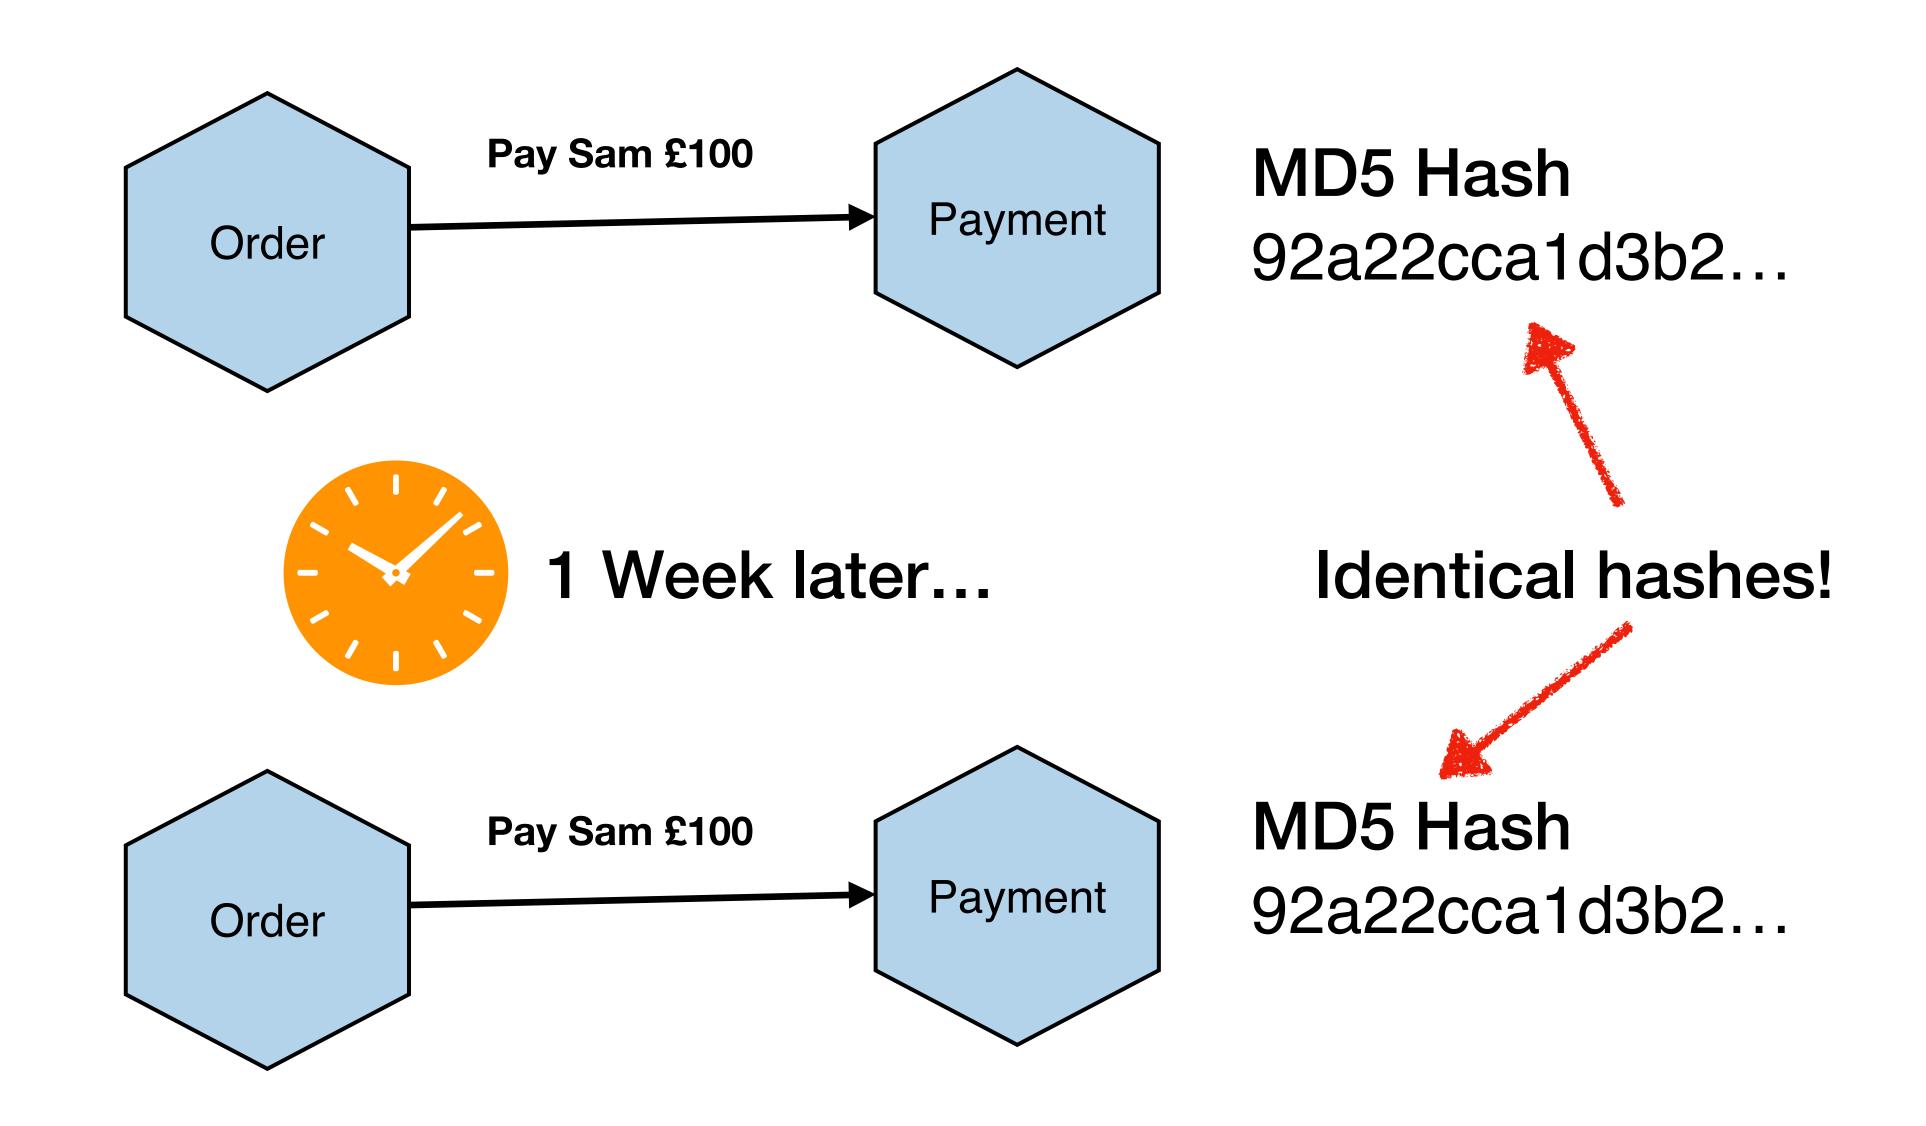
timestamp: 1728468810

customer-id: 123

payment-amount: 100

current: GBP

timestamp. 1728468810



| Fingerprint | Result | Expiry |
|-------------|---------|----------|
| 92a22cca1 | Success | 12:01:03 |
| 997d0bb0 | Failed | 12:01:17 |
| 040cfd437 | Success | 12:03:45 |

Getting retention right when using fingerprints is very important!



Can still be a problem if you want to do the same thing multiple times in a short space of time



If you can change the client-server protocol, then idempotency keys are the easiest way to make an operation idempotent

Otherwise consider server side request fingerprinting, but be aware of the downsides

Using idempotency keys and server side request fingerprinting together?

Can be useful to catch client errors

Pay Sam £100 accepted" here could **Idempotency Key: 456** confuse the client! 200 Payment **Accepted** Order Payment Pay Sam £200 **Idempotency Key: 456** If the server sees a retry with changing values, sending an **409 Conflict** error can be more clear

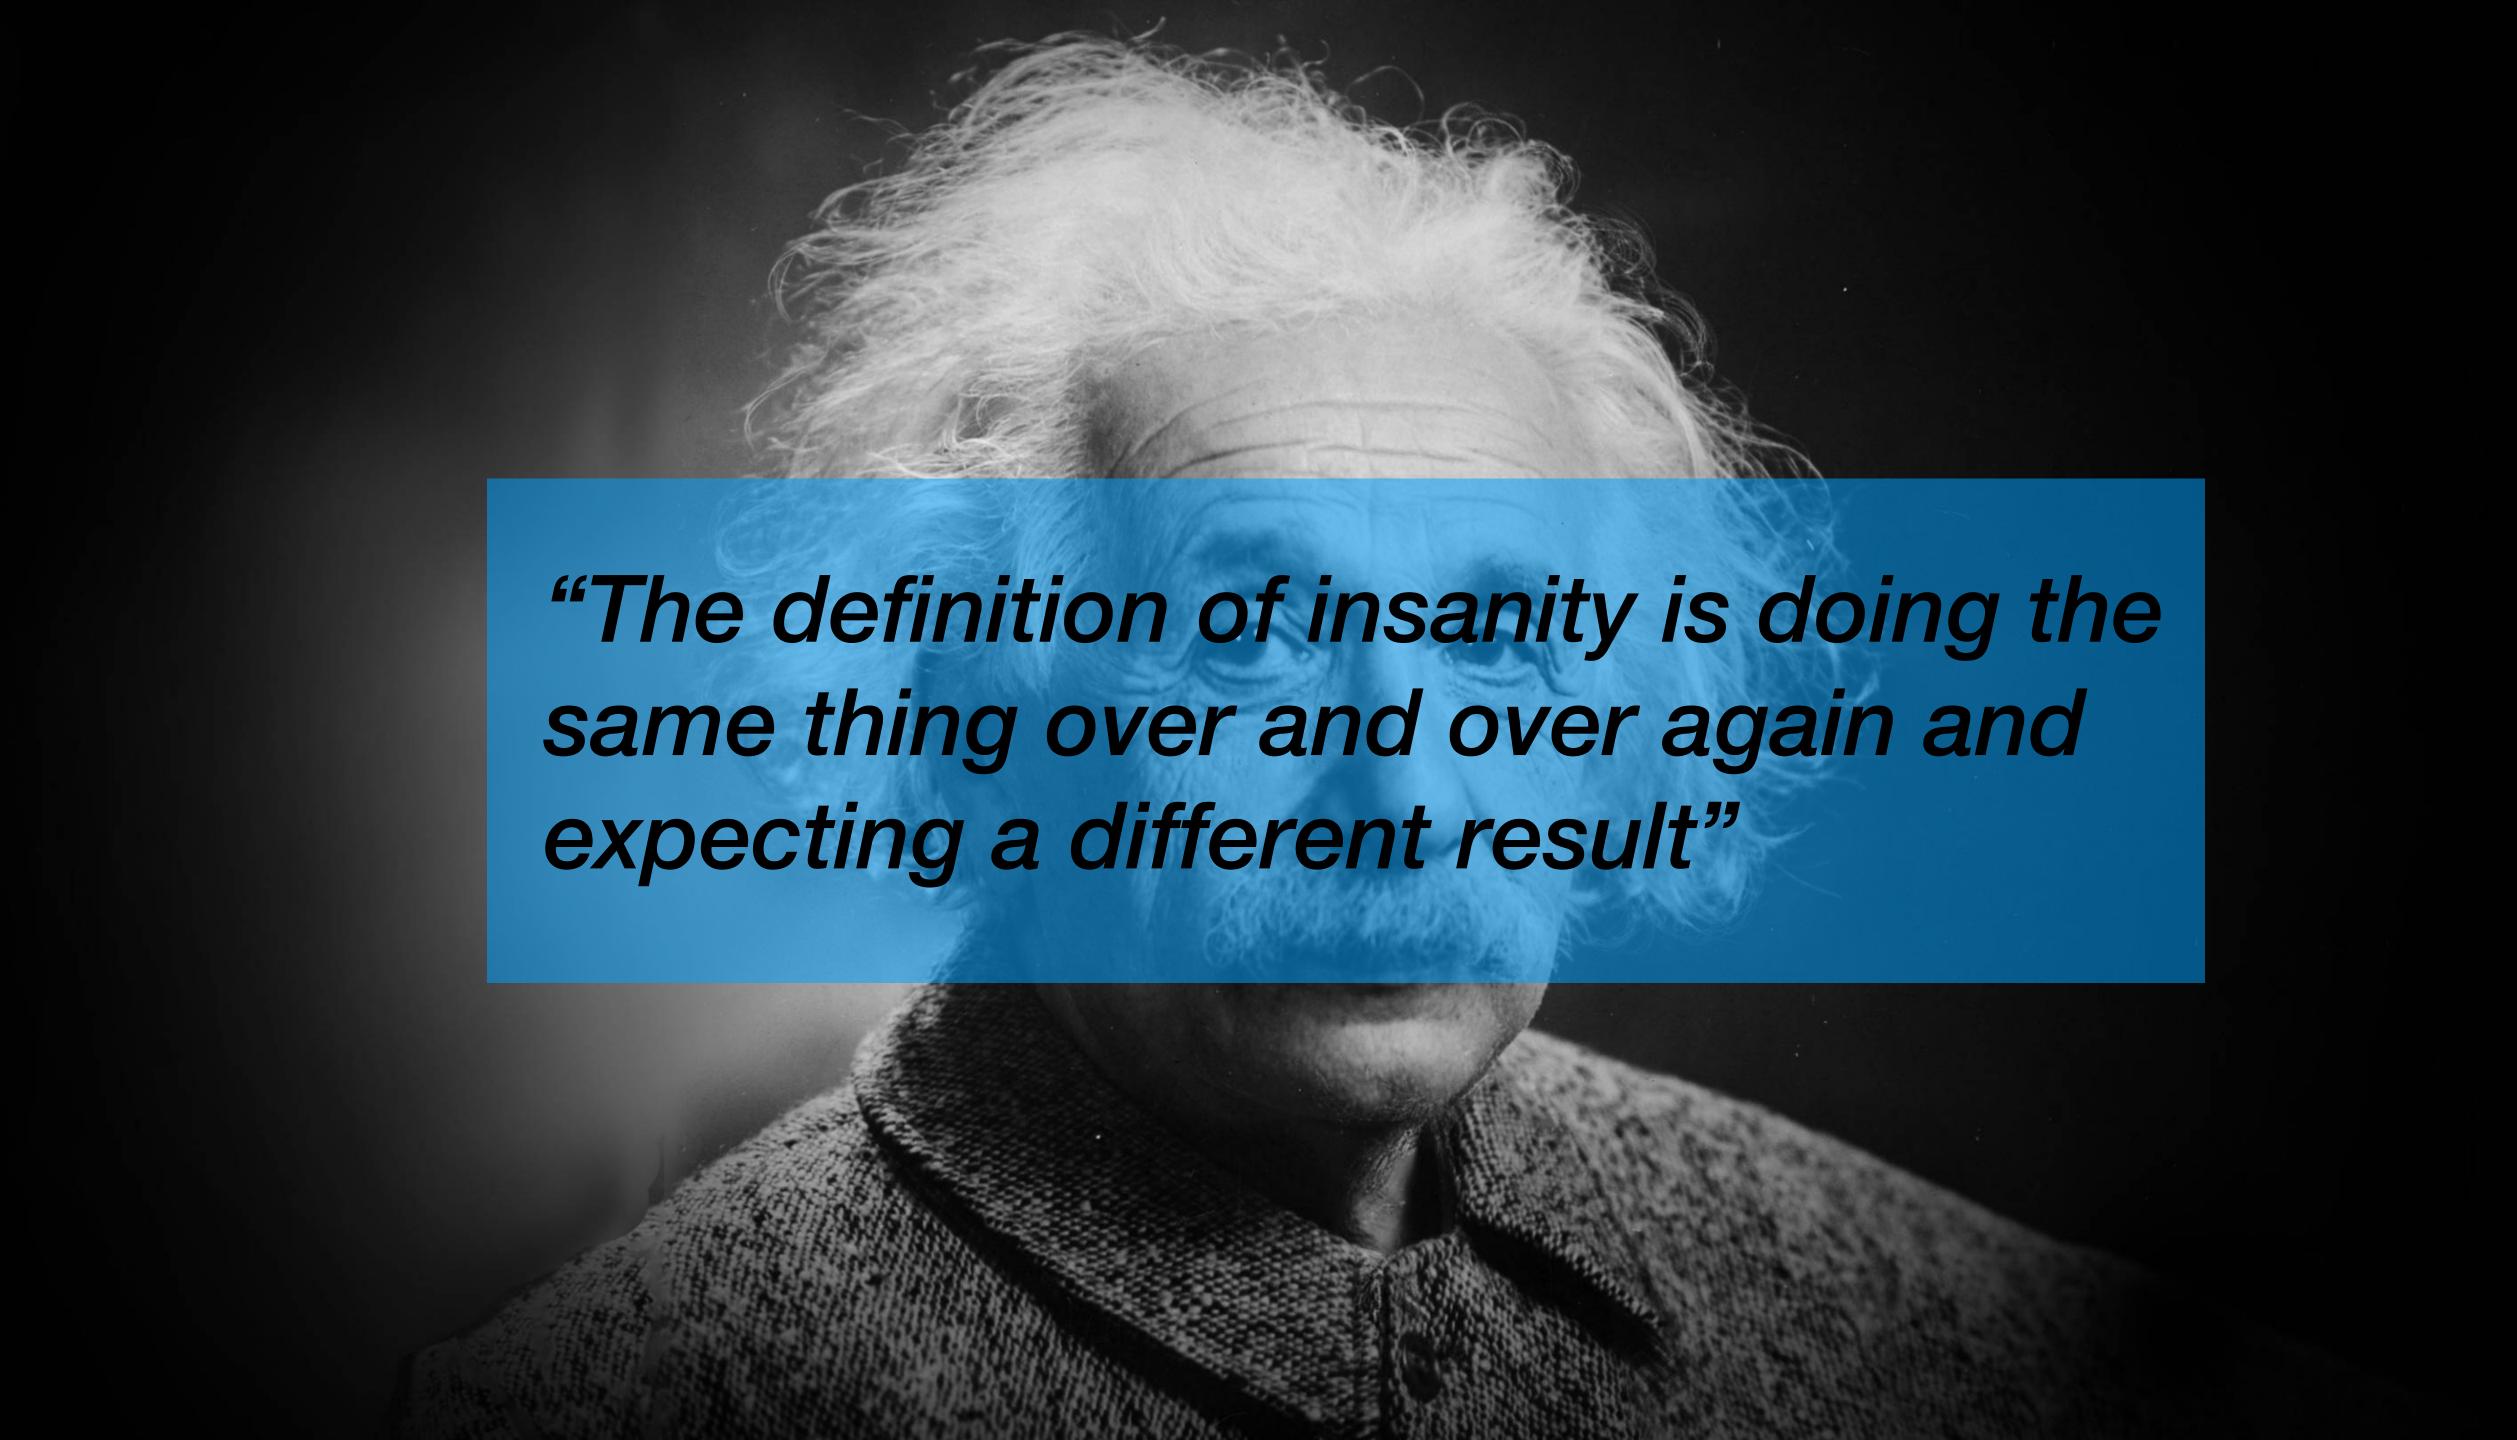
Resending "Payment

| Key | Fingerprint | Result | Expiry |
|-----|-------------|---------|----------|
| 456 | 92a22cca1 | Success | 12:01:03 |
| 789 | 997d0bb0 | Failed | 12:01:17 |
| 123 | 040cfd437 | Success | 12:03:45 |

Just store the key and the fingerprint together

So doing the same thing over and over again might be sensible...

...but only if you can do so safely



Einstein didn't say this quote

THE THREE GOLDEN RULES OF DISTRIBUTED COMPUTING

1. You can't beam information between two points instantly

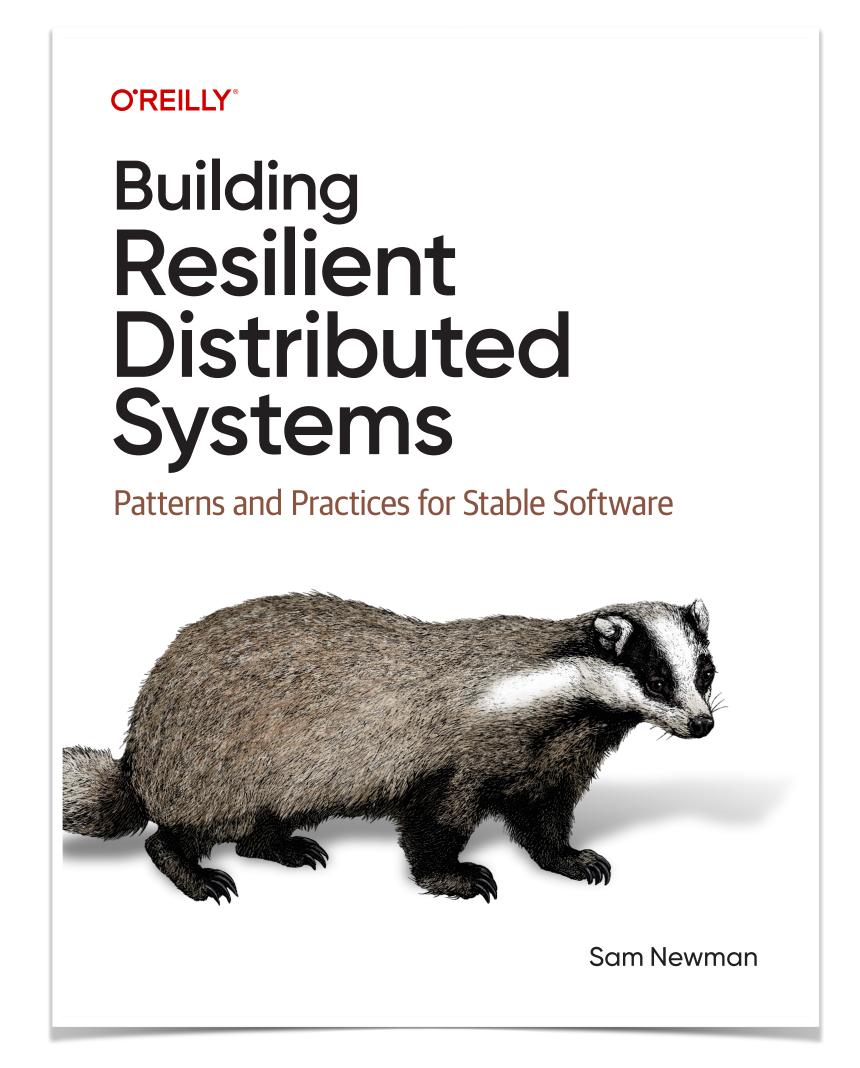
2. Sometimes, you can't reach the thing you want to talk to

3. Resources are not infinite

If at first you don't succeed, retrying can make sense...to a point!

If you do want to retry, making your operations idempotent is vital

NEW BOOK IN EARLY ACCESS!





Book details and slides

https://samnewman.io/books/building-resilient-distributed-systems/