

Lost in transaction?

Über (In-)Konsistenz in
verteilten Systemen





Benjamin Hoffmann

Technical Consultant
Camunda



Berlin, Germany




benjamin.hoffmann@camunda.com

REST, SOAP,
Cloud, SaaS,
Microservices, SCS,
FaaS, Serverless,

...

Distributed systems






Failure will happen.
Accept it!

But keep it local!
Be resilient.

Let's start with a simple example



Live demo



Circuit Breaker



Photo by CITYEDV, available under Creative Commons CC0 1.0 license.

Live demo



Fail fast
is important



Photo by Tookapic, available under Creative Commons CC0 1.0 license.

Ihre Bordkarten

Beim Versenden der Bordkarte ist ein Fehler aufgetreten.

Ihr Buchungscode **08HHSS**


Hinflug

BERND RUECKER

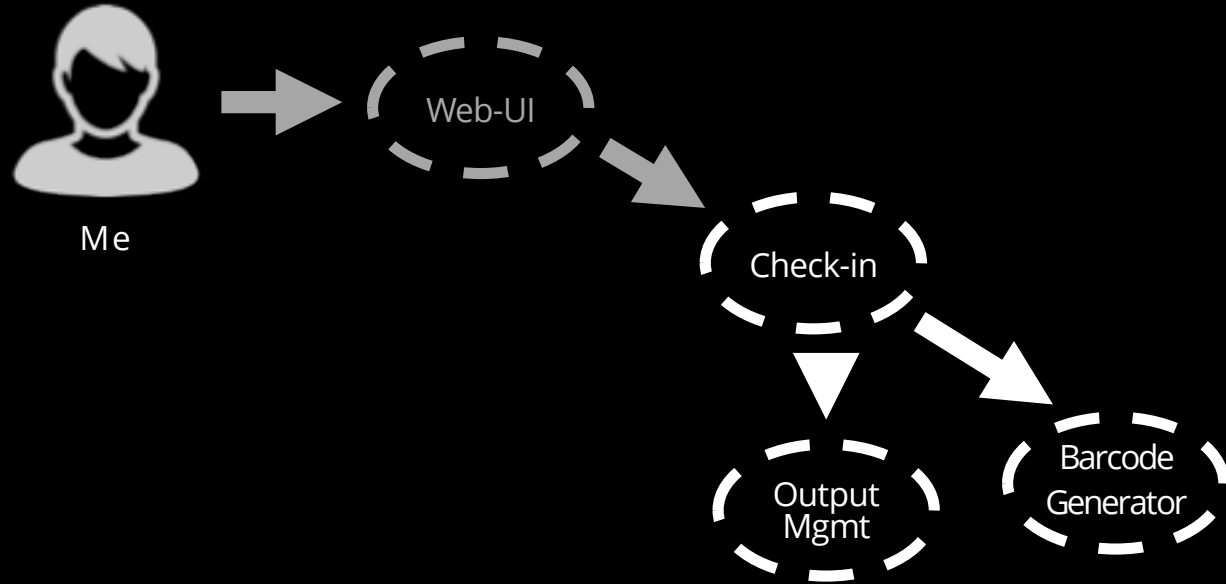
Stuttgart (STR) - London-Stansted (STN)

„There was an error while sending your boarding pass“

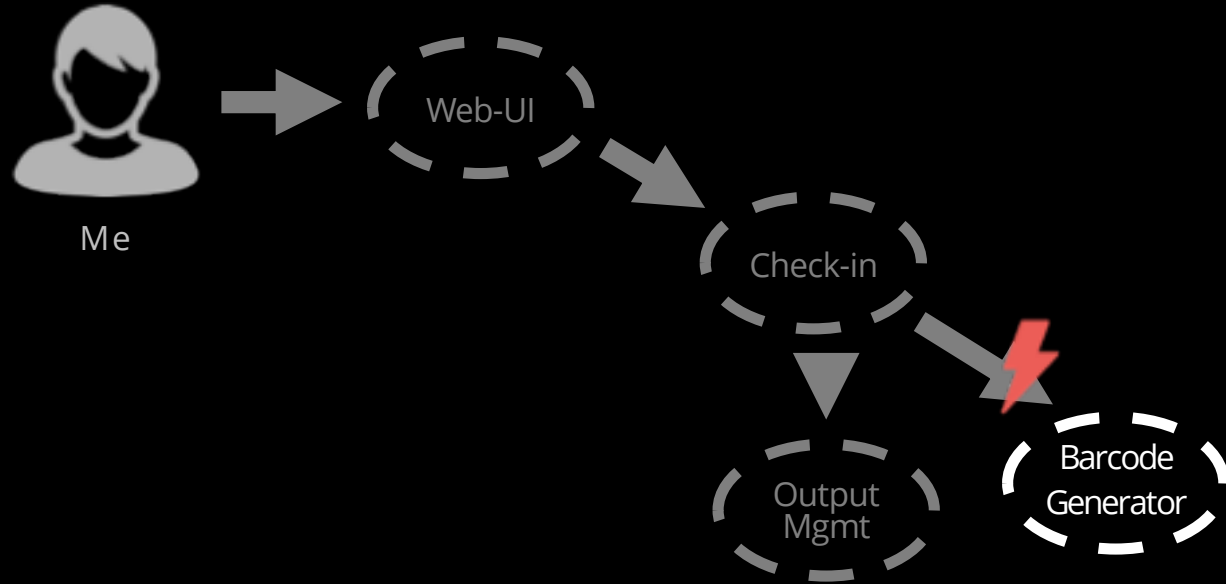
Current situation



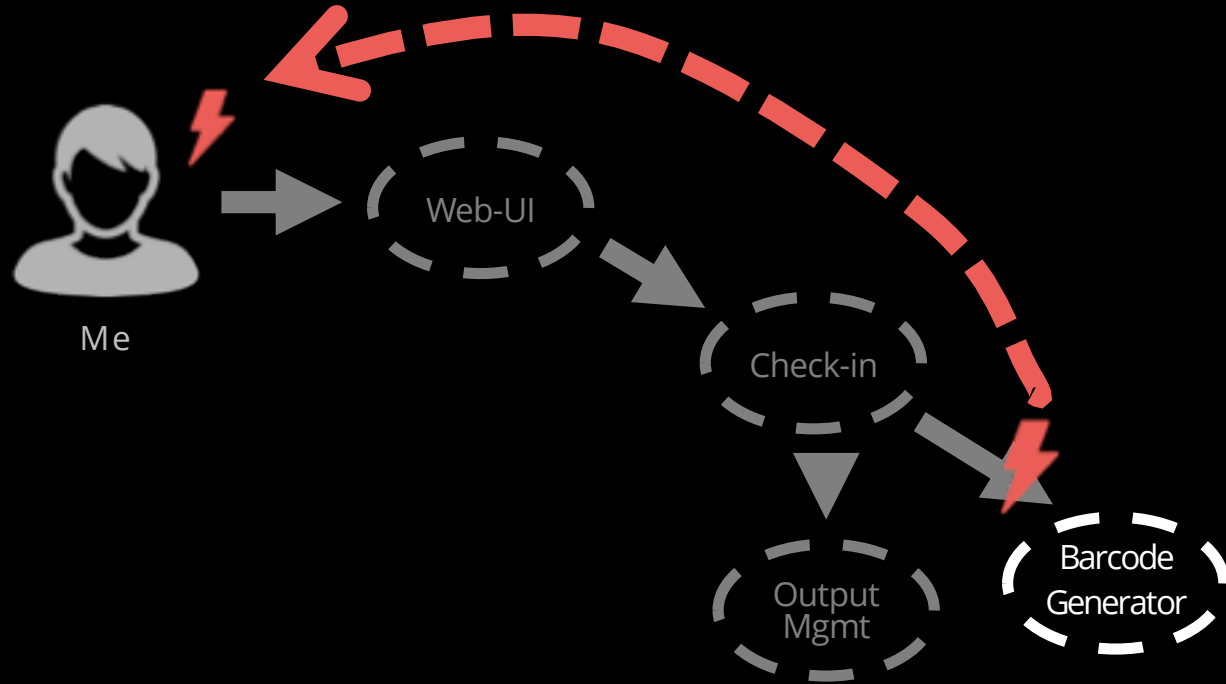
Current situation



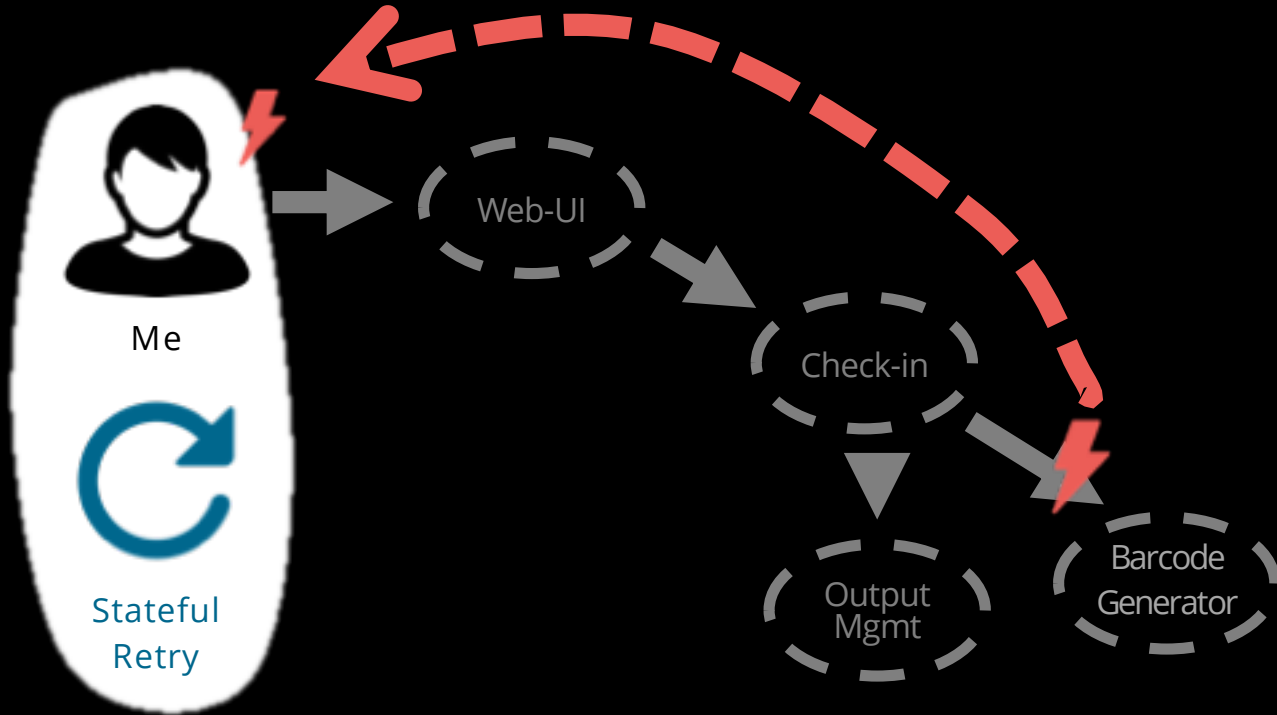
Current situation – the bad part



Current situation – the bad part



Current situation – the bad part



Ihre B

Ihr Buchu

Hinflug

BERND RUEC

easyJet

We're sorry

We are having some technical difficulties at the moment.

Please log on again via www.easyjet.com

If that doesn't work, please try again in five minutes.

We do actively monitor our site and will be working to resolve the issue, so there's no need to call

[Go to easyJet.com](http://www.easyjet.com)

Fail fast
is important

Fail fast
is important
but not enough!

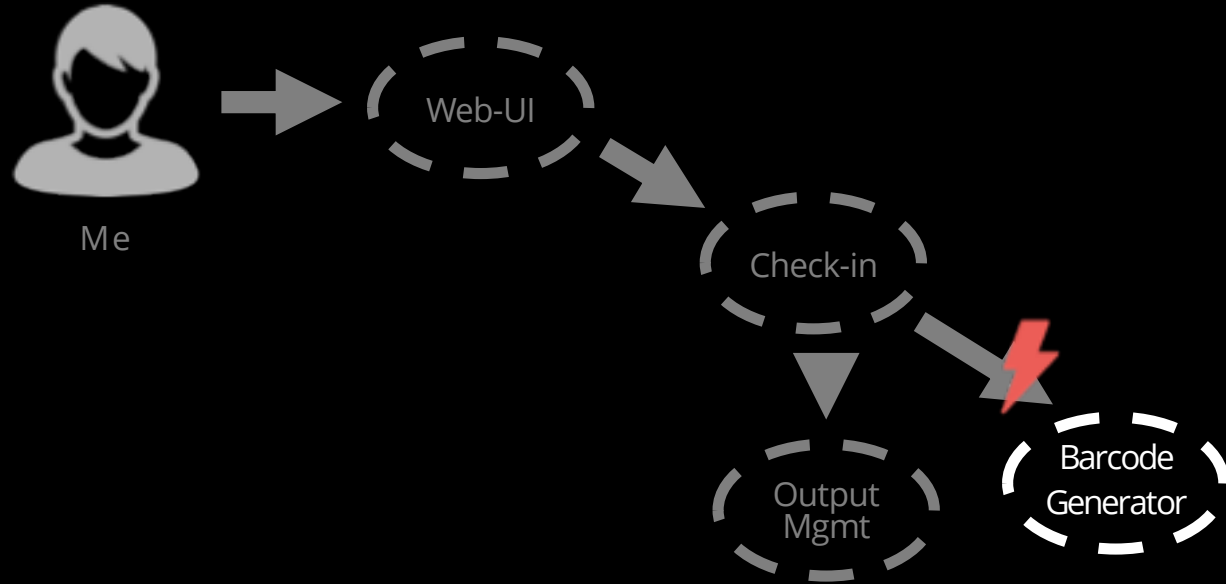
...I just made this up...

We're sorry

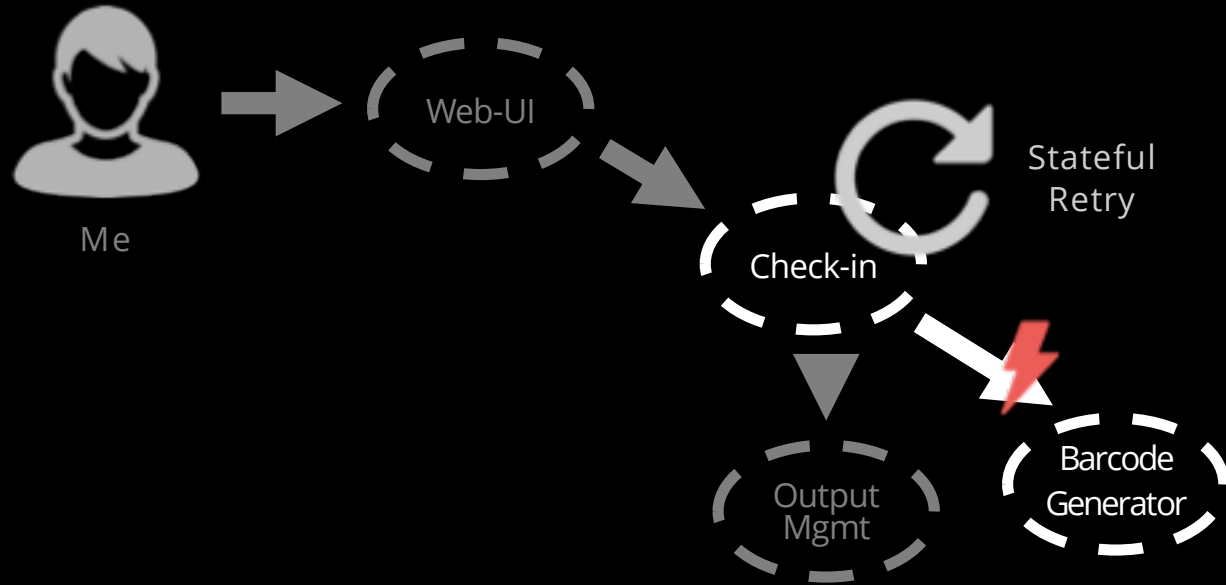
We are having some technical difficulties and cannot present you your boarding pass right away.

But we do actively retry ourselves, so lean back, relax and we will send it on time.

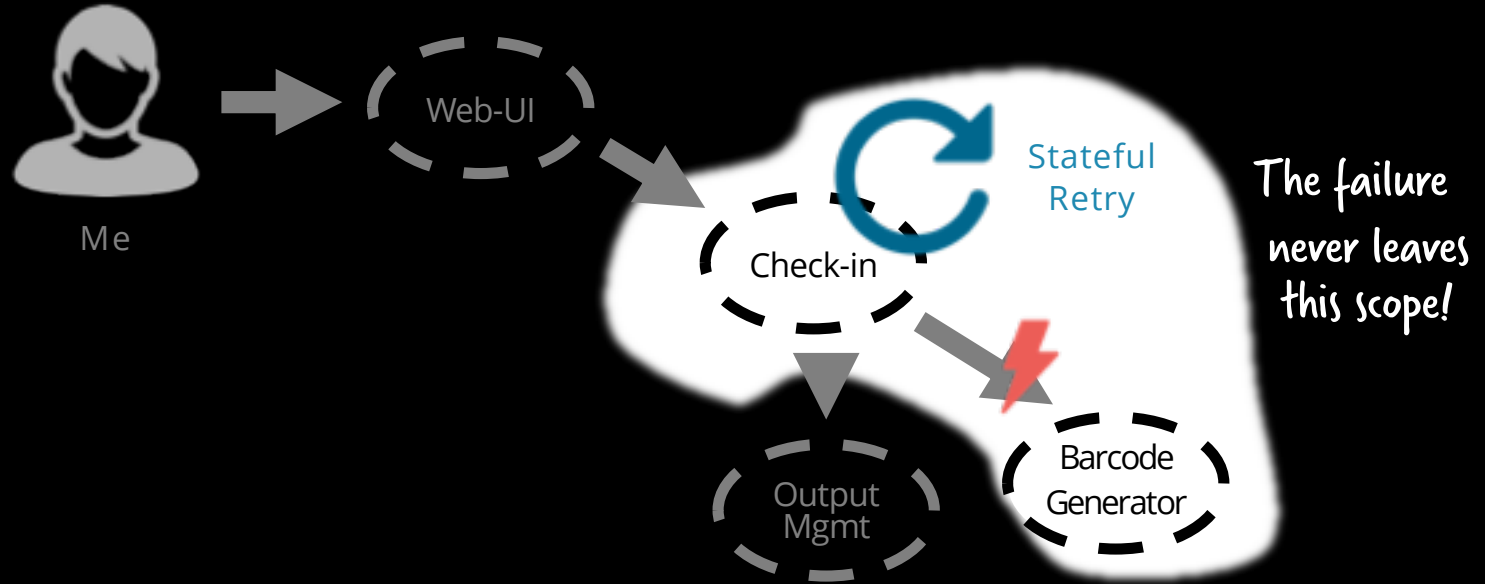
Possible situation – much better!



Possible situation – much better!



Possible situation – much better!



Handling State




Persist thing
(Entity, Document, Actor, ...)

Typical
concerns

DIY = effort,
accidental
complexity



Scheduling, Versioning,
operating, visibility,
scalability, ...



State machine or
workflow engine

Complex, proprietary,
heavyweight, slow,
difficult to use,
developer adverse



Workflow engines,
state machines



AWS Step
Functions

It is
relevant
in modern
architectures



Workflow engines,
state machines



Silicon valley
has recognized



NETFLIX OSS



Workflow engines,
state machines

UBER

CADENCE

There are
lightweight open source
options



camunda



Activiti



AWS Step
Functions



conductor

NETFLIX OSS



Workflow engines,
state machines

UBER

CADENCE

also at scale



zeebe

by Camunda



AWS Step
Functions



camunda



jBPM



conductor

NETFLIX OSS



Activiti



Workflow engines,
state machines




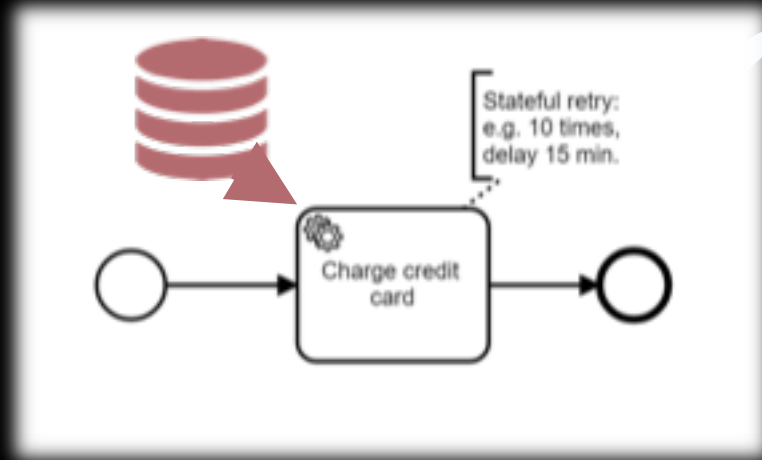
for today's demo



Live demo



Now you have a state machine!

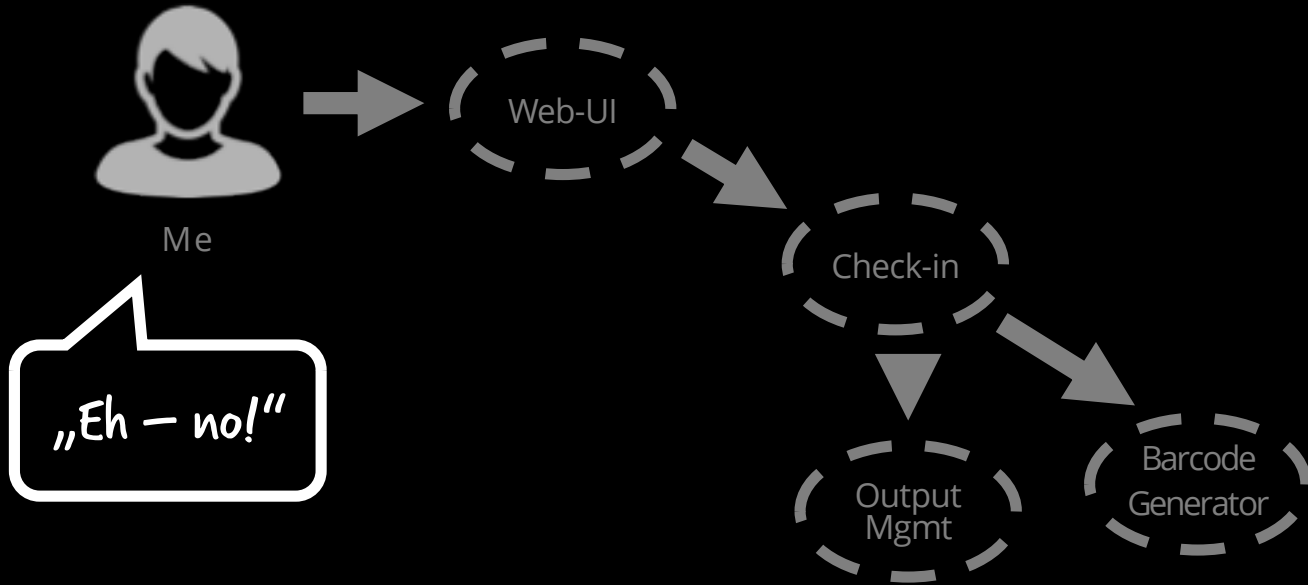


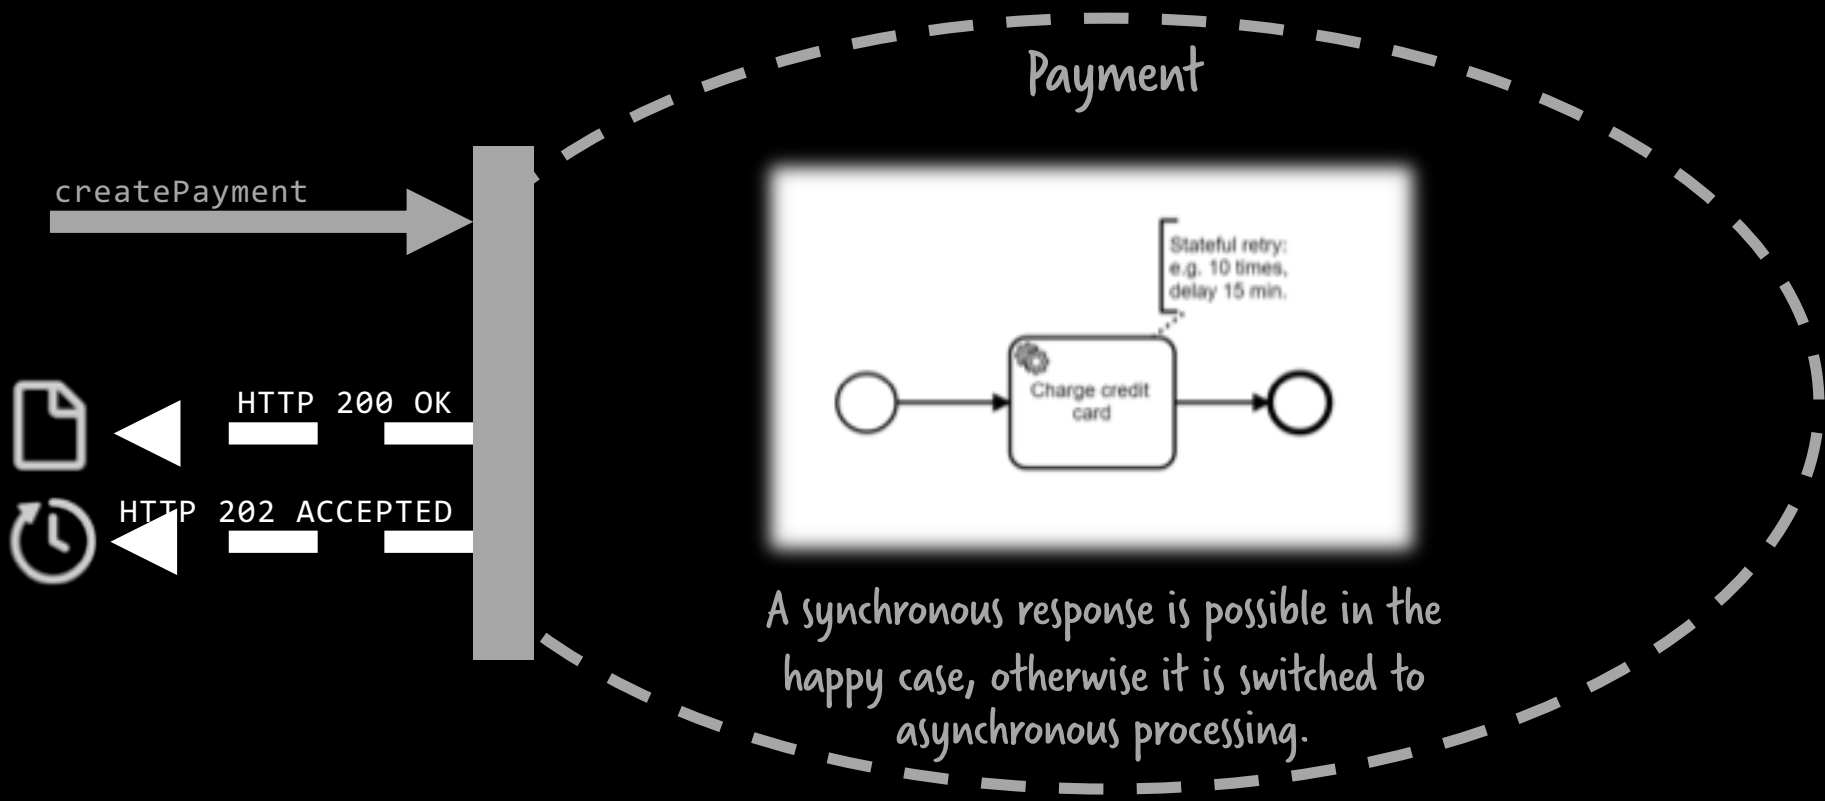
Client

has to implement

Retry

„The customer wants a synchronous response“





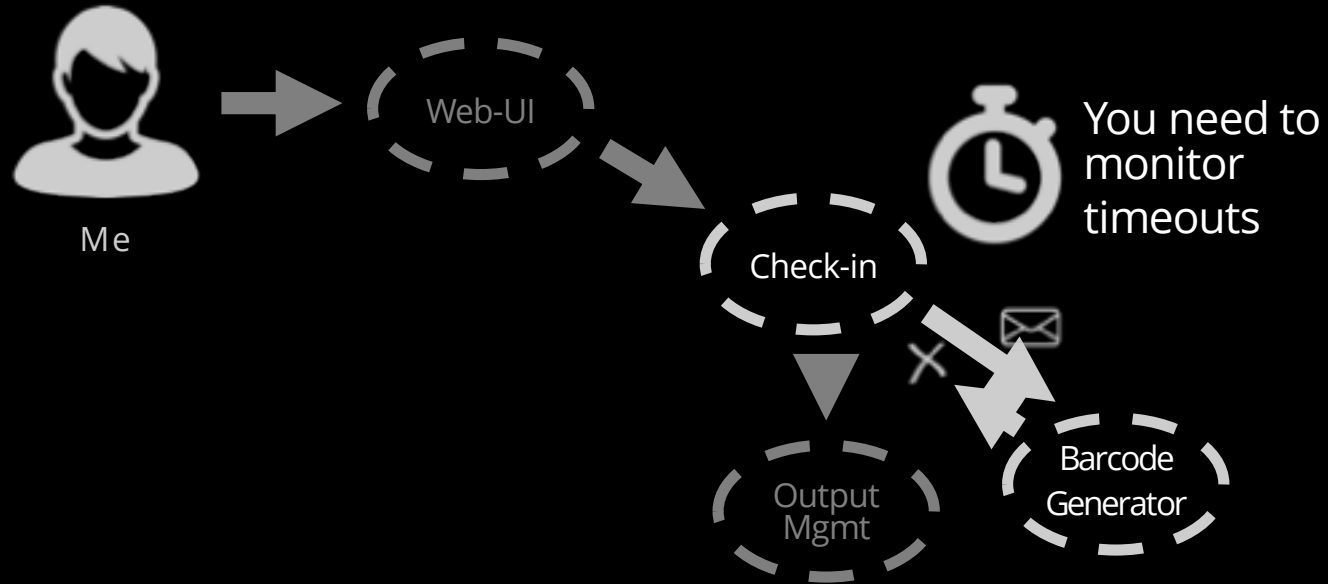
Live demo



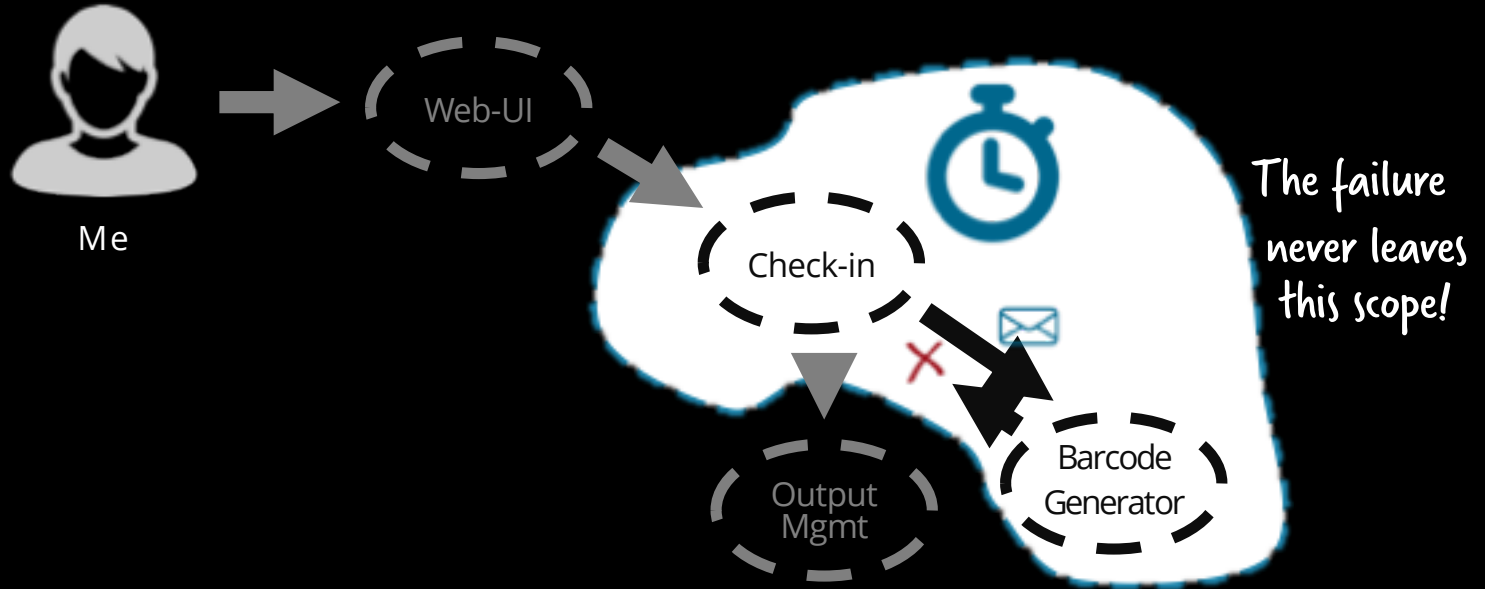
Synchronous communication
is the crystal meth of
distributed programming

Todd Montgomery and Martin Thompson
in "How did we end up here" at GoTo Chicago 2015

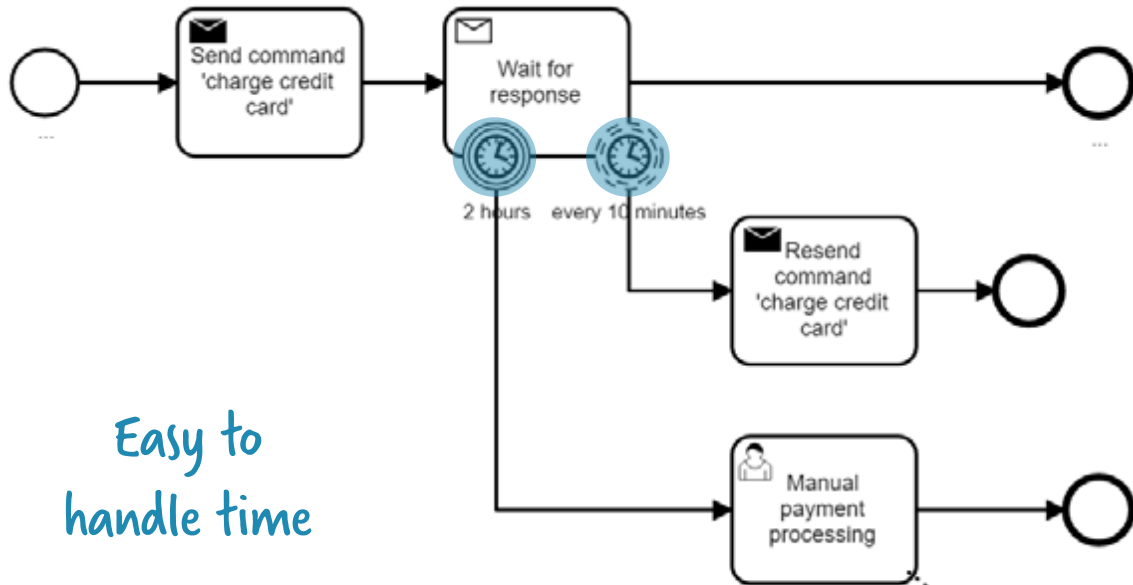
Asynchronous communication



Remember...



Workflow ...



Easy to handle time

Probably makes no sense in this example :-)

BPMN

Business Process
Model and Notation

ISO Standard



Client

has to implement


Retry,
Timeout

Distributed systems


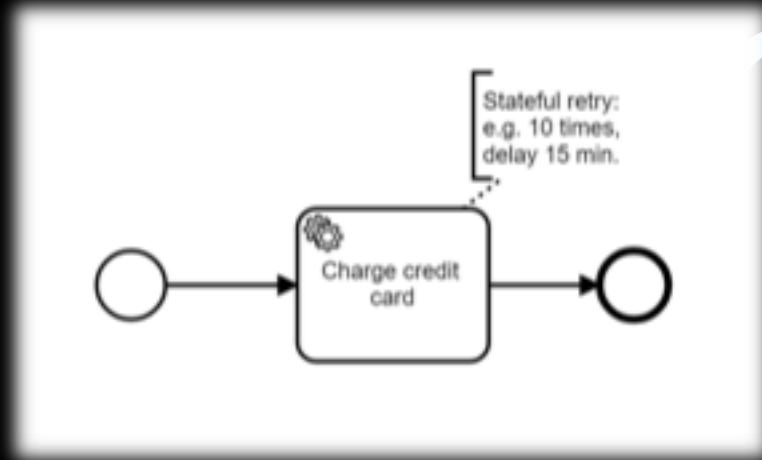


It is impossible to differentiate certain failure scenarios.

Independent of communication style!



What does it mean?



Duplicates
Duplicates



Client

has to implement

Retry,
Timeout

Service Provider

has to implement

Idempotency

Idempotency
is a business
problem!



We are processing your payment.
Do not leave this page.

And for god's sake – do not
reload!

Idempotency
is a business
problem!



We are currently processing your request.
Don't worry, it will happen safely –
even if you loose connection.
Feel free to reload this page at any time!

But we wanted to
talk about consistency!

Distributed systems

2007

Life beyond Distributed Transactions: an Apostate's Opinion

Position Paper

Pat Helland
Amazon.Com
705 Fifth Ave South
Seattle, WA 98104
USA

PHelland@Amazon.com

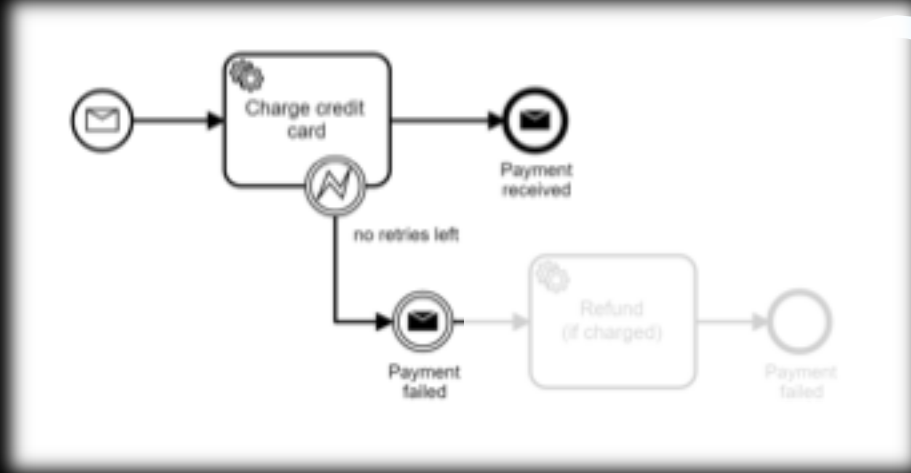
The positions expressed in this paper are personal opinions and do not in any way reflect the positions of my employer Amazon.com.

ABSTRACT


Many decades of work have been invested in the area of distributed transactions, including protocols such as 2PC. This paper explores approaches to queuing and the current practice of

Instead, applications are built using different techniques which do not provide the same transactional guarantees but still meet the needs of their businesses. This paper explores the current practice of

Distributed systems introduce complexity you have to tackle!

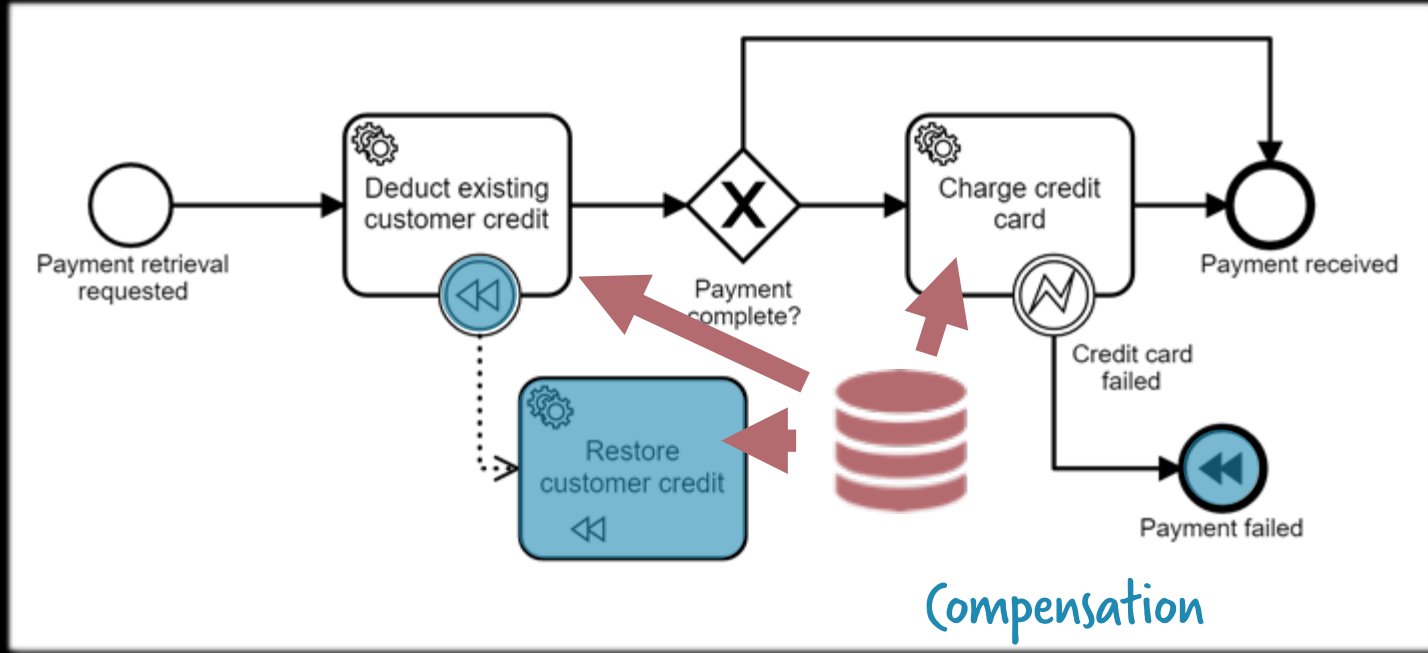


Distributed systems introduce complexity you have to tackle!



Distributed transactions using compensation *

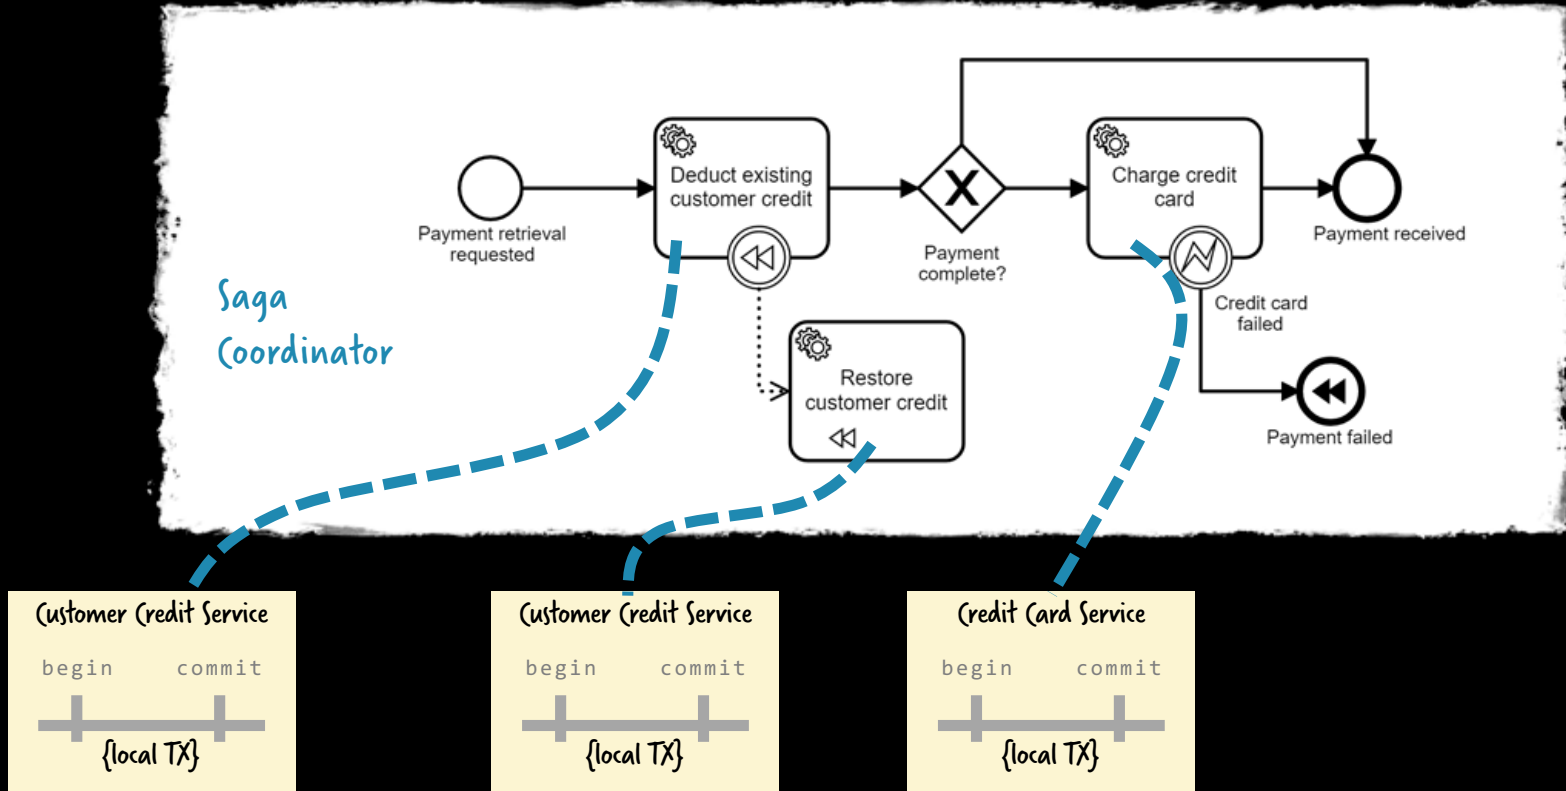
* aka Saga pattern



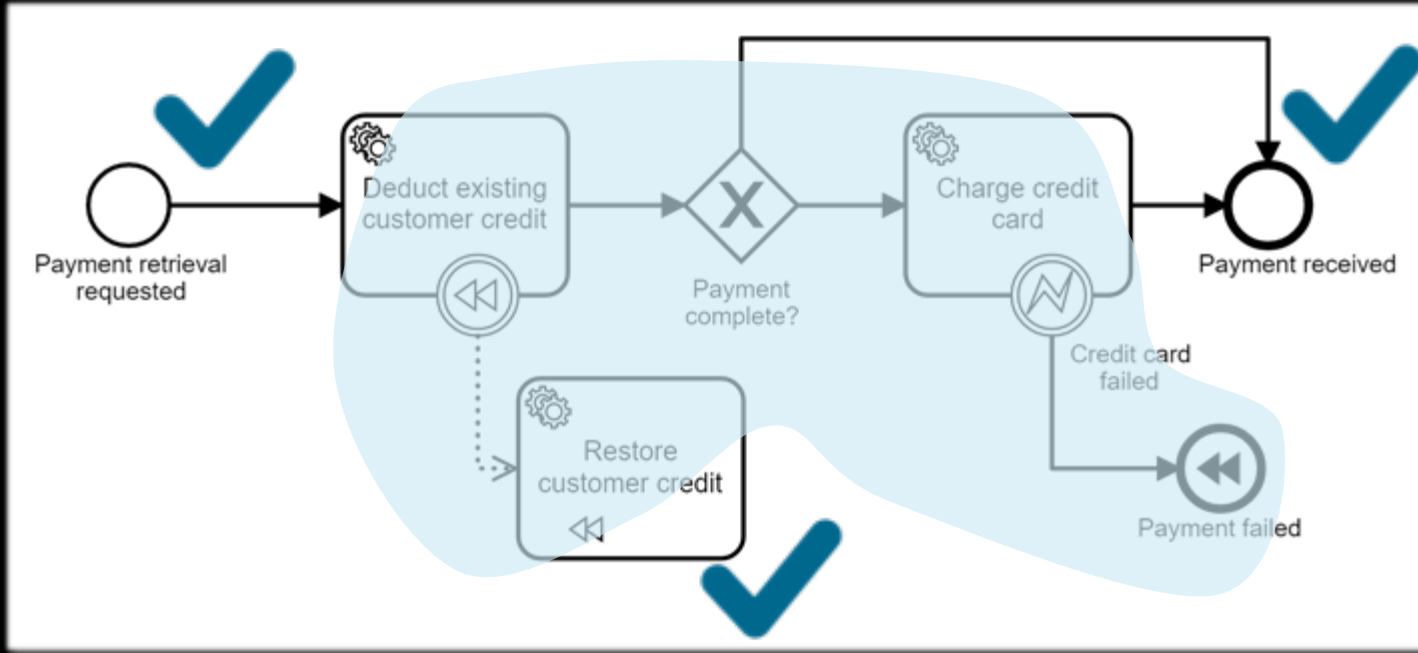
Live demo



A workflow engine can serve as Saga coordinator



Eventual consistency



Temporarily
inconsistent state

But only
temporary

No Isolation
(as in ACID)

Apologize!



PatHelland's WebLog



Memories, Guesses, and Apologies

Rate this article ★★★★★

 Pat Helland May 15, 2007

 Share 3  12  0

 5

Well, here I am blogging on the bus with my newly installed Windows Live Writer!!!
This blog is a text version of a five minute "Gong Show" presentation I did at CIDR (Conference on Innovative Database Research) on Jan 8,2007.
All computing can be considered as "M...
how computer...

<https://blogs.msdn.microsoft.com/pathelland/2007/05/15/memories-guesses-and-apologies/>

Client

has to implement

Timeout, Retry,
Compensation

Service Provider

has to offer

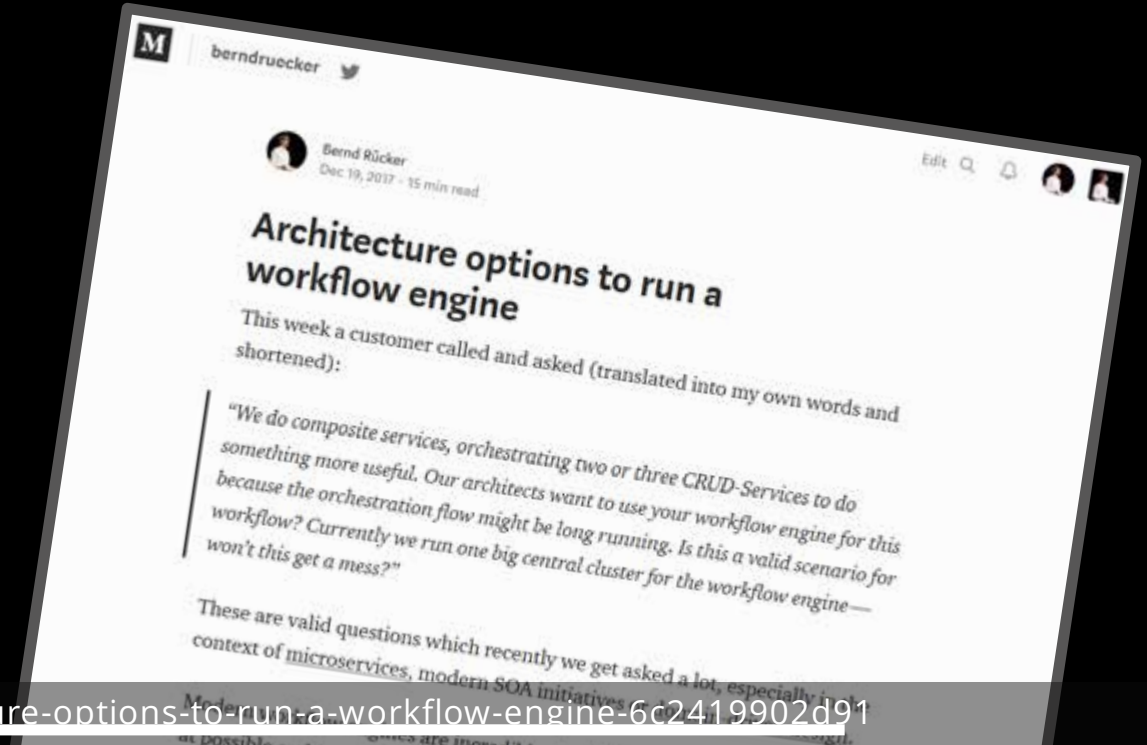
Compensation
has to implement
Idempotency

And: don't forget to apologize sometimes...



A few words
about architecture ...



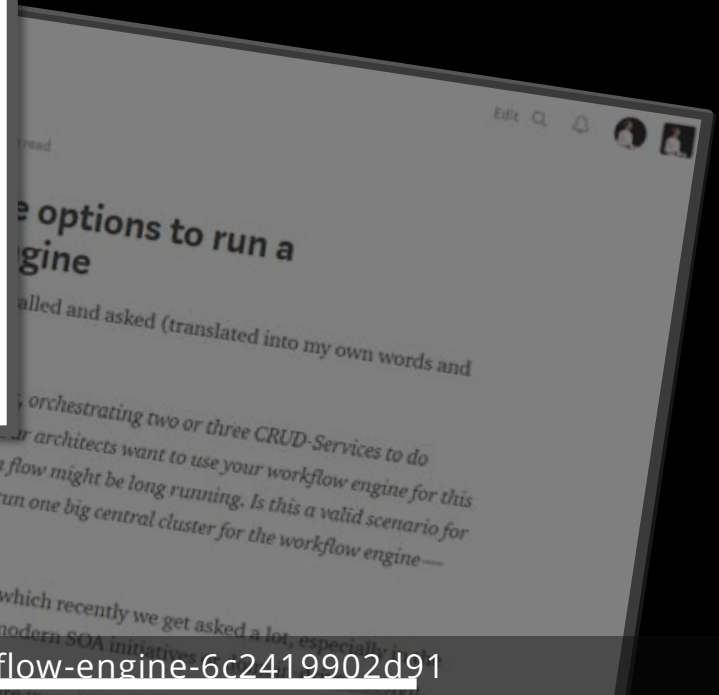
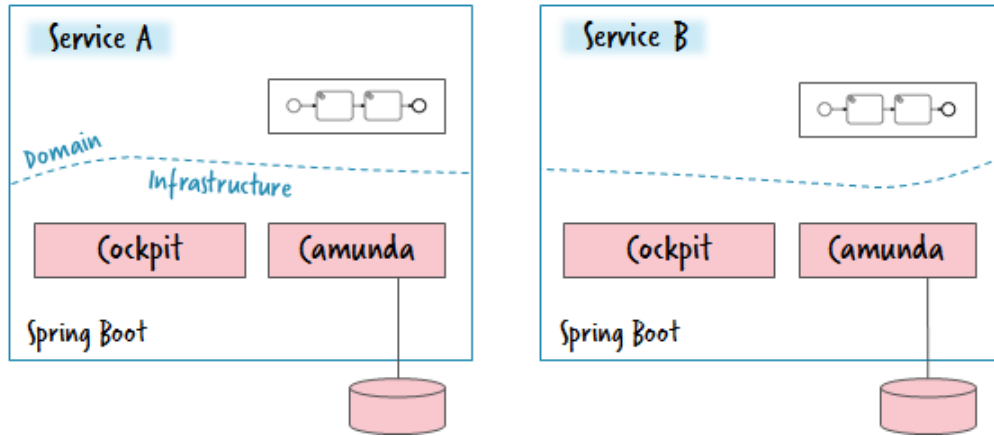
Architecture



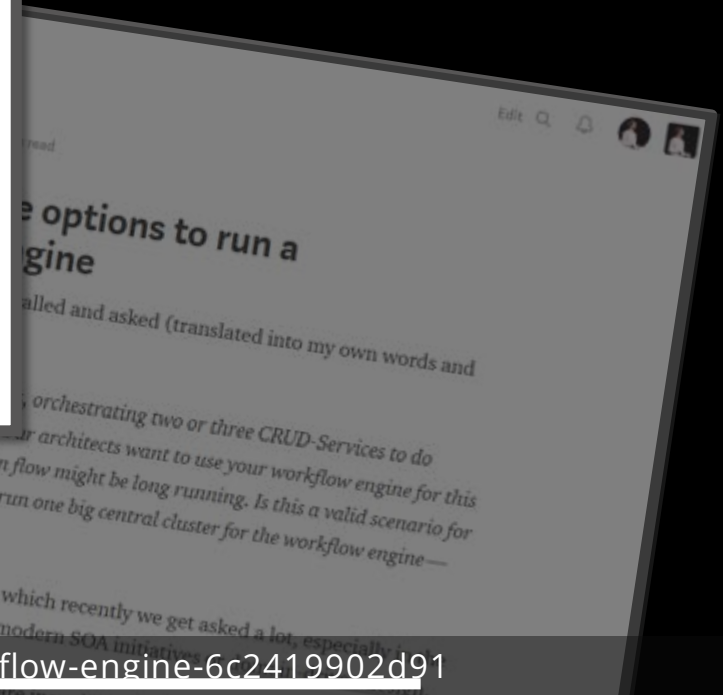
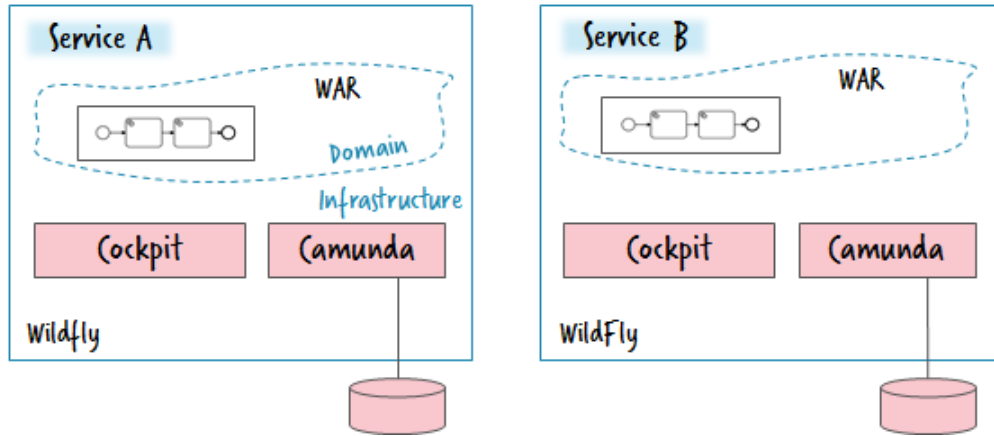
Workflows live inside service boundaries



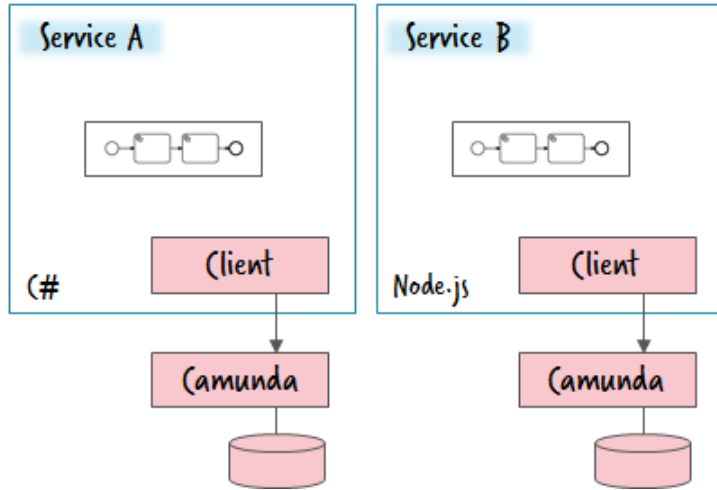
Architecture



Architecture



Architecture



Options to run a workflow engine

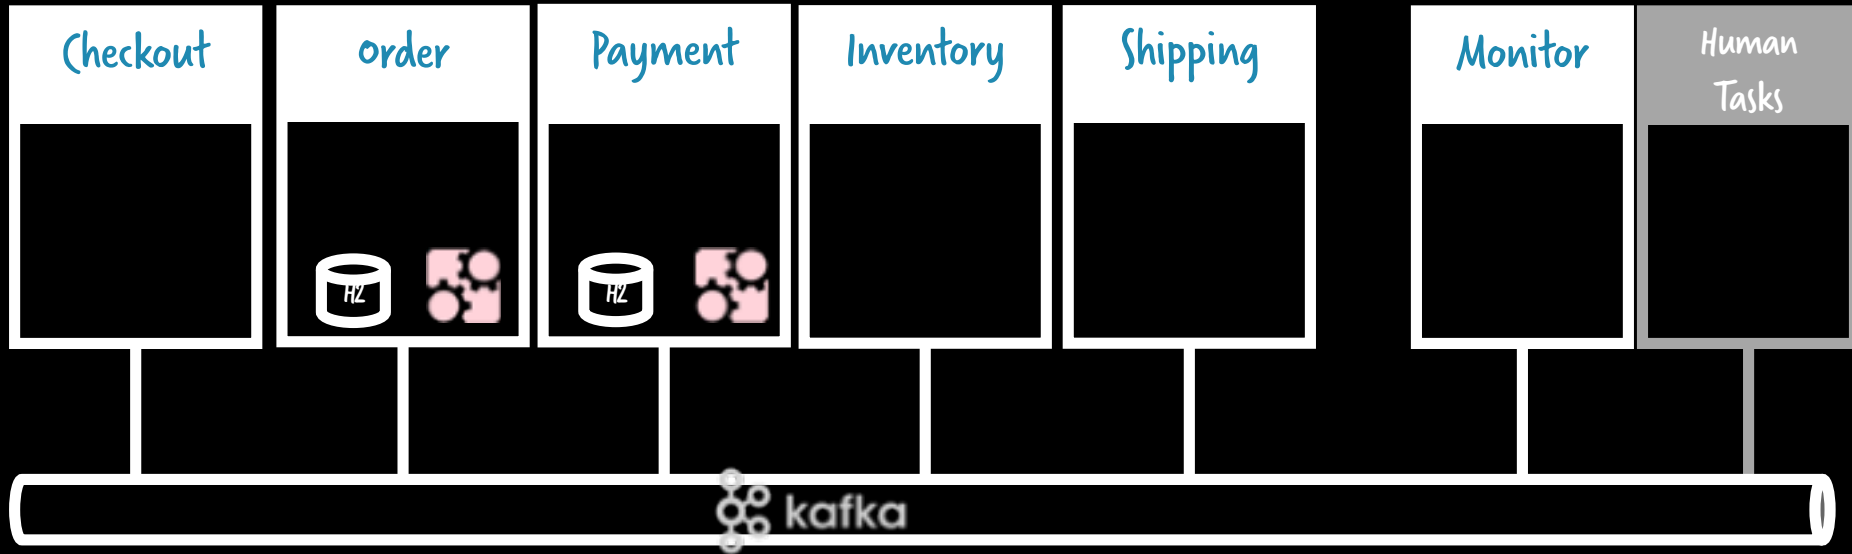
called and asked (translated into my own words and

orchestrating two or three CRUD-Services to do
or architects want to use your workflow engine for this
how might be long running. Is this a valid scenario for

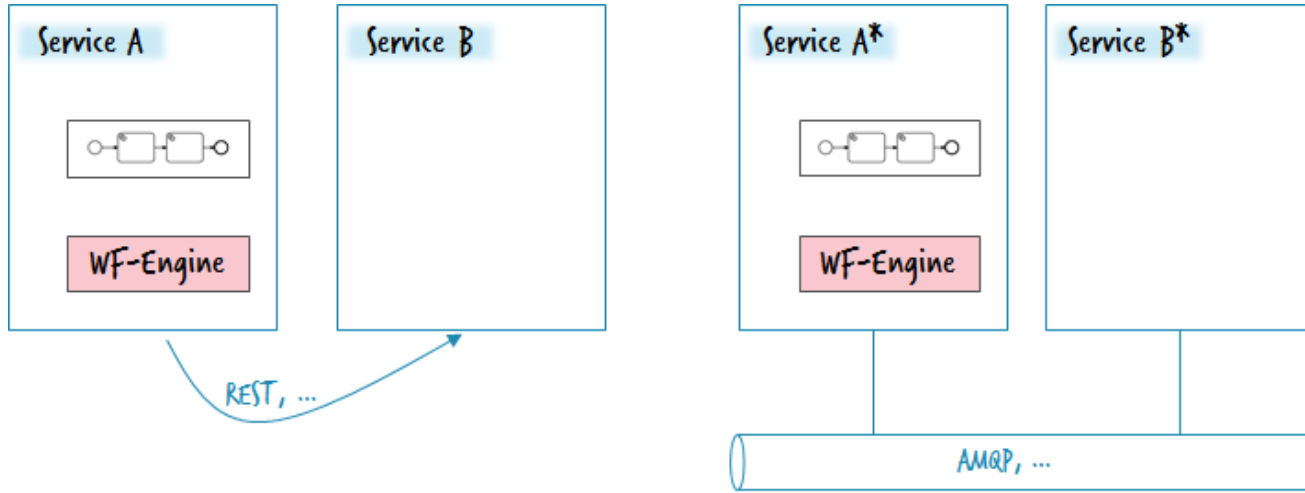
Currently we run one big central cluster for the workflow engine—
won't this get a mess?"

These are valid questions which recently we get asked a lot, especially
context of microservices, modern SOA initiatives

Event-driven example also available



Architecture



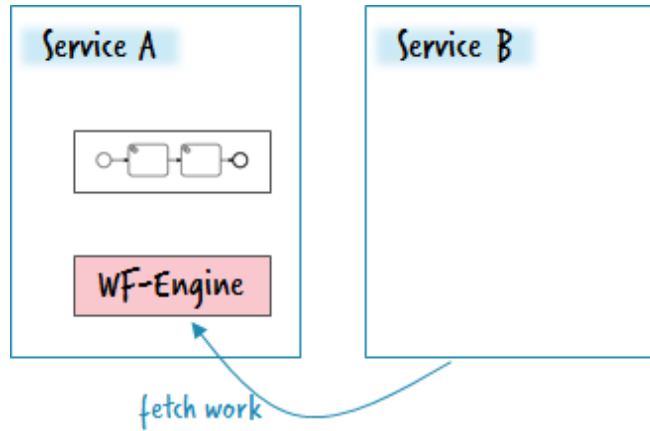
... questions which recently we get asked a lot, especially in the context of microservices, modern SOA initiatives. Modern SOA initiatives are more and more moving towards a message-oriented architecture (MOA) as opposed to a REST-based architecture. At possible, ... are more and more moving towards a message-oriented architecture (MOA) as opposed to a REST-based architecture.

Who uses a message bus?

Who has no problems
operating a message bus?

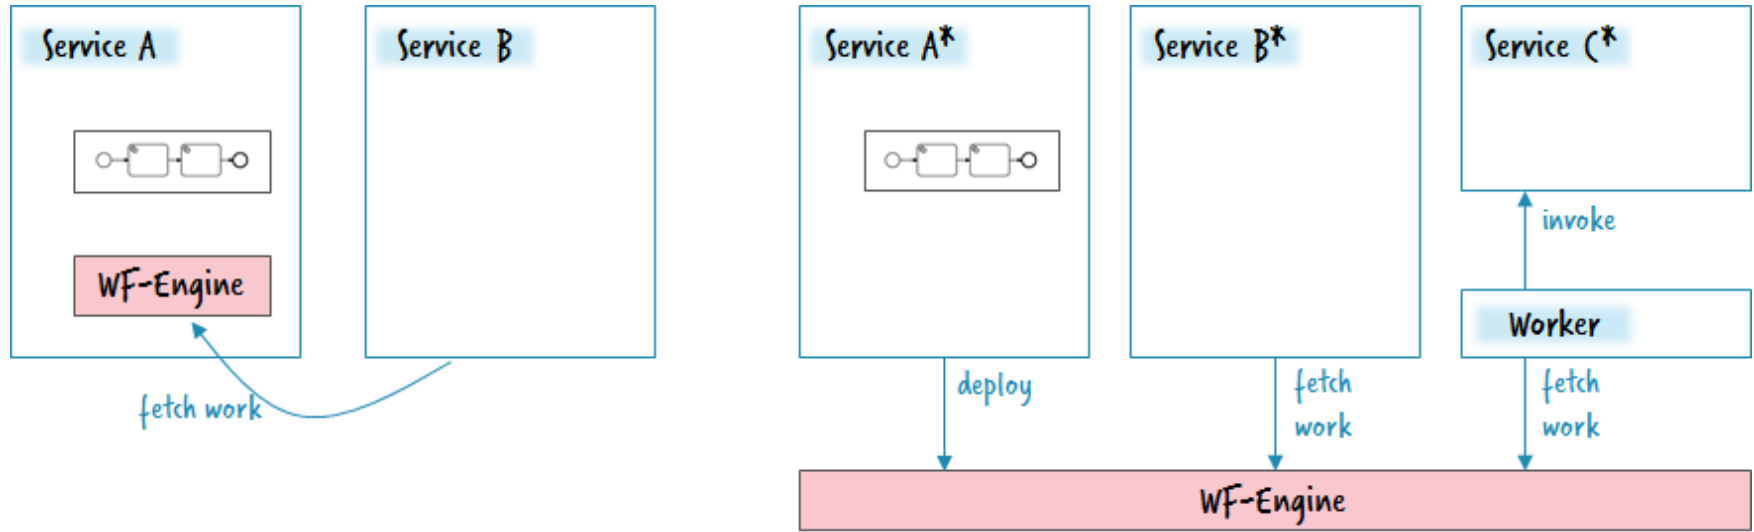
Dead messages | No context | Inaccessible payload | Hard to redeliver |
Home-grown message hospitals | ...

Architecture



...questions which recently we get asked a lot, especially in the
context of microservices, modern SOA initiatives

Architecture



...questions which recently we get asked a lot, especially in the context of microservices, modern SOA initiatives...

Live demo



Reality check



”

Before mapping processes explicitly with BPMN, the truth was buried in the code and nobody knew what was going on.

Jimmy Floyd, 24 Hour Fitness

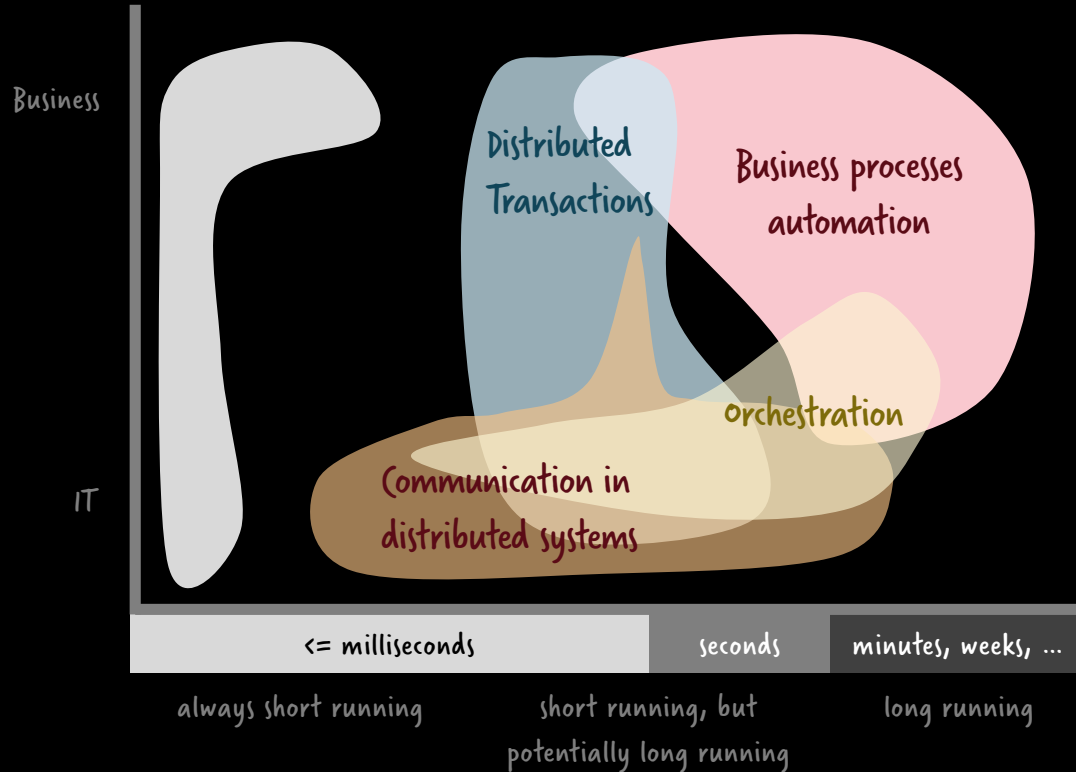
CONTRIBUTED / TOP STORIES / GLOBAL

5 Workflow Automation Use Cases You Might Not Have Considered

9 Apr 2018 3:00am, by Bernd Rucker



Use cases for workflow automation



Be aware of complexity of distributed systems

Know strategies and tools to handle it

e.g. (circuit breaker (*Hystrix*))

Workflow engine for stateful retry, waiting, timeout
and compensation (*Camunda*)

Thank you!

