



Cloud Computing Patterns @ OOP 2017

Case Study and Discussion

<http://www.cloudcomputingpatterns.org/oop17/> ←

GET SLIDES HERE!

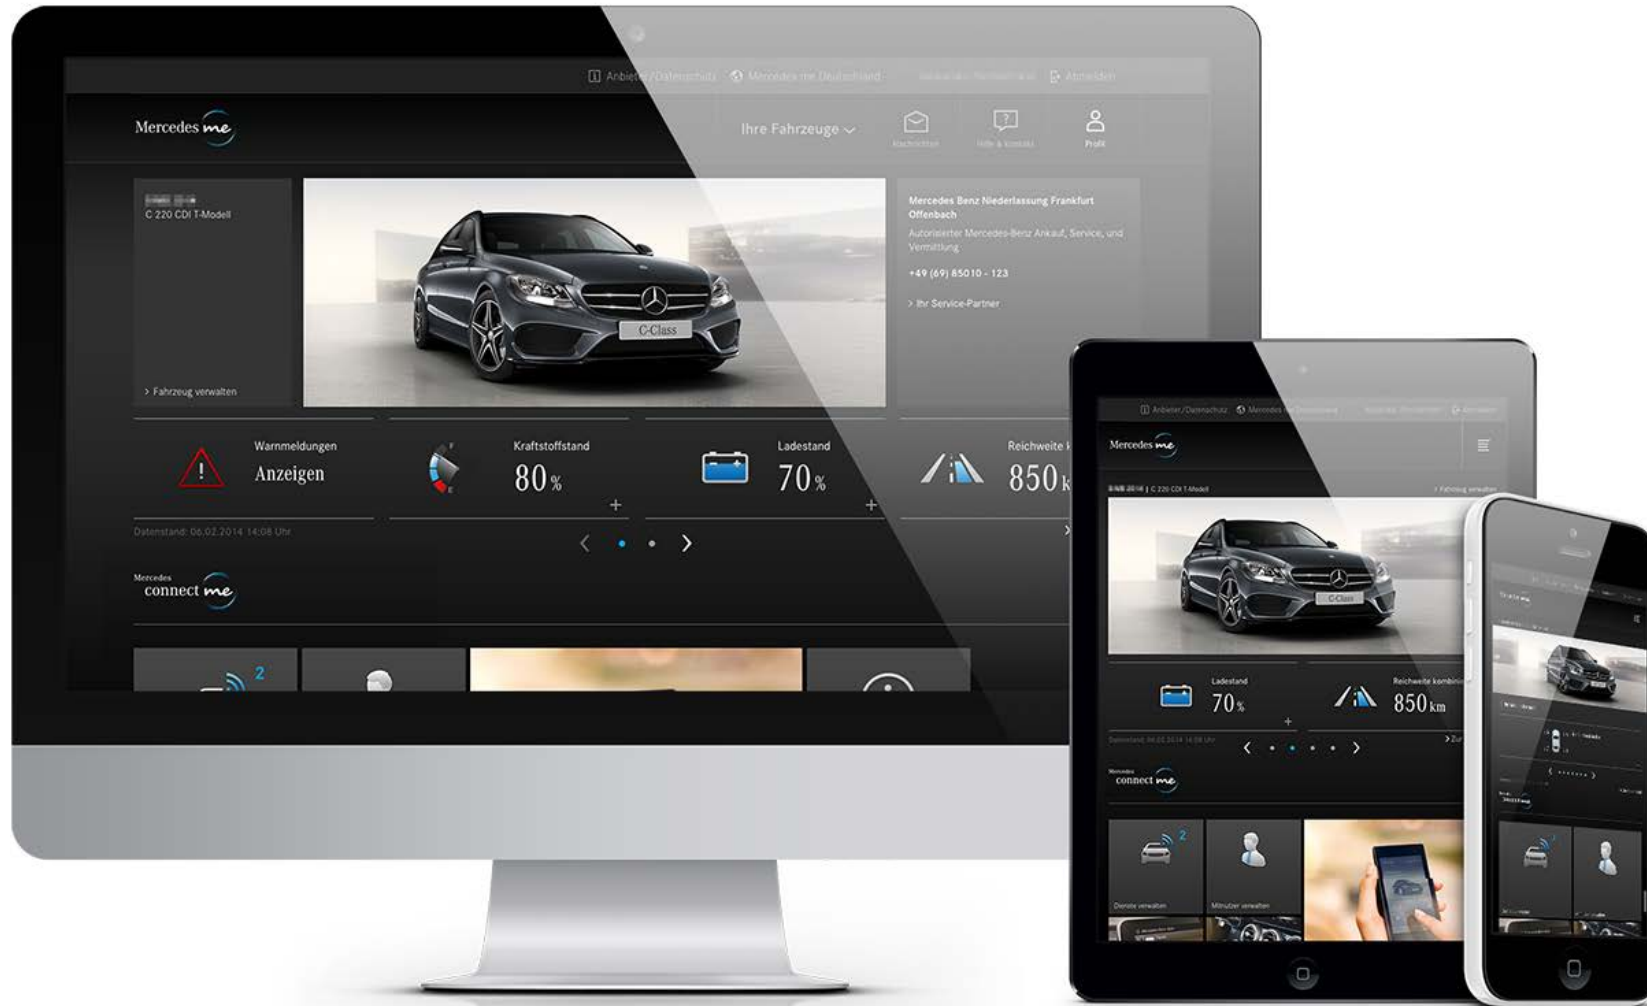
Dr. Christoph Fehling

christoph.fehling@daimler.com

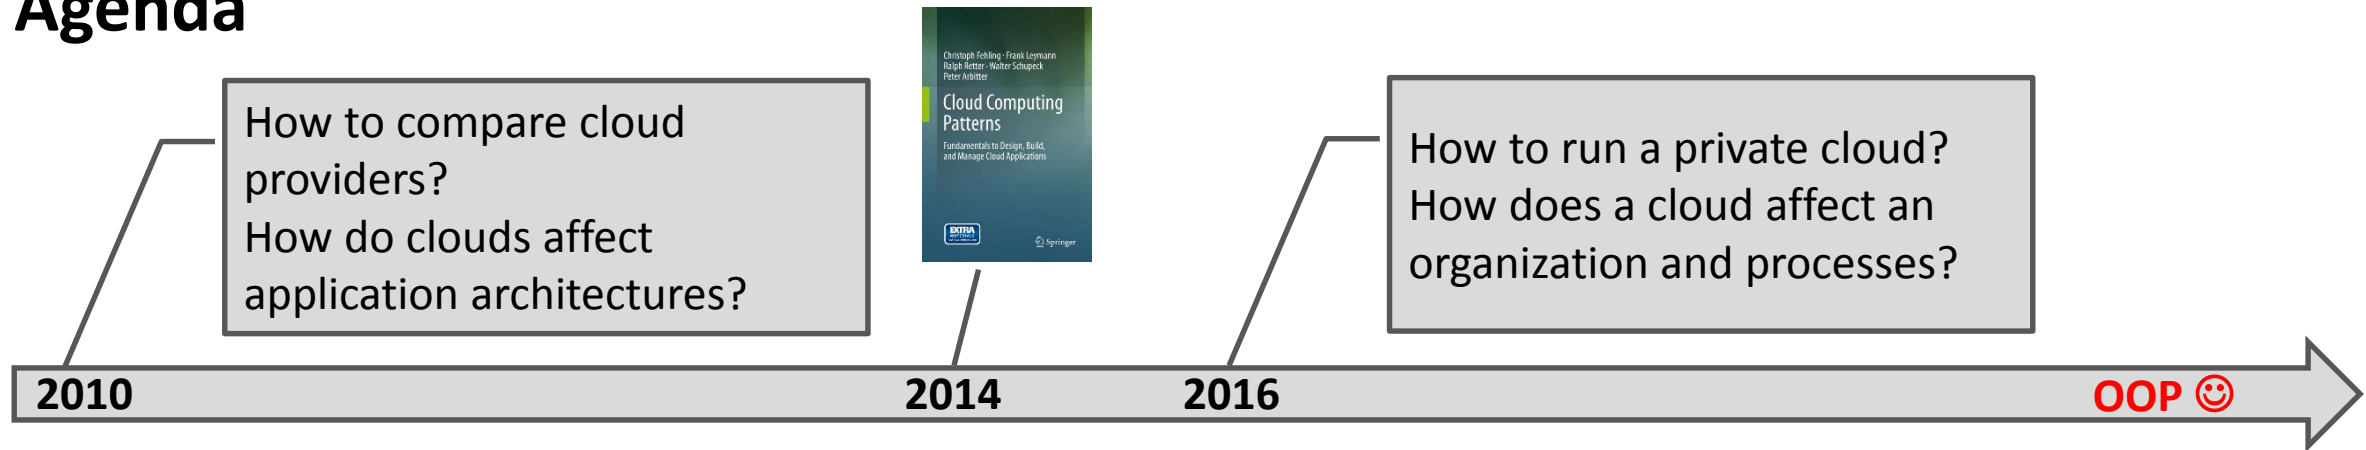
Daimler AG

@ccpatterns

Mercedes Me



Agenda



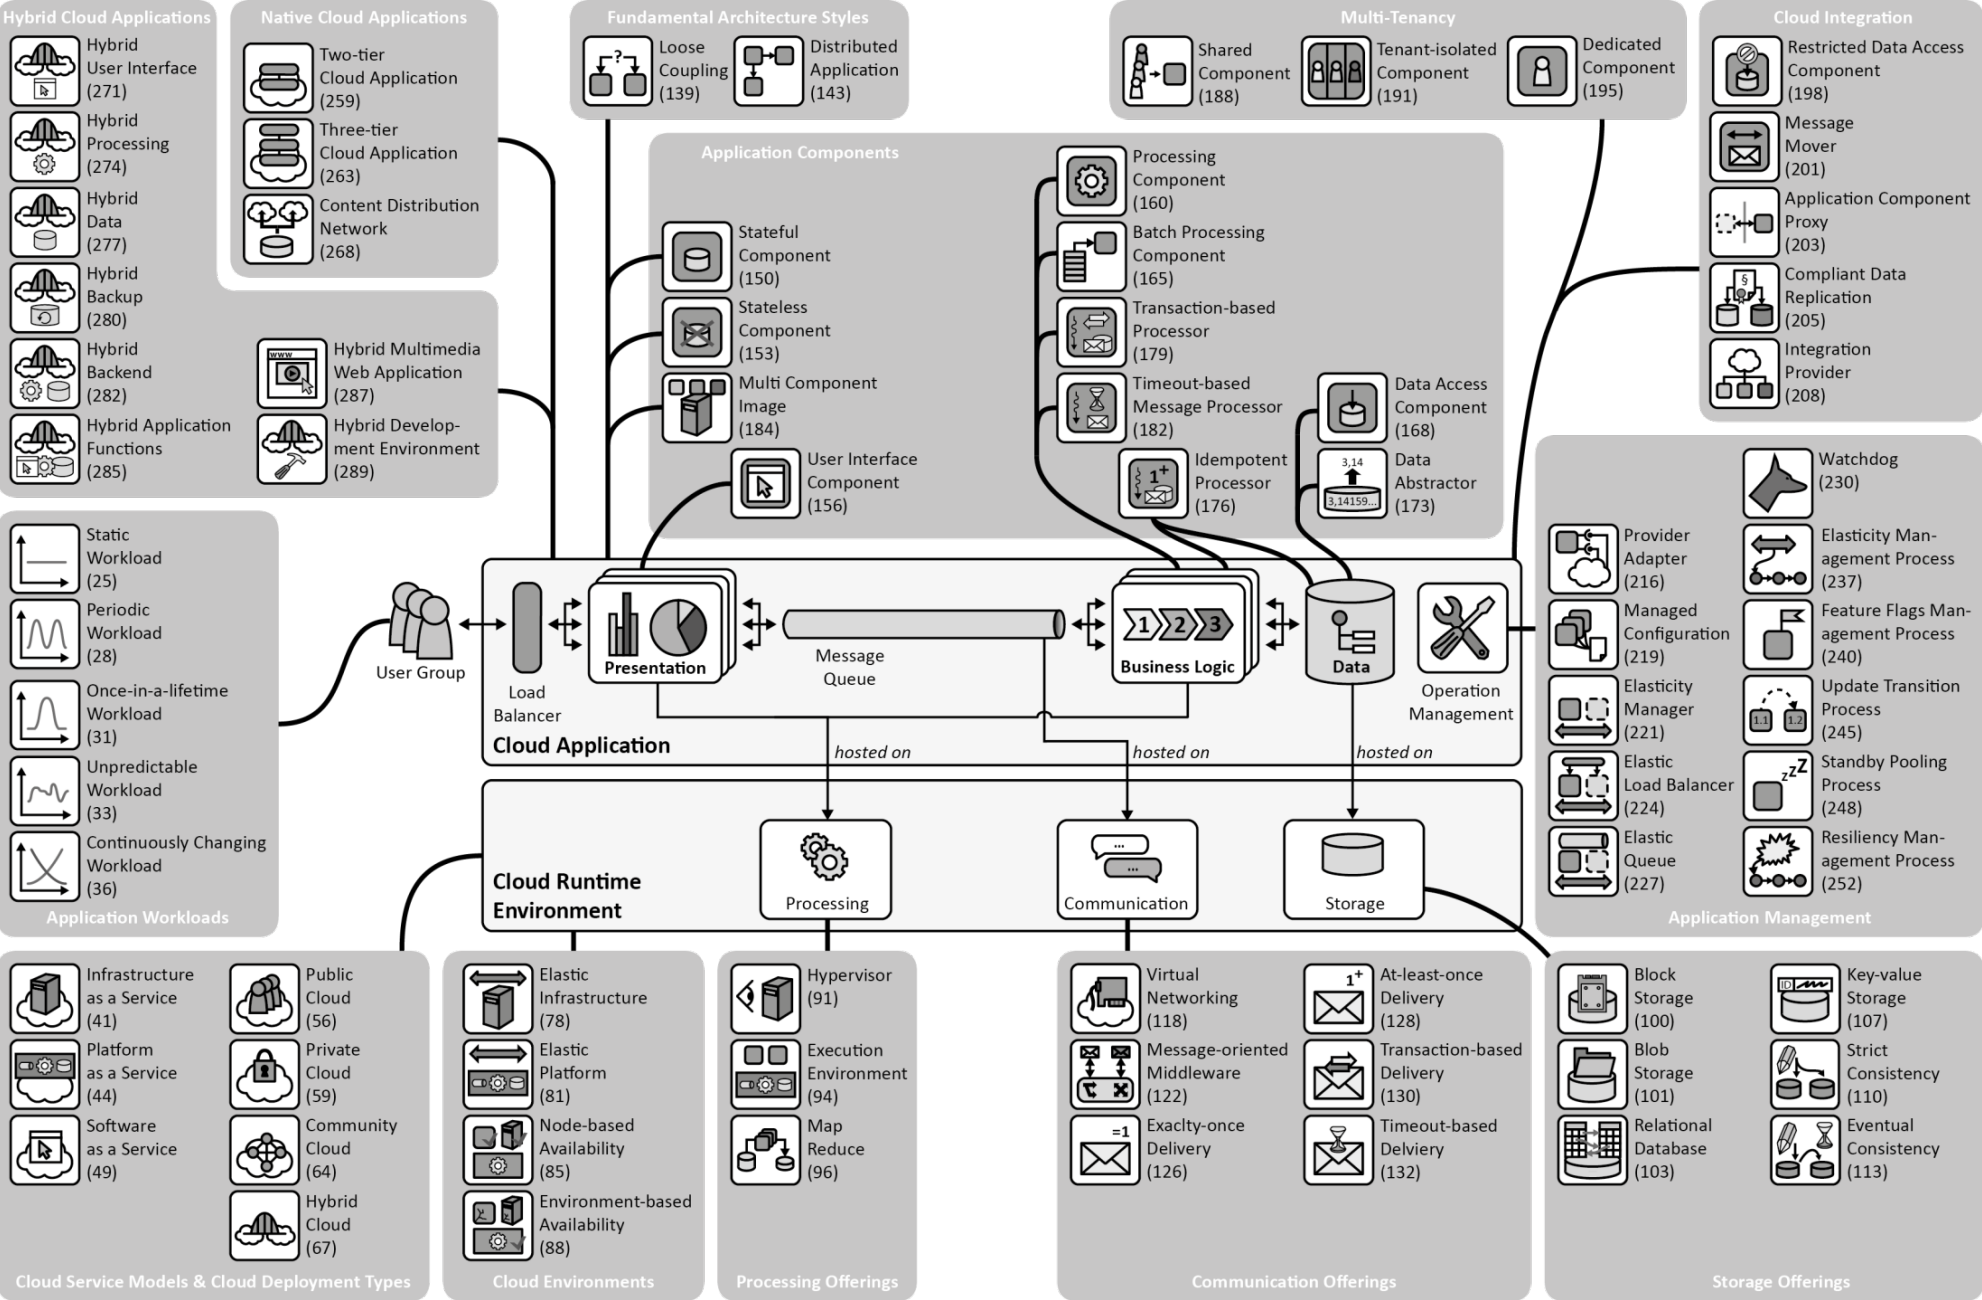
Part 1: Cloud Computing Patterns @ Mercedes Me

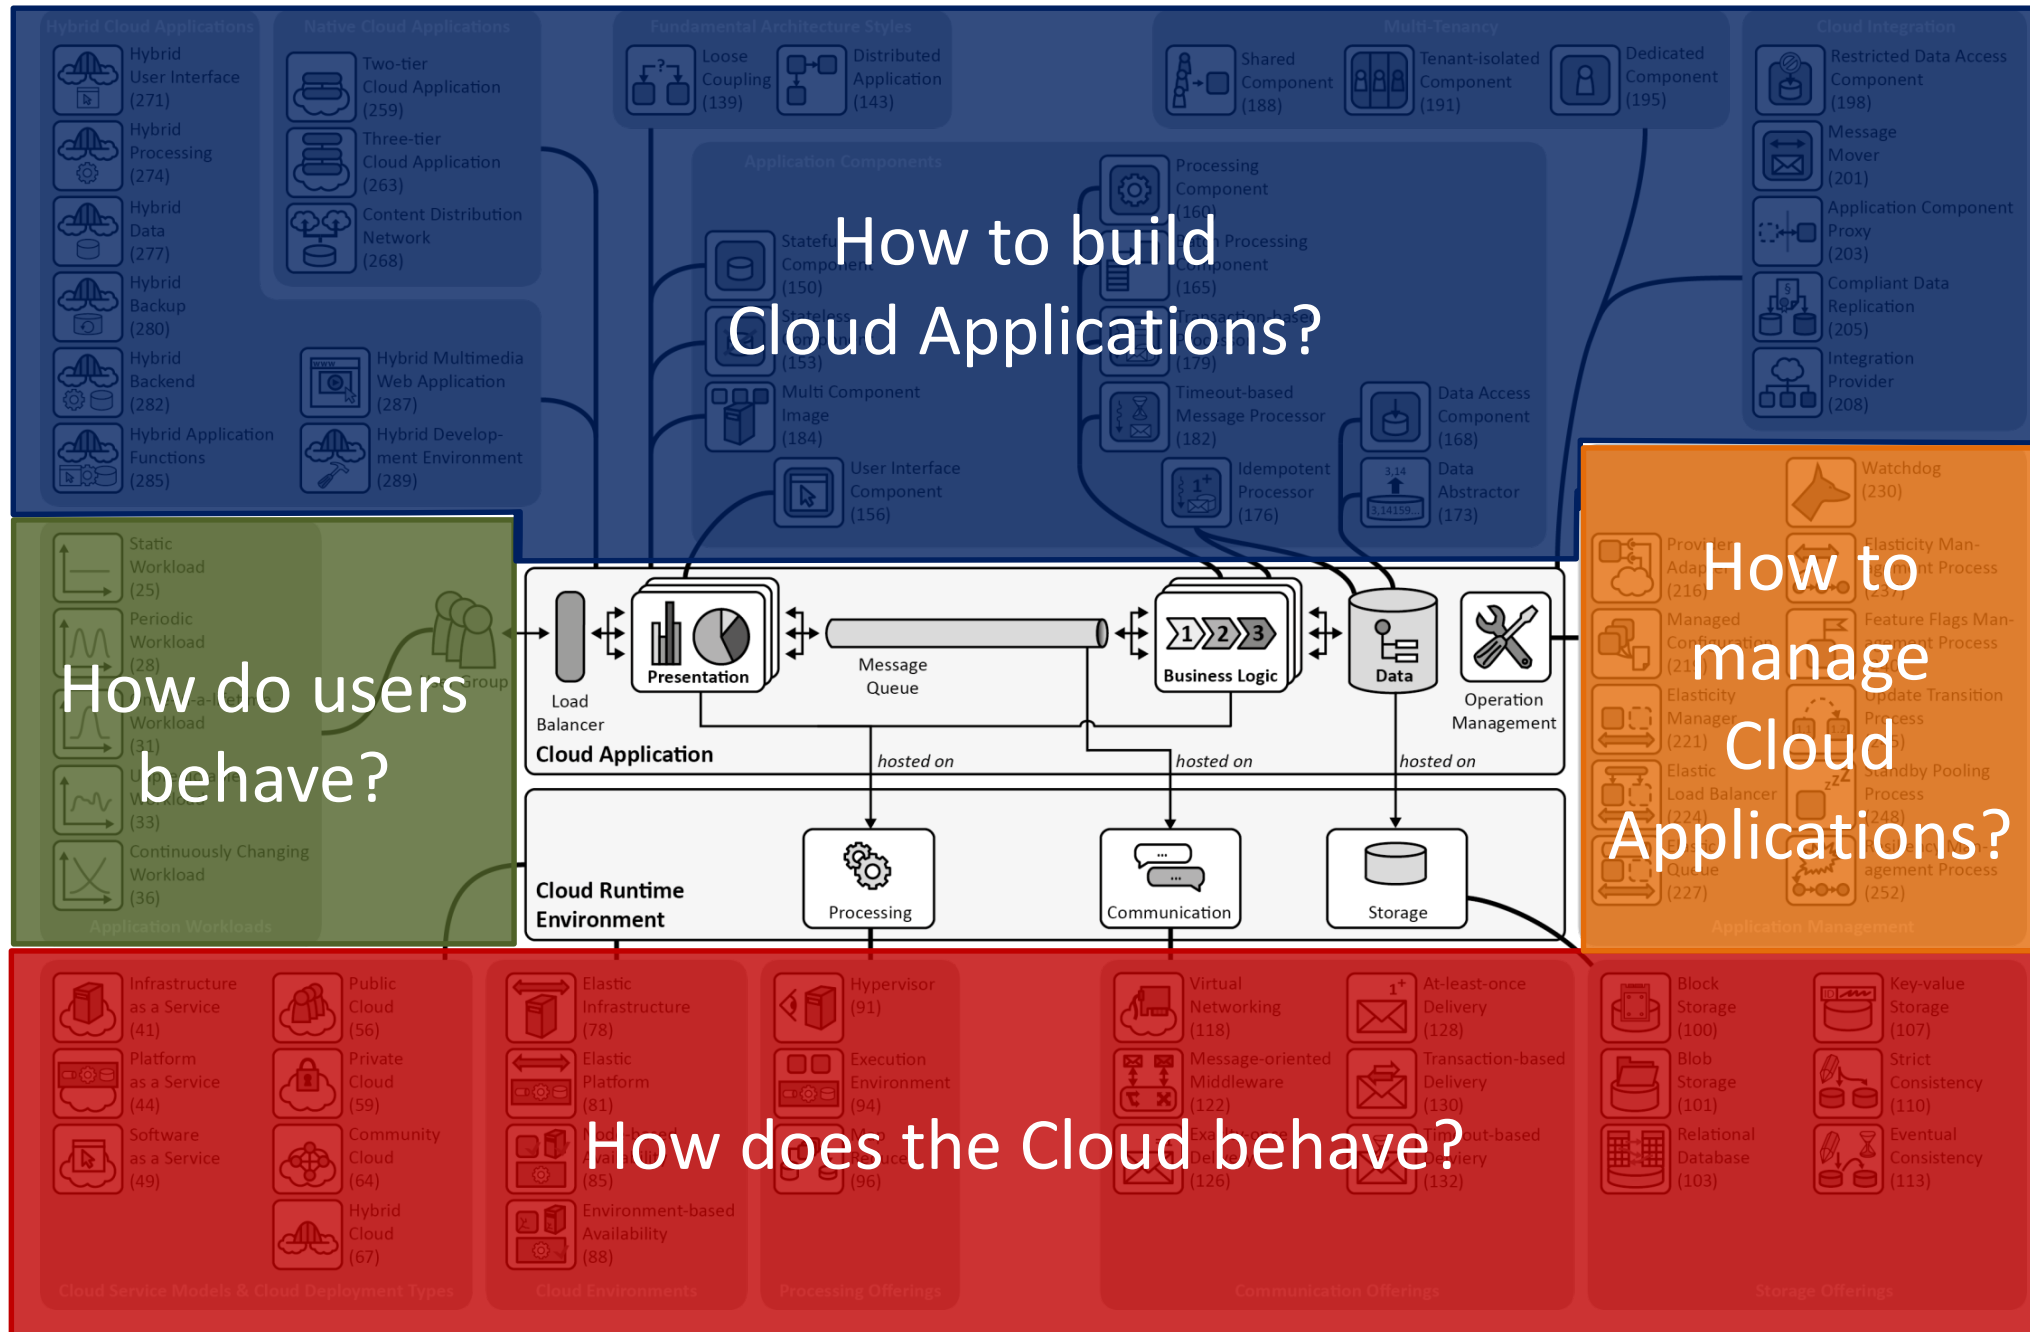
- What are the **cloud computing patterns**?
- What are the **properties of a cloud-native application**?
- How does a cloud-native application look like at **Mercedes Me**?

Part 2: the Non-technical “Stuff” (Discussion)


- How does cloud computing affect **procurement processes**?
- Why are **licenses** of cloud products so problematic?
- How does cloud computing affect **organizational hierarchies**?









IDEAL Cloud Application Properties




Isolated State: most application components should be *stateless*. They do not handle:
Session State: state of the communication with the application.
Application State: business data handled by the application.




Distribution: cloud applications are split up into multiple components...
... to utilize multiple cloud resources.
... because the cloud itself is a large distributed system.



Elasticity: cloud applications are scaled by adjusting resource numbers (*scaling out*) – not by *scaling up*:
Scale out: Increase performance by adding more resources.
Scale up: Increase performance by improving existing resources.



Automated Management: management tasks during runtime have to be handled quickly.
Example: Cost reduction by adjusting pay-per-use resource numbers automatically.
Example: automatic reaction to resource failures.




Loose Coupling: application components should not influence each other regarding factors such as availability, data format, data exchange rate.
Example: failure of one application component does not cause failure of other components.




Part 1: Cloud Computing Patterns @ Mercedes Me

Mercedes Me




Workload Patterns @ Mercedes Me




Static Workload

IT resources with an **equal utilization over time** experience static workload.




Periodic Workload

IT resources with a **peaking utilization at reoccurring time intervals** experience periodic workload.




Once-in-a-Lifetime Workload

IT resources with an equal utilization over time disturbed by a **strong peak occurring only once** experience once-in-a-lifetime workload.



Unpredictable Workload

IT resources with a **random and unforeseeable utilization** over time experience unpredictable workload.



Continuously Changing Workload


IT resources with a **utilization that grows or shrinks constantly** over time experience continuously changing workload.



Cloud-Native Application

- ☐ Isolated State
- ☐ Distribution
- ☒ Elasticity
- ☒ Automated Management
- ☐ Loose Coupling


Car Status Updates during a single Day



- **Peak workload:** beginning and end of each day
→ Rush hour
- **Low workload:** during each night
→ People are sleeping



Data Handling and Data Abstraction @ Mercedes Me




Strict Consistency

Data is stored at different locations to improve response time and to avoid data loss in case of failures while **consistency of replicas is ensured at all times**.




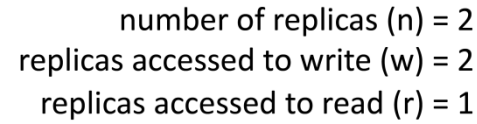
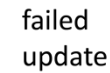
Eventual Consistency

Performance and availability of data in case of network partitioning are enabled by ensuring **data consistency eventually and not at all times**.





Cloud-Native Application

- ☒ Isolated State
- ☐ Distribution
- ☐ Elasticity
- ☐ Automated Management
- ☒ Loose Coupling


$$n \leq r + w$$

$$n > r + w$$


 Loose Coupling








Data Abtractor

Data is **abstracted to inherently support eventually consistent data** storage through the use of abstractions and approximations.

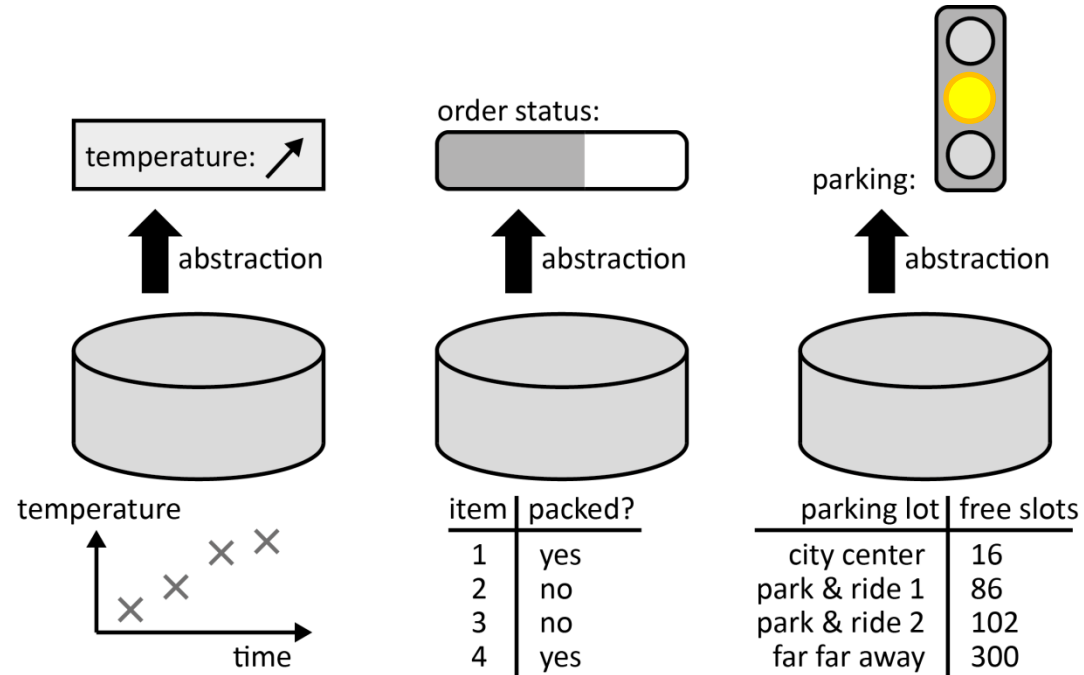


Cloud-Native Application

- ☒ Isolated State
- ☐ Distribution
- ☐ Elasticity
- ☐ Automated Management
- ☒ Loose Coupling




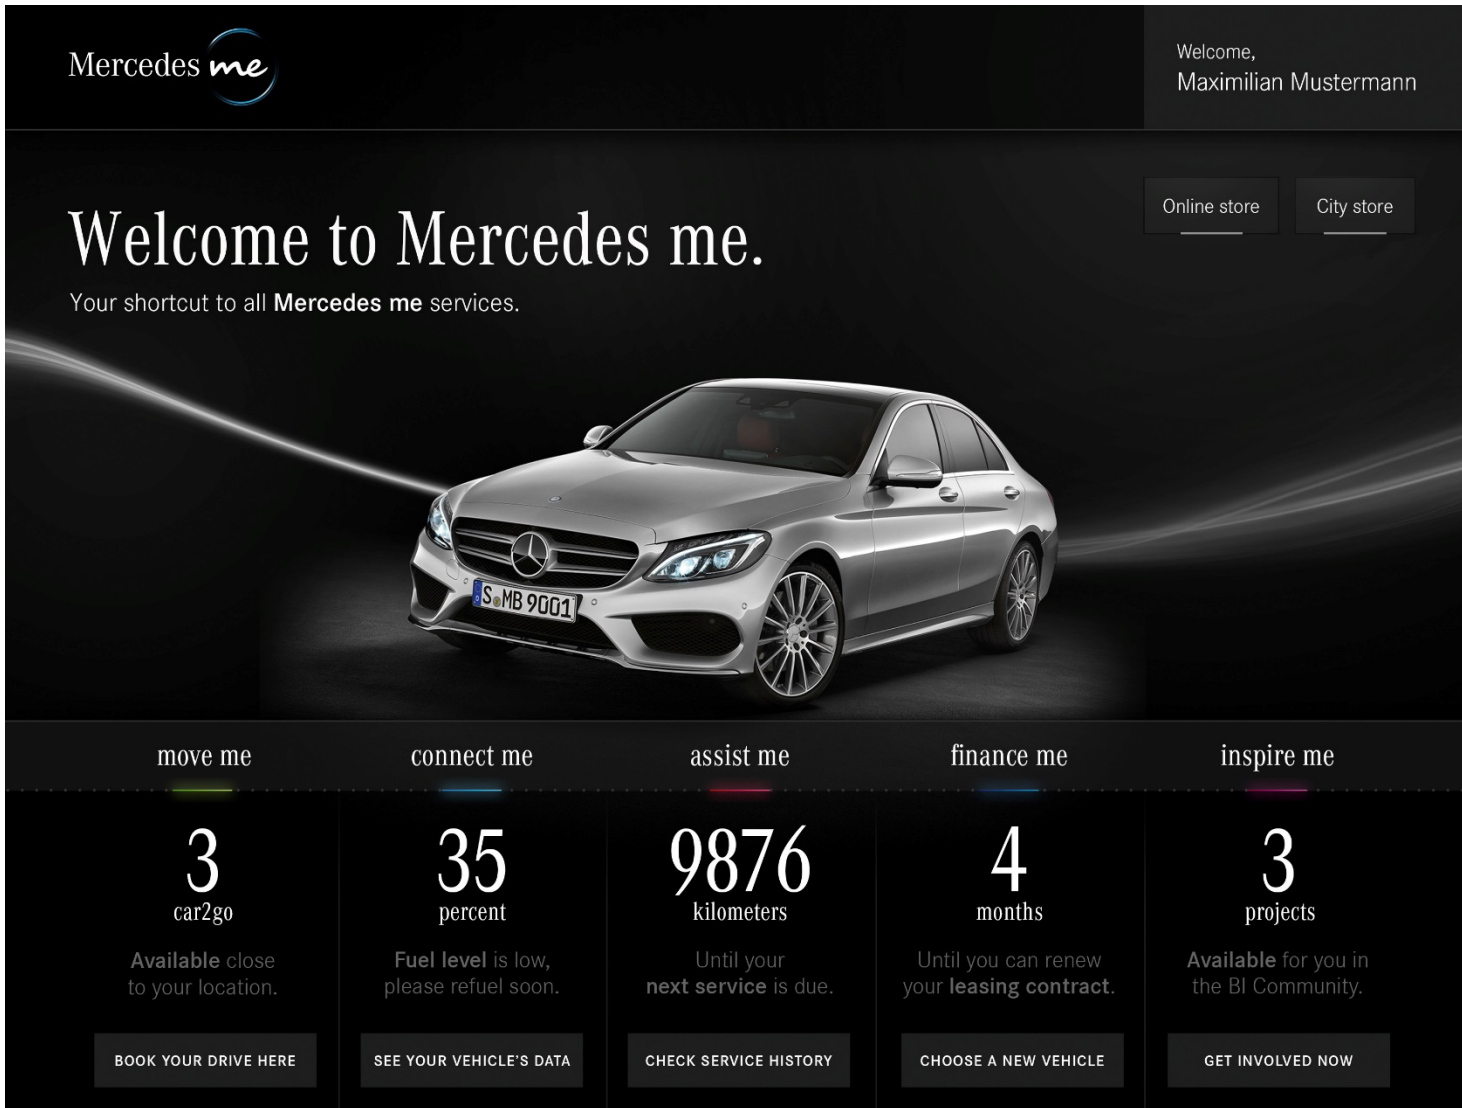
Data Abtractor



Cloud-Native Application


- ☒ Isolated State
- ☐ Distribution
- ☐ Elasticity
- ☐ Automated Management
- ☒ Loose Coupling

Mercedes Me Portal



A lot of eventual consistent data:

- Odometer
- Fuel level
- Service intervals
- Errors shown in the car


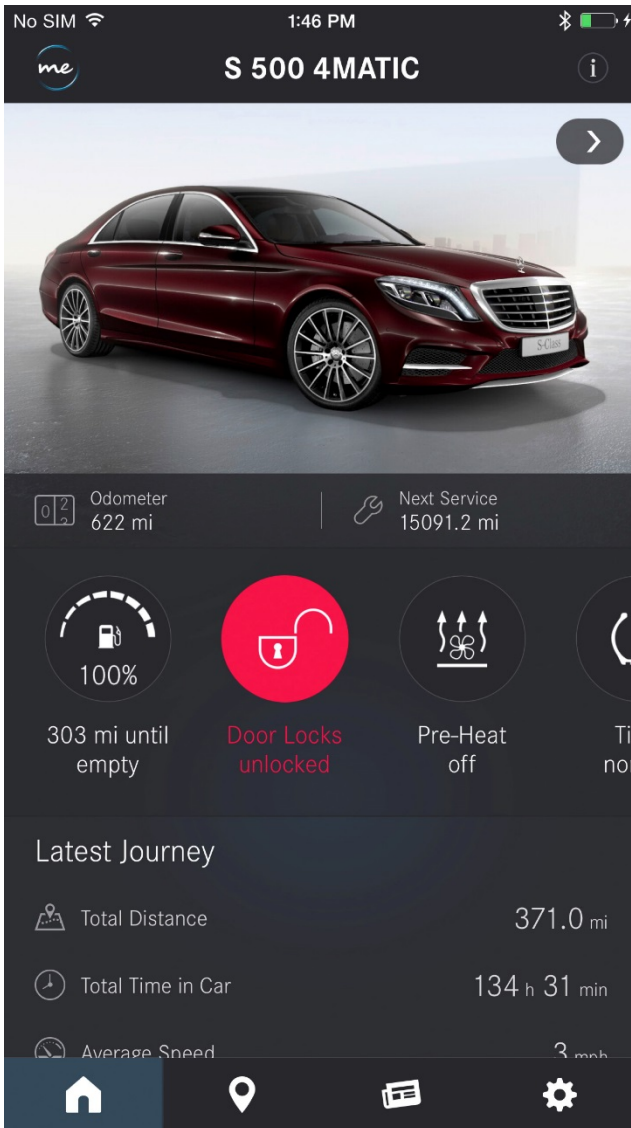


If the connectivity of the car is limited (parking garage), an obsolete status is displayed!

But can all data be treated this way?




Mercedes Me App



Same as in the portal:

- Odometer
- Fuel level
- Service intervals
- Errors shown in the car




But certain interactions **must not be abstracted!**

- Door lock status
- Heating status
- Engine status (cooling)


→ **Data abstraction and eventual consistency has to be evaluated for each data element!**

Mercedes Me Microservice Template





Node-based Availability

A cloud provider guarantees the **availability of nodes**, such as individual virtual **servers**, **middleware** components or hosted **application components**.



Environment-based Availability


A cloud provider guarantees the **availability of the environment hosting** individual nodes, such as virtual **servers** or hosted **application components**.






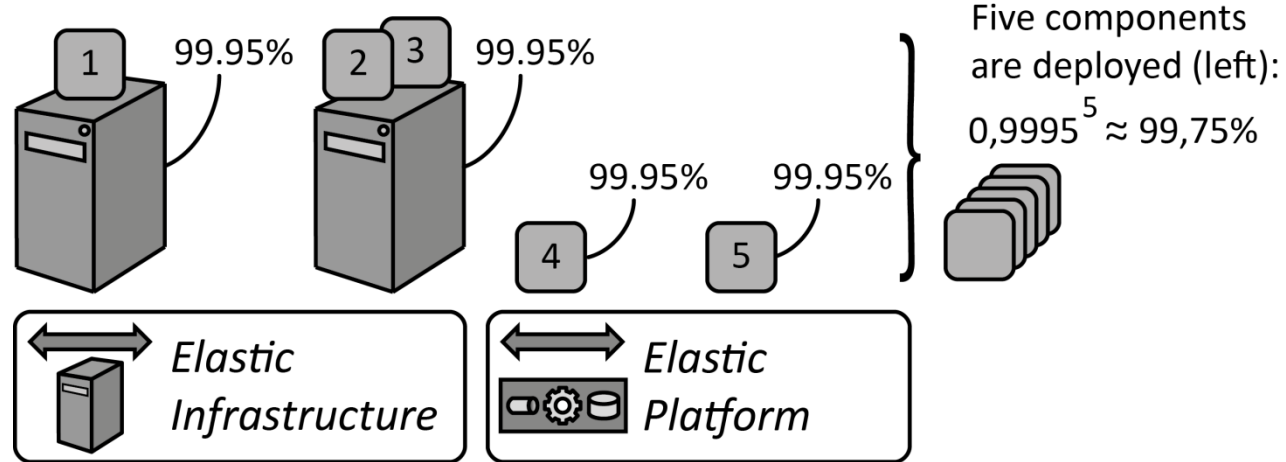
Cloud-Native Application

- ☐ Isolated State
- ☒ Distribution
- ☐ Elasticity
- ☒ Automated Management
- ☐ Loose Coupling






Node-based Availability






Environment-based Availability




Cloud-Native Application

- ☐ Isolated State
- ☒ Distribution
- ☐ Elasticity
- ☒ Automated Management
- ☐ Loose Coupling



User Interface Component





Synchronous user interfaces are accessed by *humans*, while *application-internal* interaction is realized *asynchronously* to ensure loose coupling.



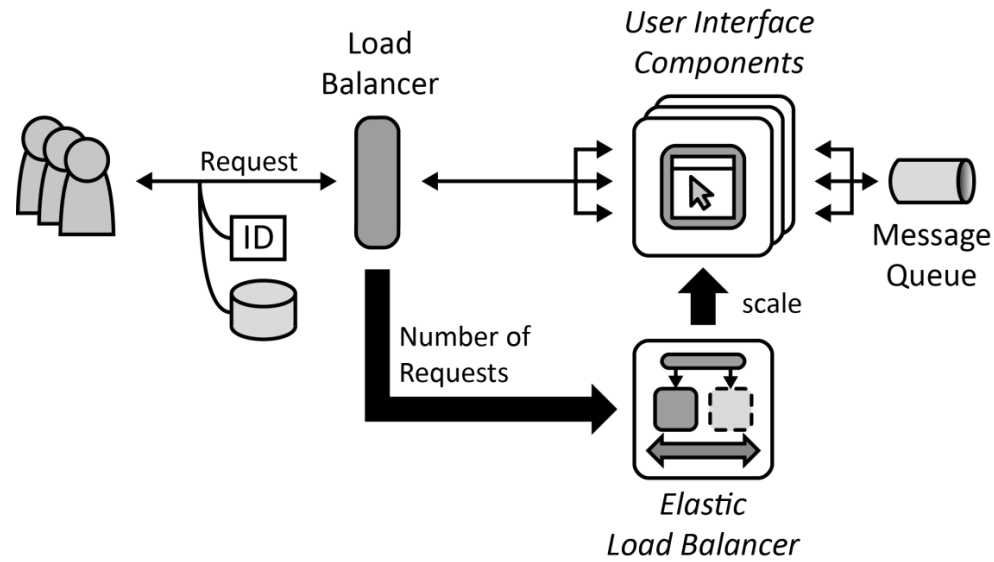
Cloud-Native Application

- ☐ Isolated State
- ☒ Distribution
- ☒ Elasticity
- ☒ Automated Management
- ☐ Loose Coupling









User Interface Component



Cloud-Native Application

- ☐ Isolated State
- ☒ Distribution
- ☒ Elasticity
- ☒ Automated Management
- ☐ Loose Coupling



Processing Component





Processing functionality is handled by [elastically scaled components](#).



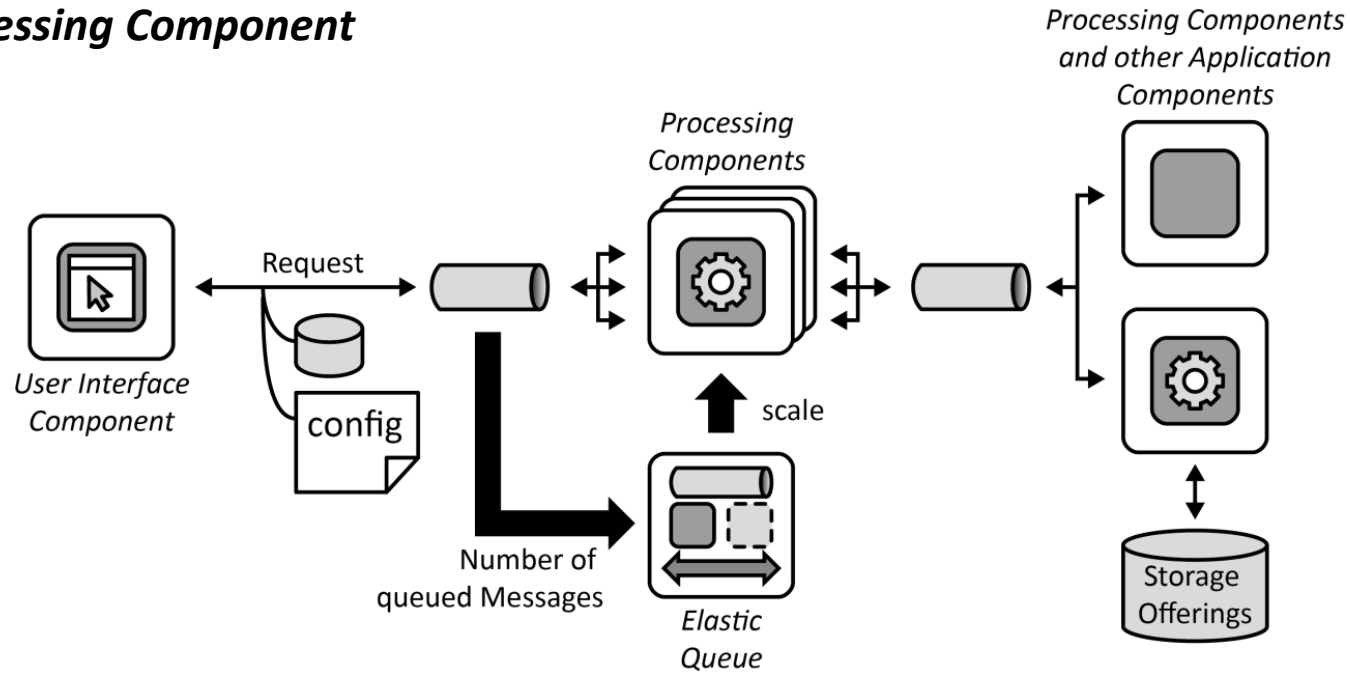
Cloud-Native Application

- ☐ Isolated State
- ☒ Distribution
- ☒ Elasticity
- ☒ Automated Management
- ☐ Loose Coupling





Processing Component

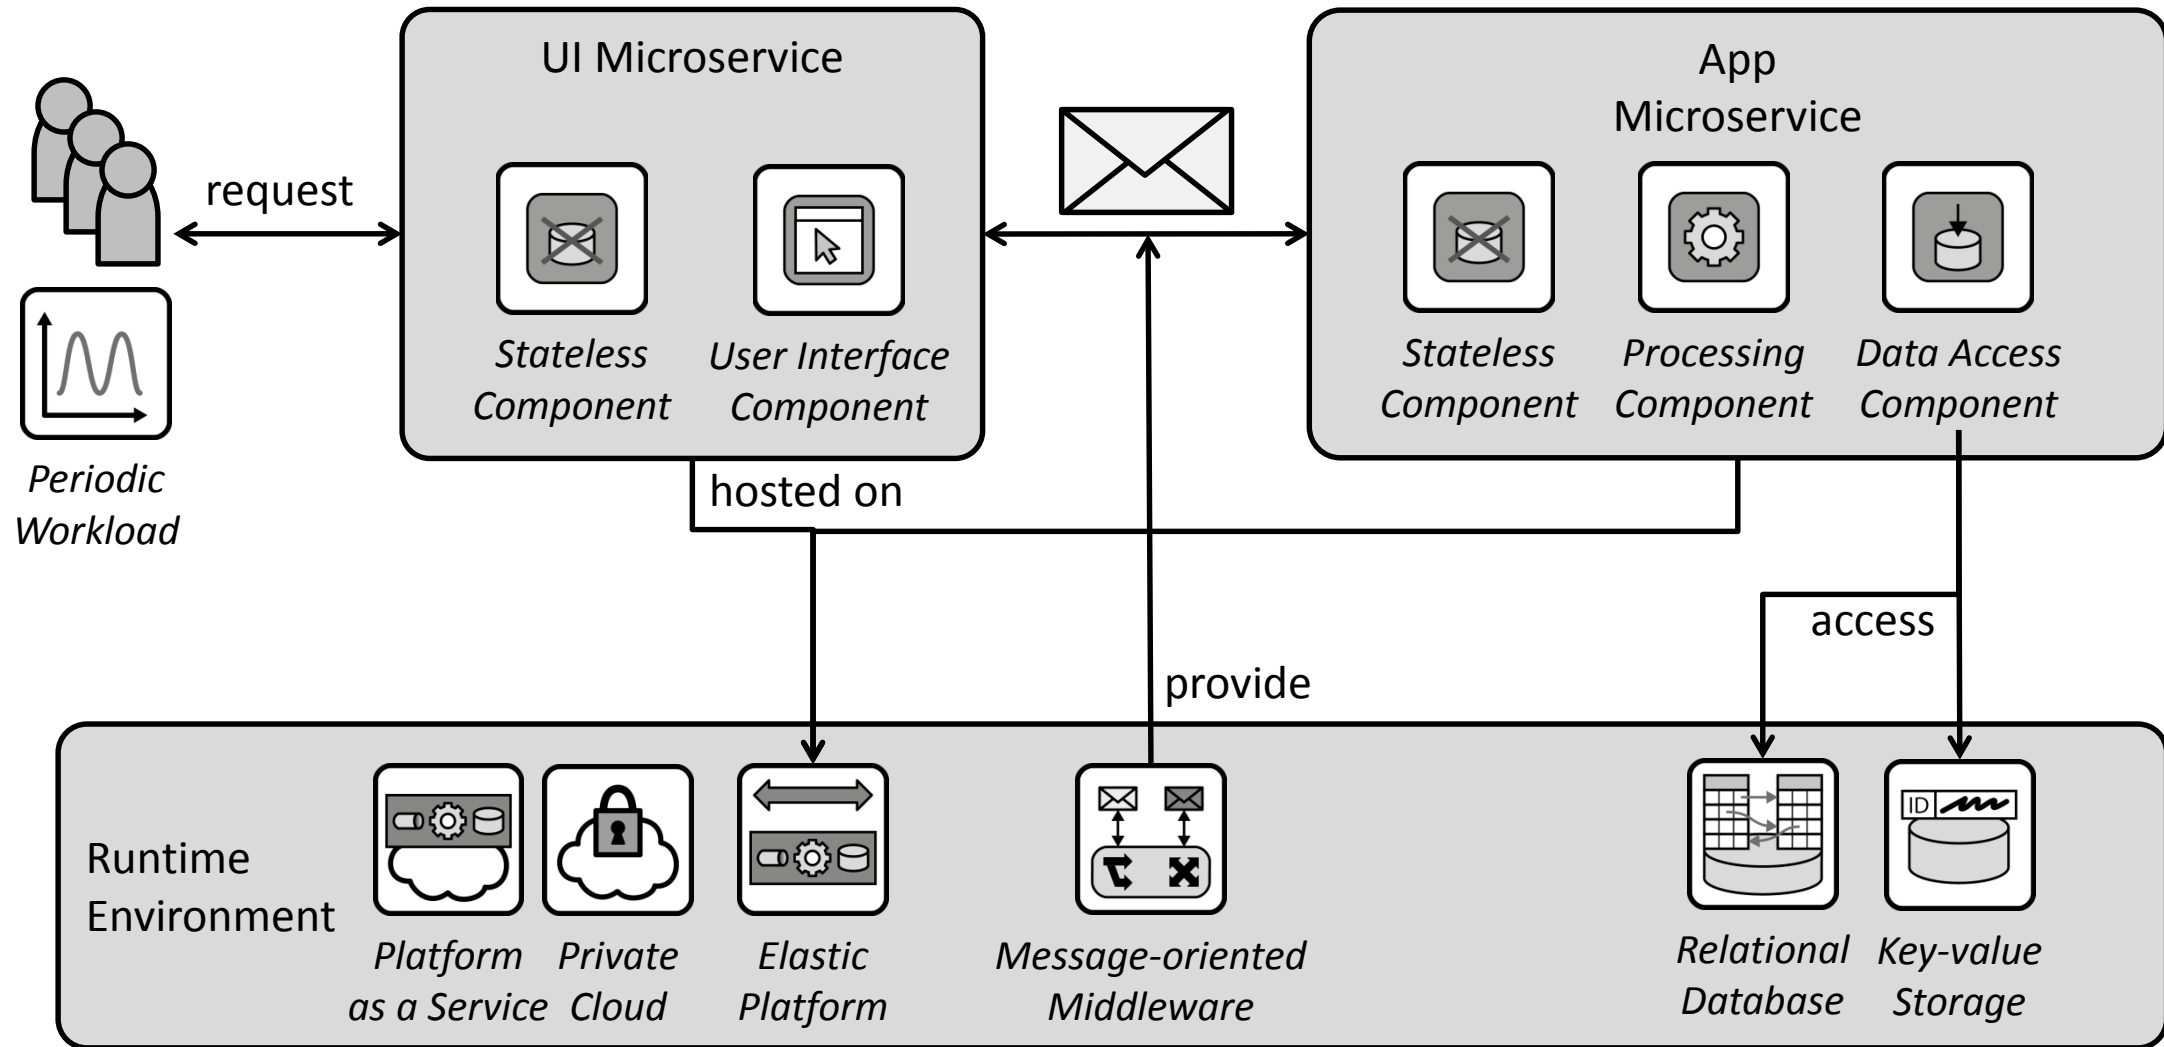


Cloud-Native Application

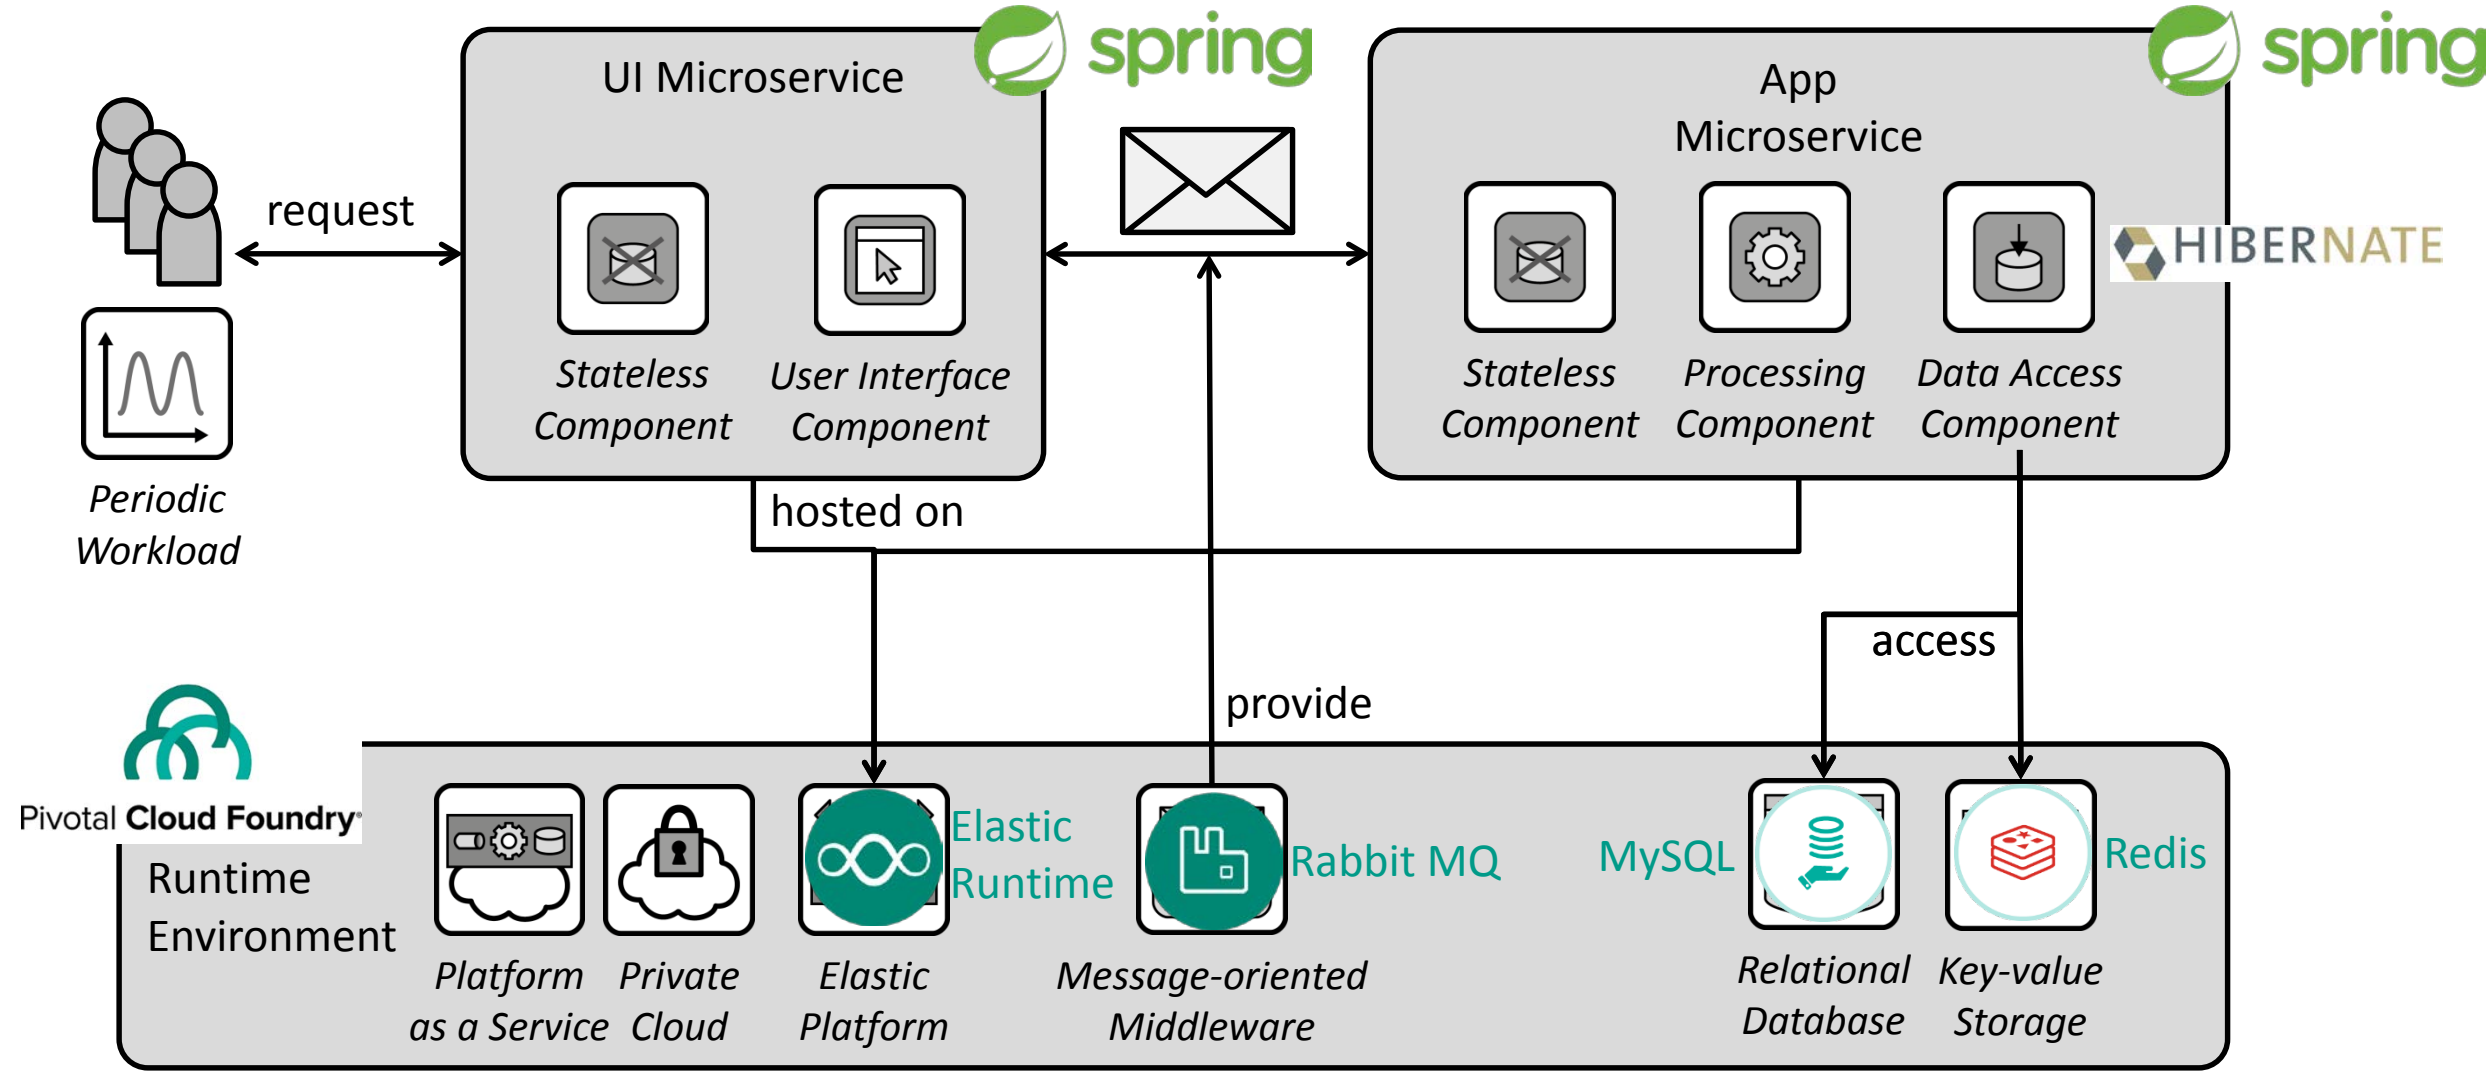
- ☐ Isolated State
- ☒ Distribution
- ☒ Elasticity
- ☒ Automated Management
- ☐ Loose Coupling



Mercedes Me Microservice Template



Mercedes Me Microservice Template



Lessons Learned

Moving from **Node-based Availability** to **Environment-based Availability** was the most **challenging factor**.

Cloud Computing is a significant enabler for **agile development**.

Soft factors (**Procurement, Licensing, Organizational Hierarchy**) pose a significant challenge!






Part 2: The Non-technical “Stuff”: Procurement, Licenses, Organizational Hierarchies

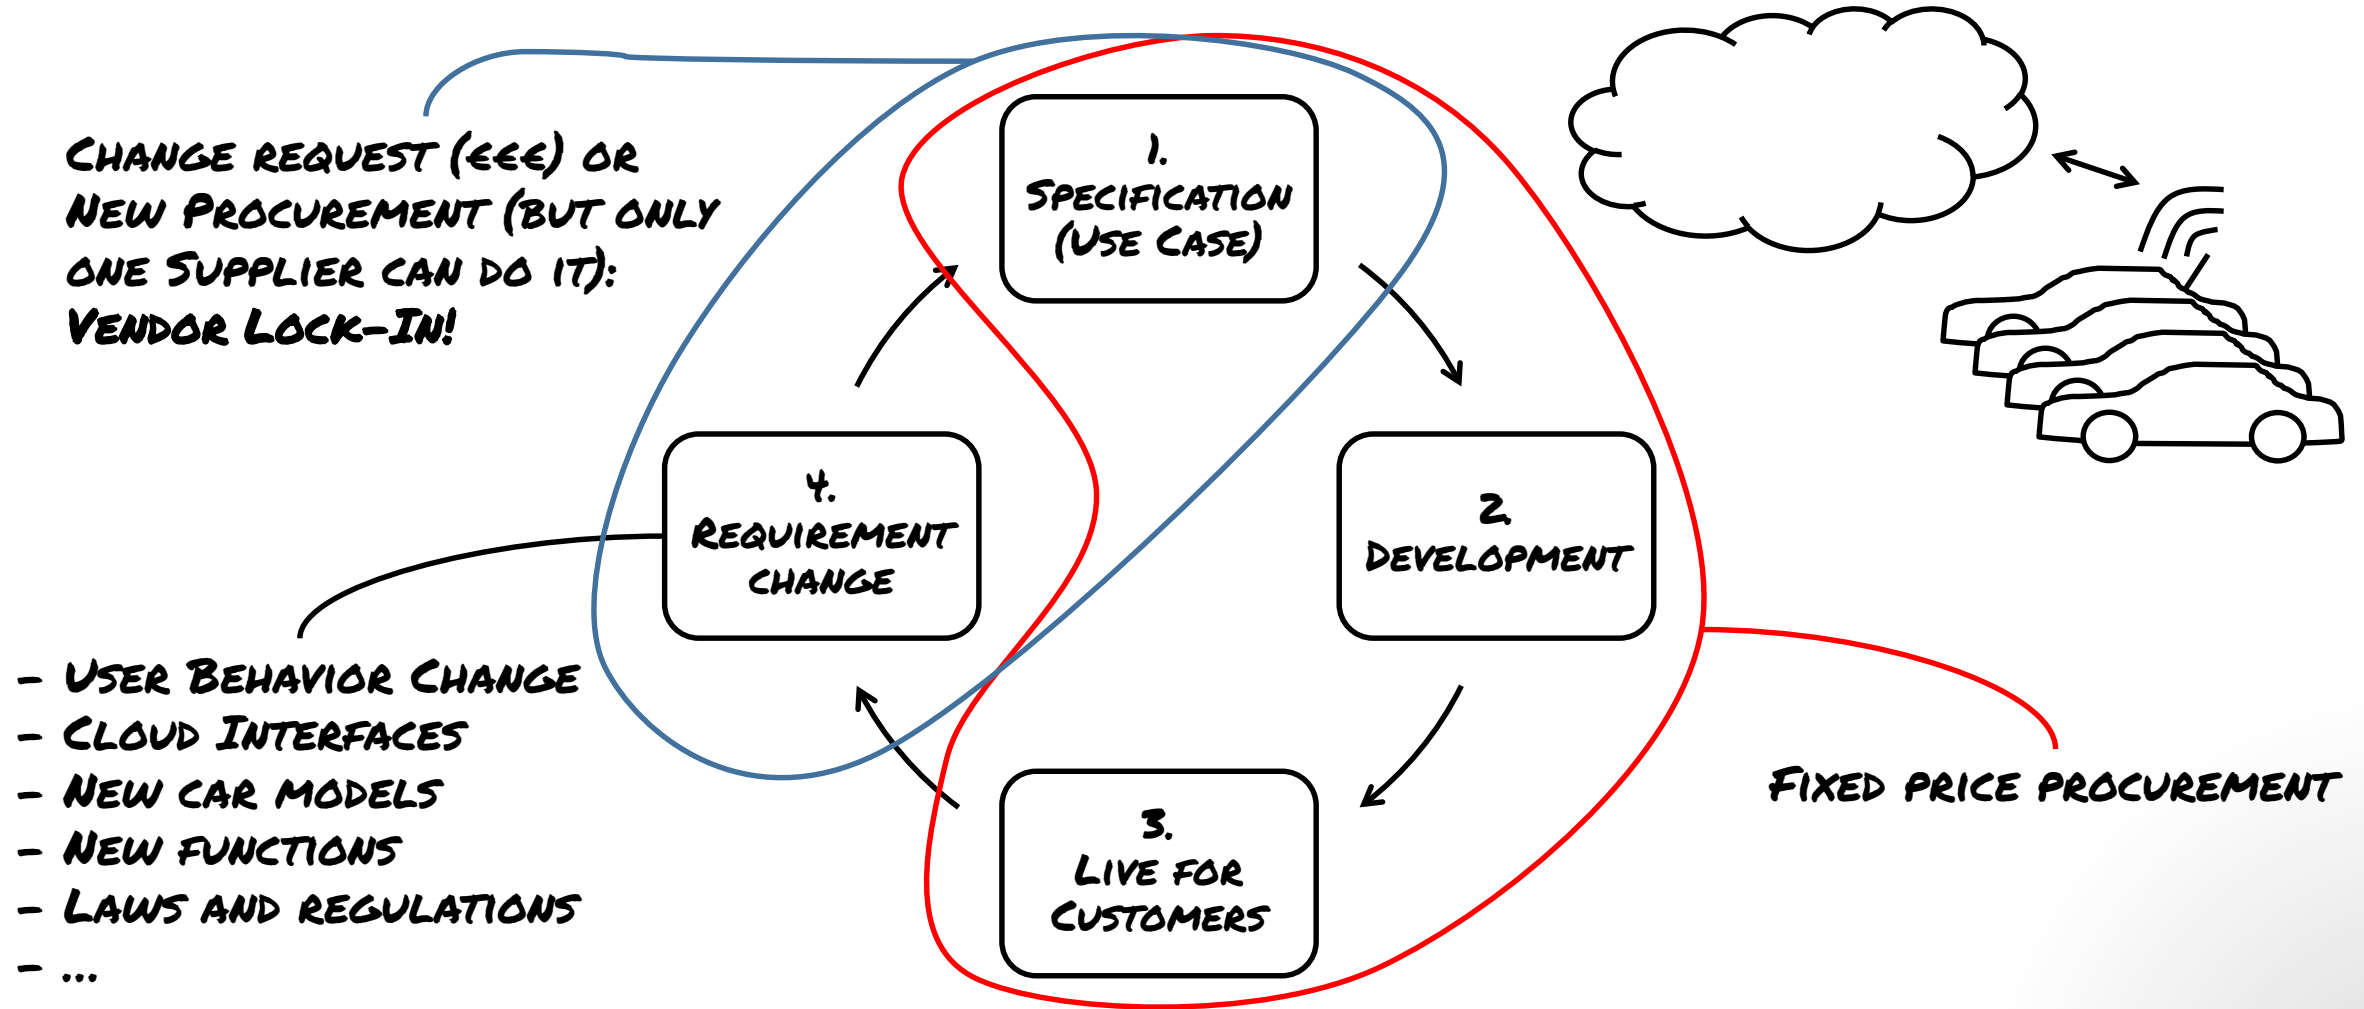
1. IMPACT OF CLOUD COMPUTING ON PROCUREMENT PROCESSES

PROCUREMENT OF A CAR PART (SIMPLIFIED VIEW)



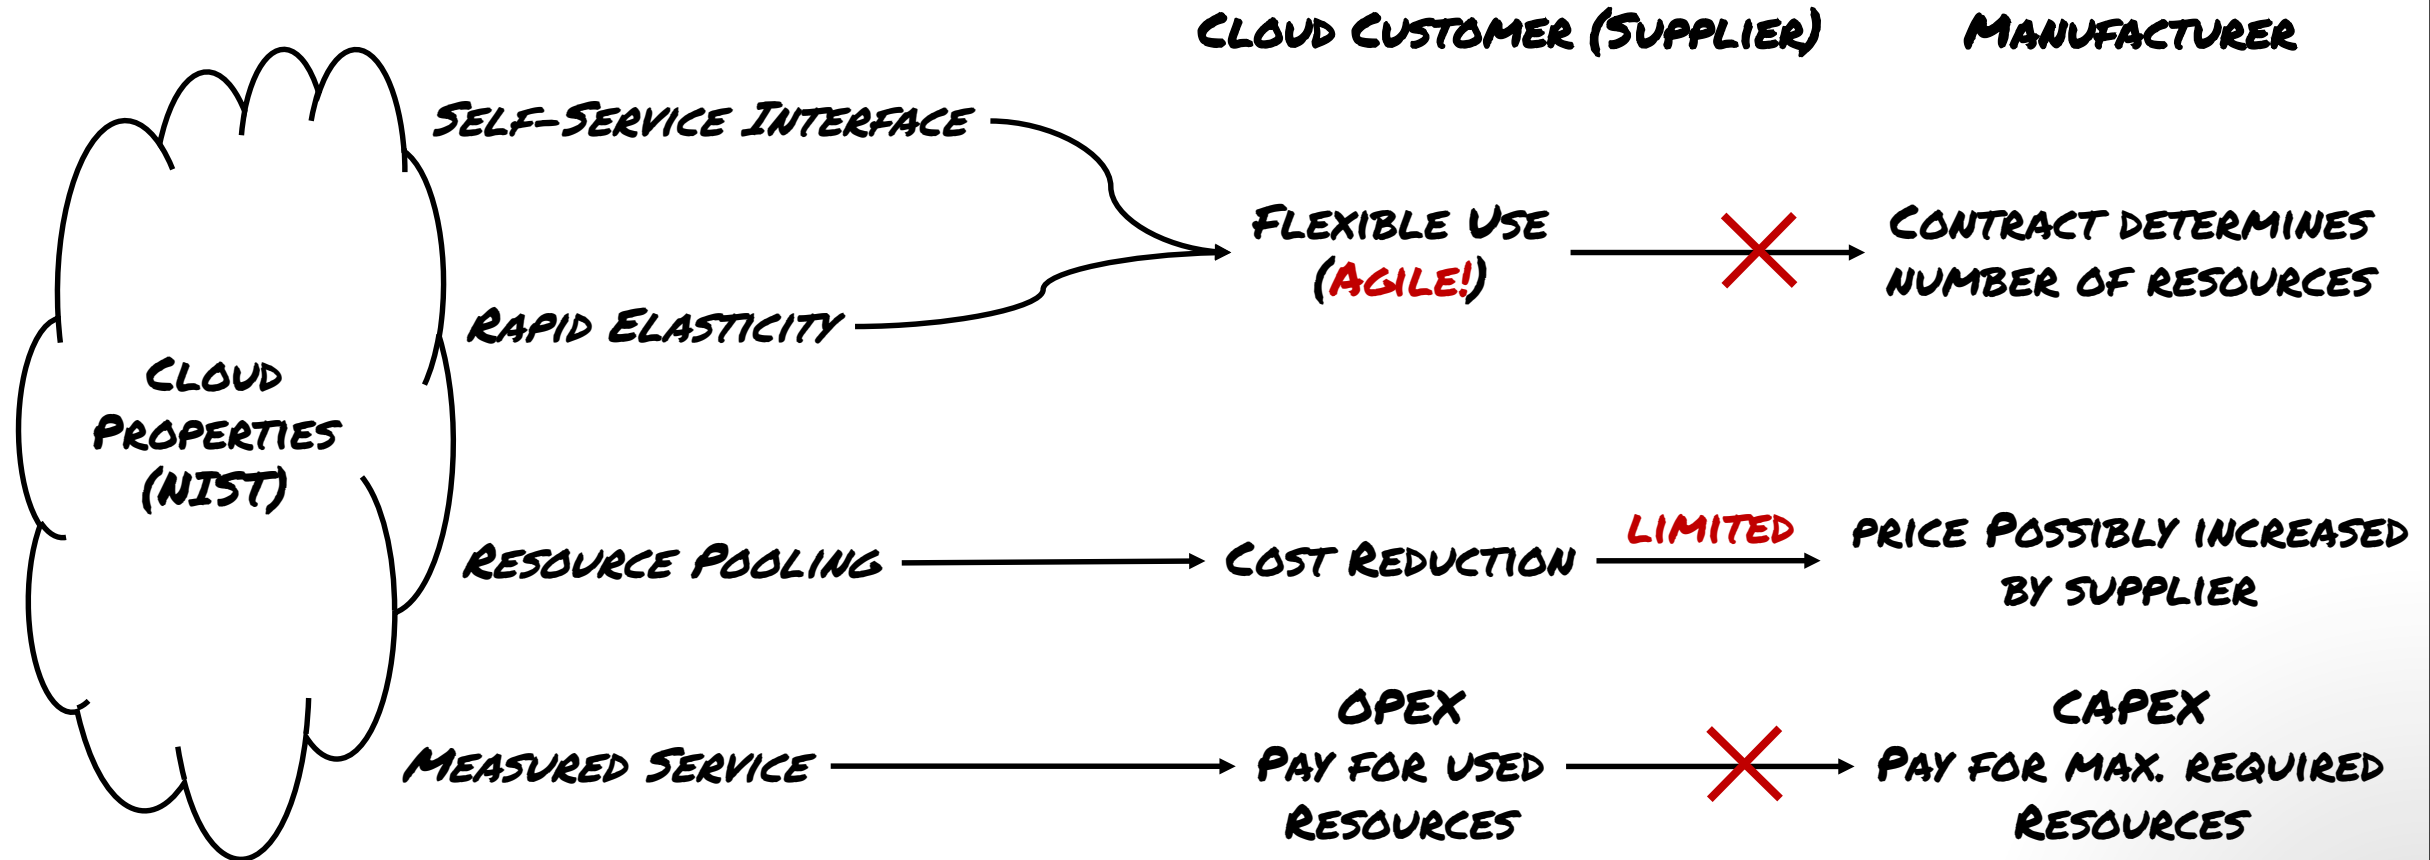
- SOFTWARE SCALES WITH THE NUMBER OF CARS AS IT IS INSTALLED IN EACH ONE OF THEM.
- IMPROVEMENTS / FIXES USUALLY INVOLVE A VISIT TO THE SHOP.

PROCUREMENT OF A CAR ~~PART~~ BACKEND SYSTEM



-> PROCUREMENT PROCESSES HAVE TO REFLECT THAT A BACKEND IS NOT A CAR PART!

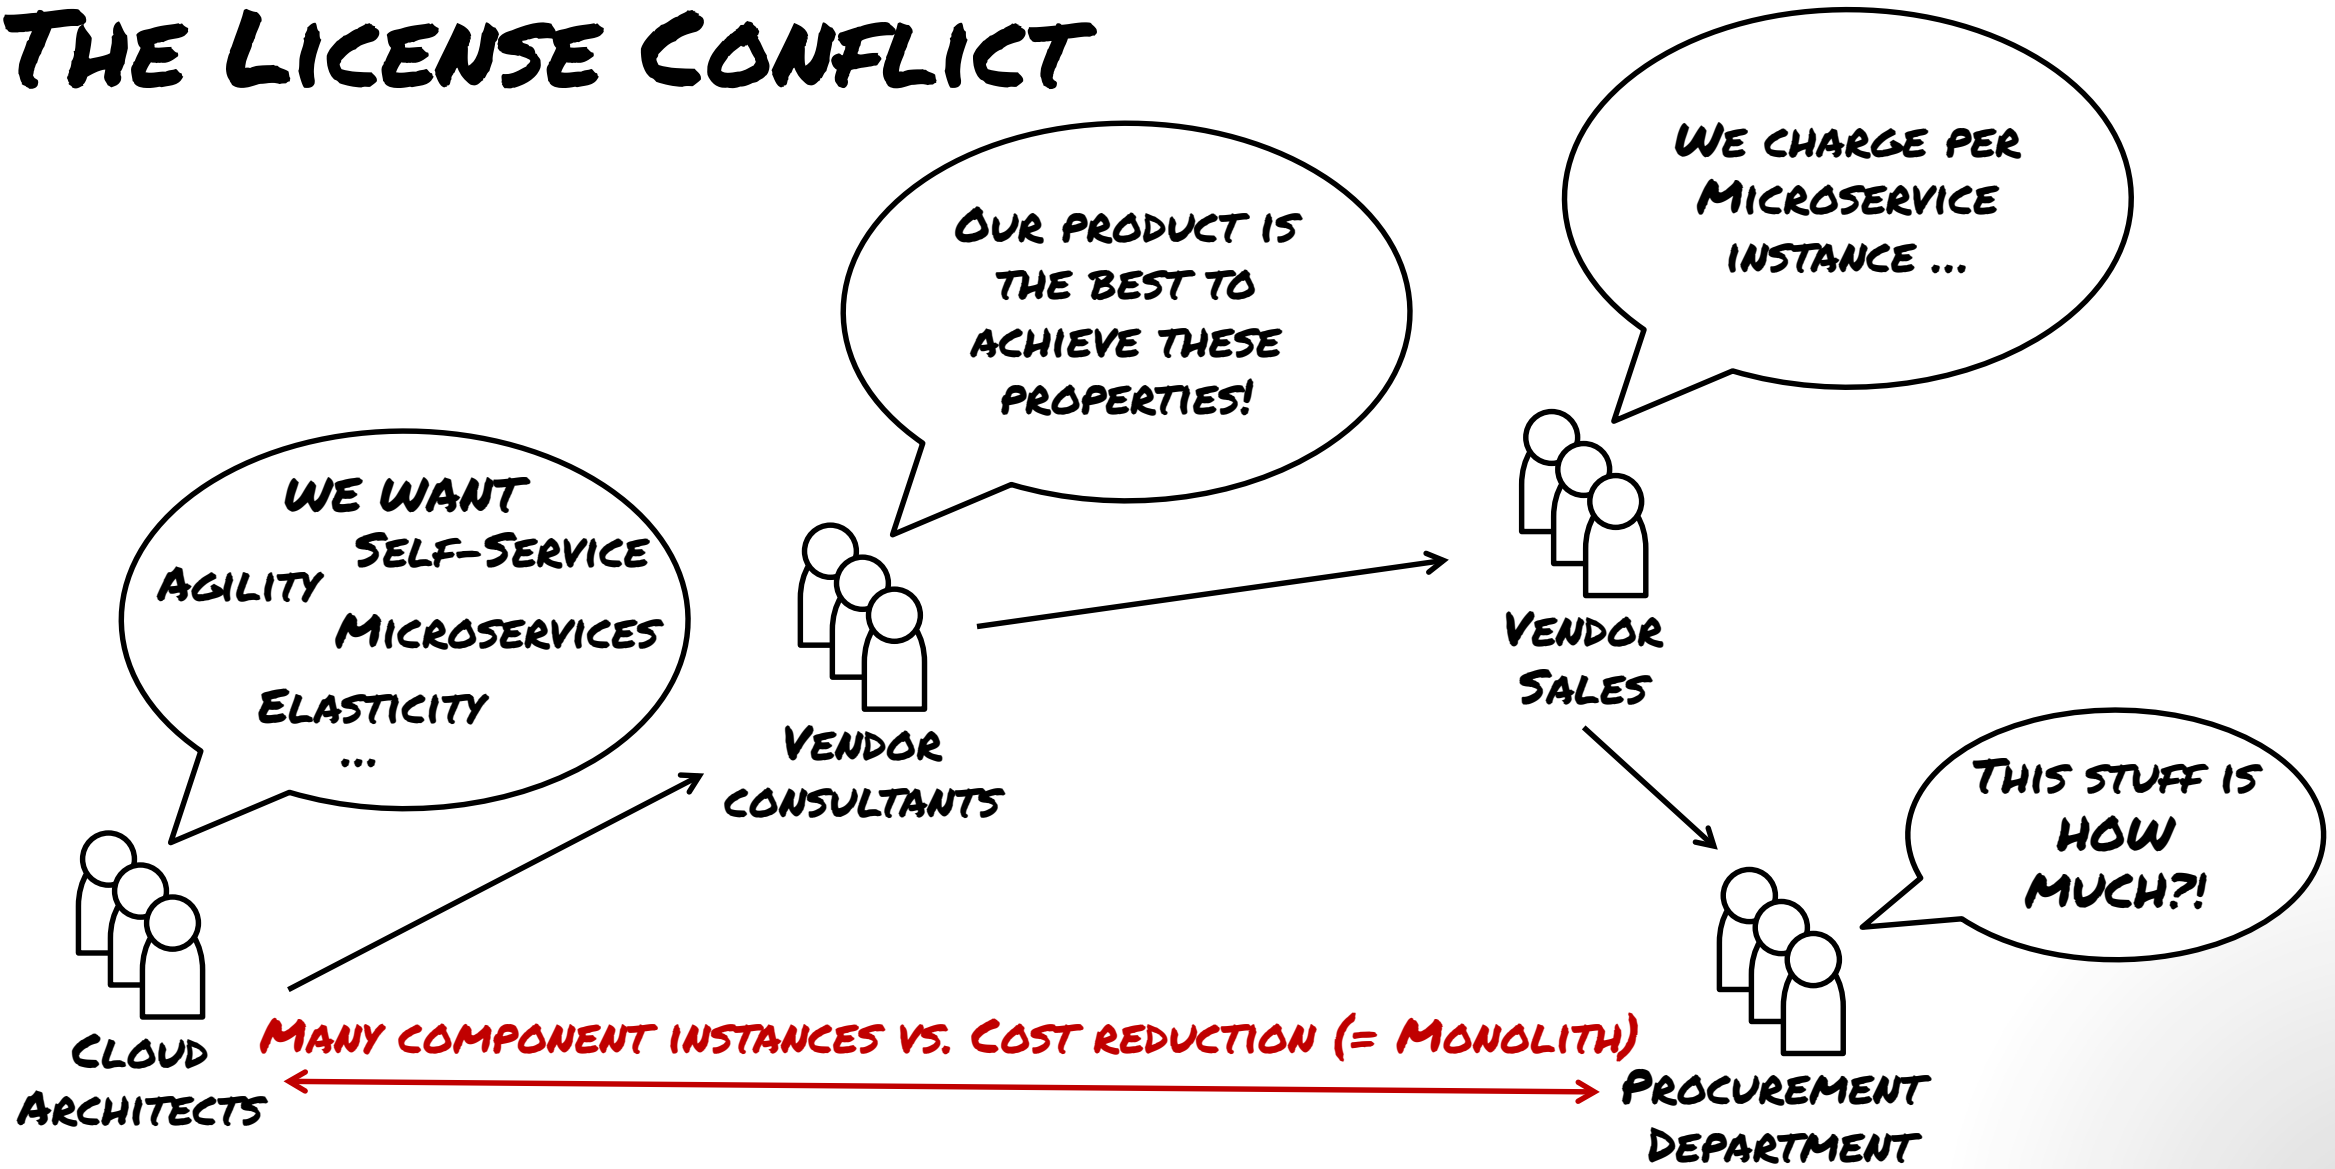
PROCUREMENT OF CLOUD-BASED BACKEND SYSTEMS



- > THE CLOUD CUSTOMER BENEFITS FROM CLOUD PROPERTIES. THIS SHOULD NOT BE THE SUPPLIER..
- > THE BENEFITS OF CLOUDS ARE HIGHLY DEPENDENT ON PROCUREMENT PROCESSES!

2. LICENSES IN THE CLOUD

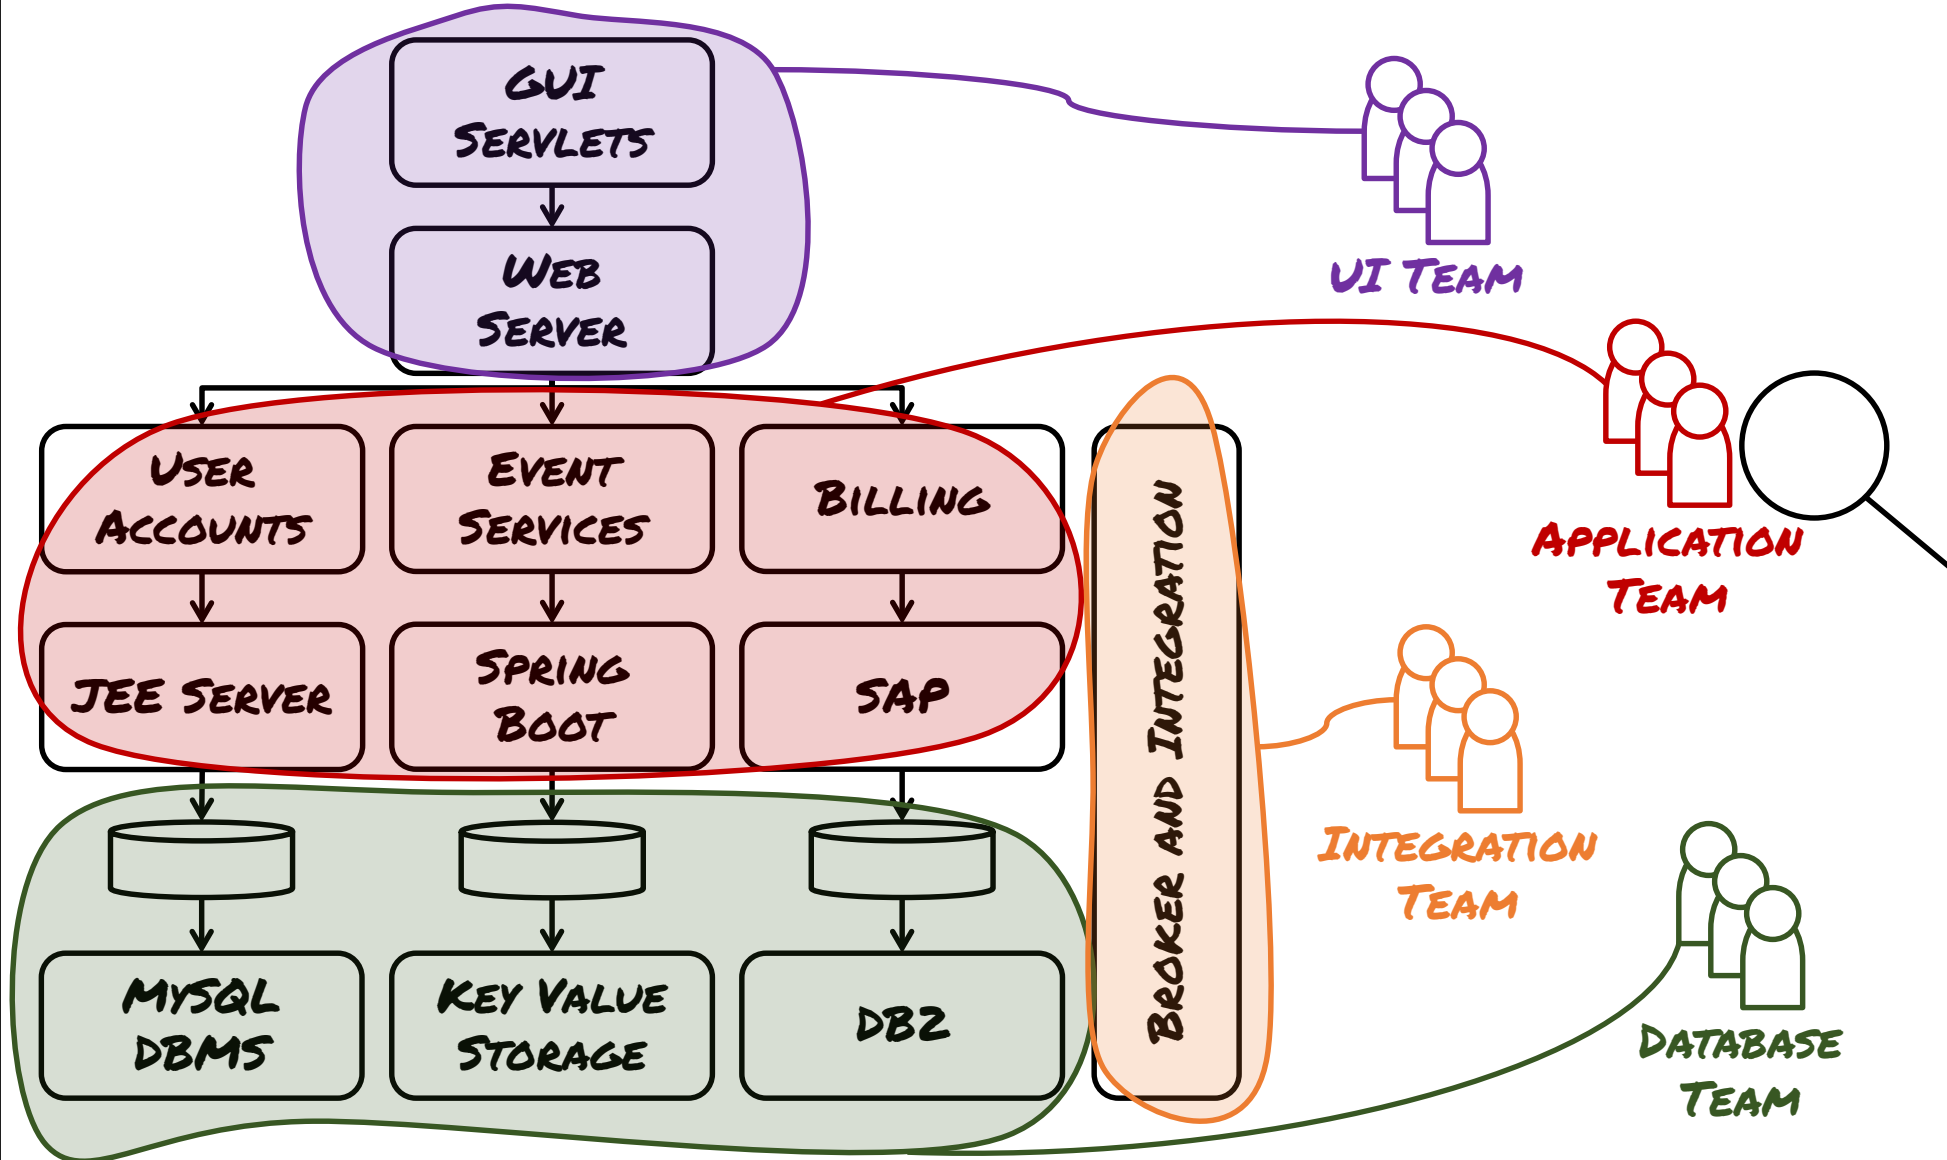
THE LICENSE CONFLICT



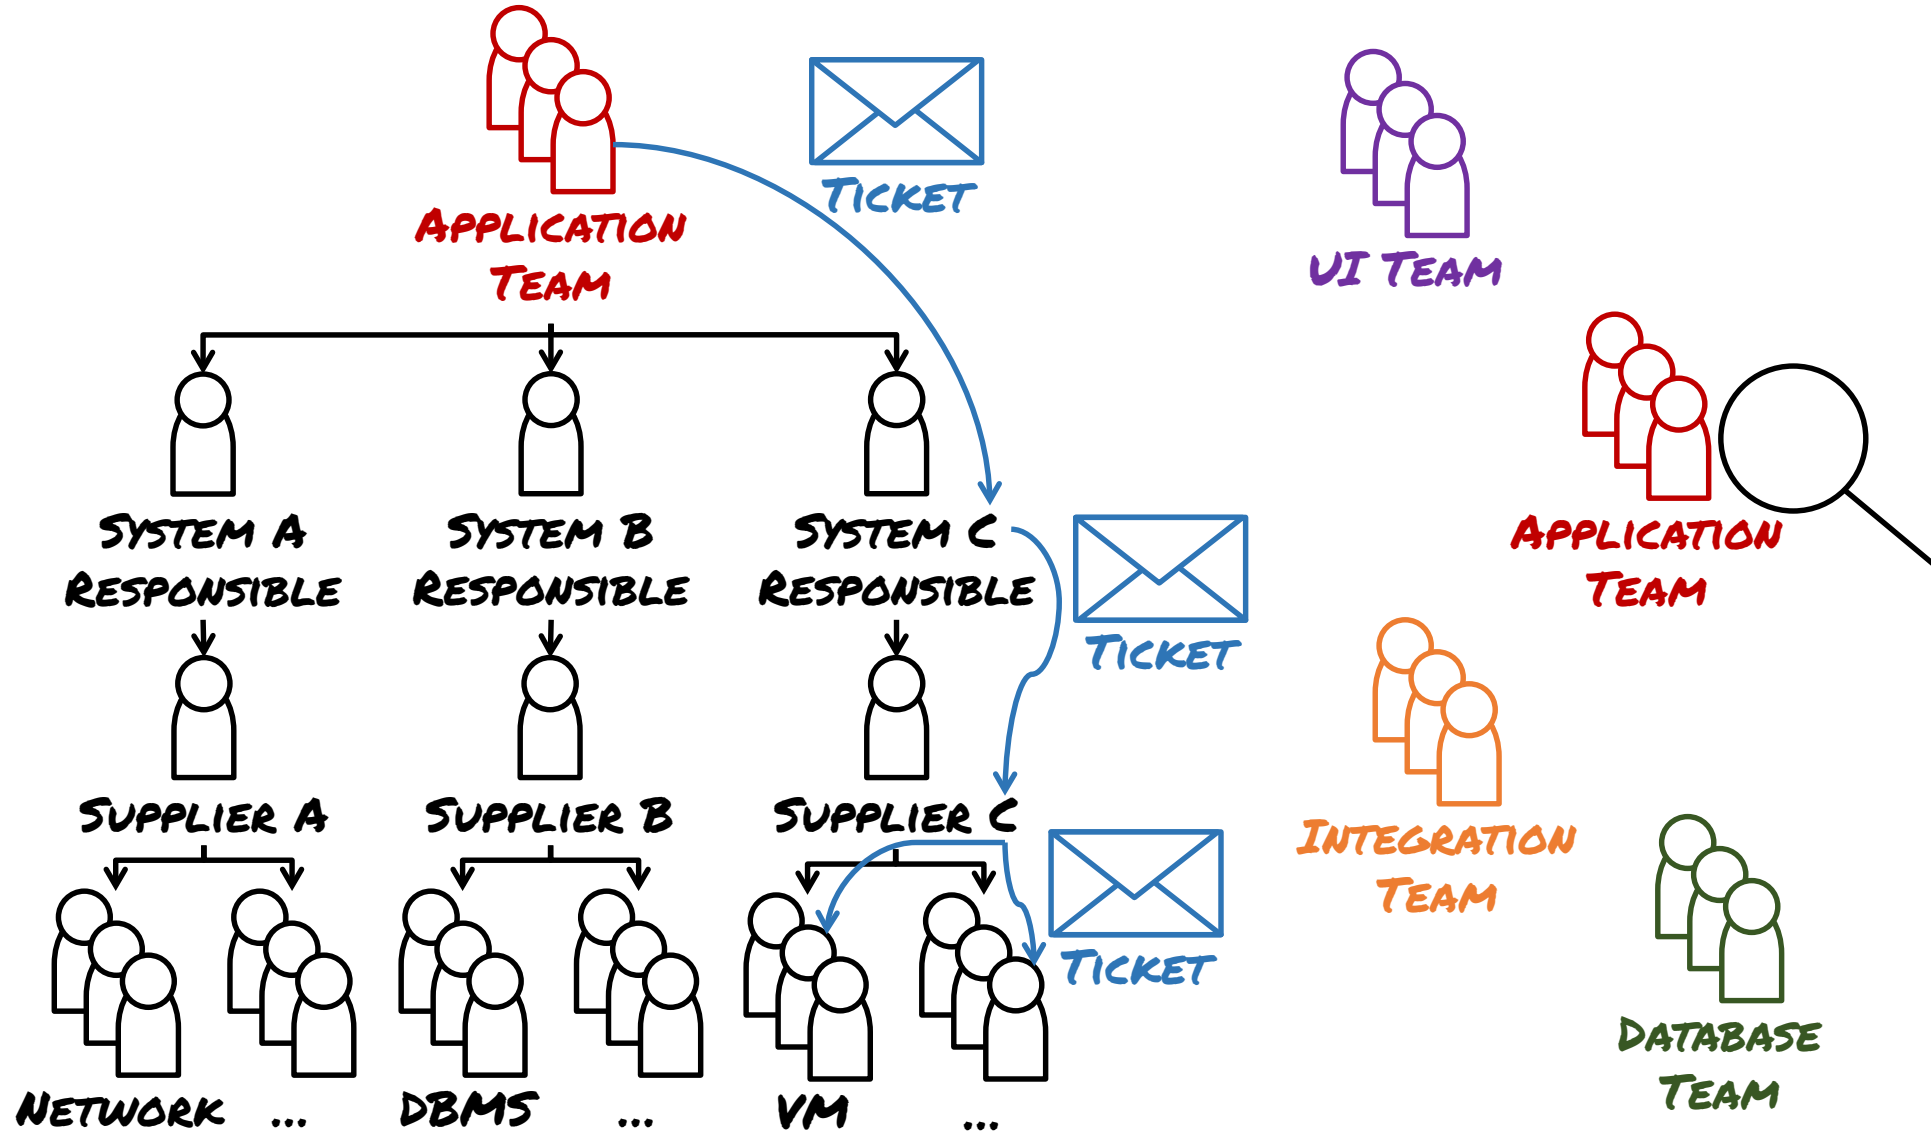
-> LICENSING MODELS OF MANY CLOUD PRODUCTS CONFLICT WITH OUR ARCHITECTURAL GOALS!

3. IMPACT OF CLOUD COMPUTING ON ORGANIZATION HIERARCHIES


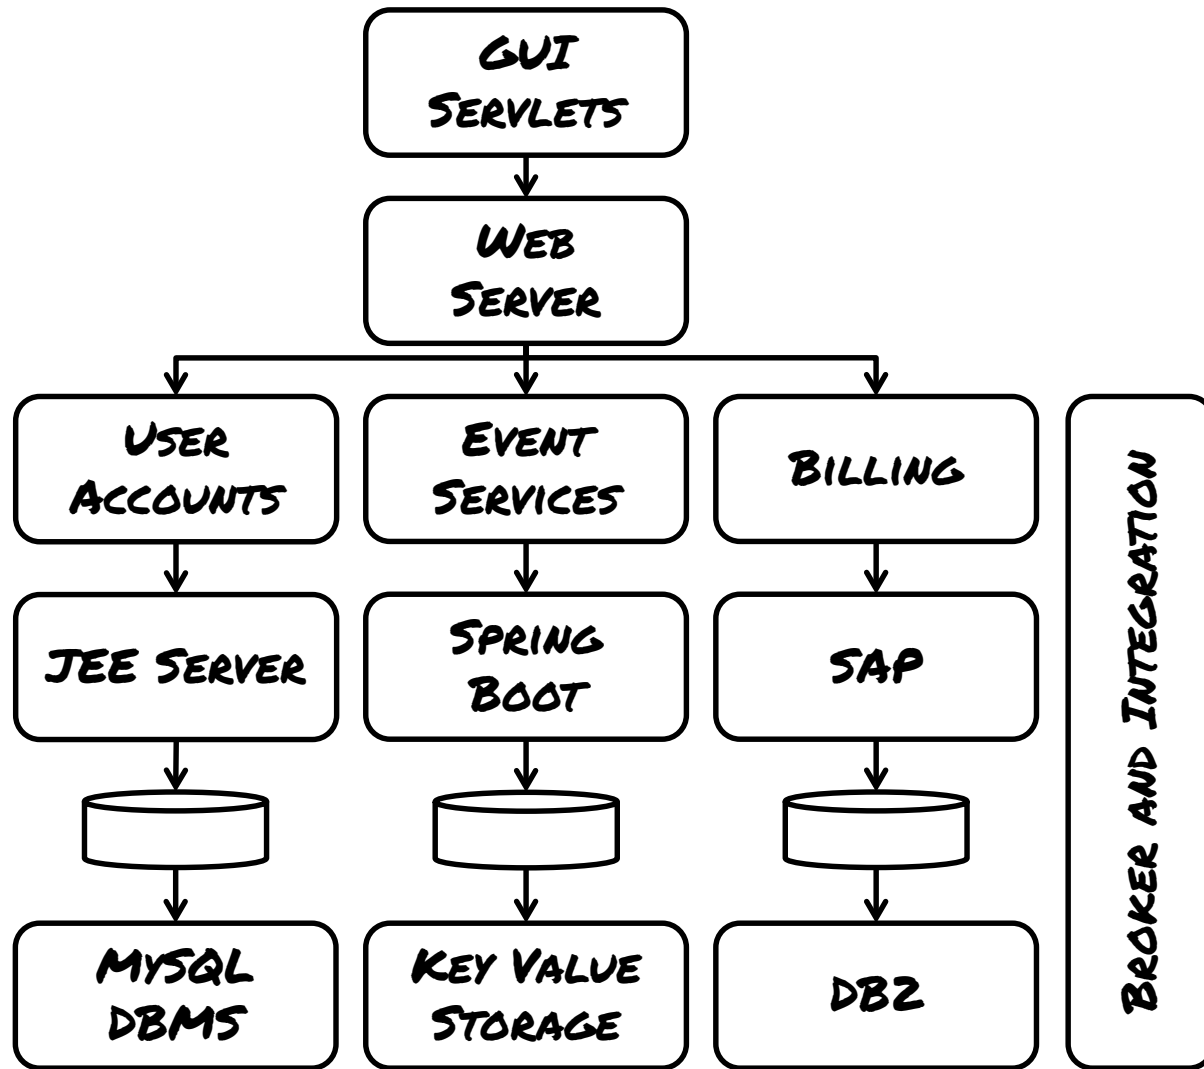
ORGANIZATION HIERARCHIES AND CLOUDS




ORGANIZATION HIERARCHIES AND CLOUDS



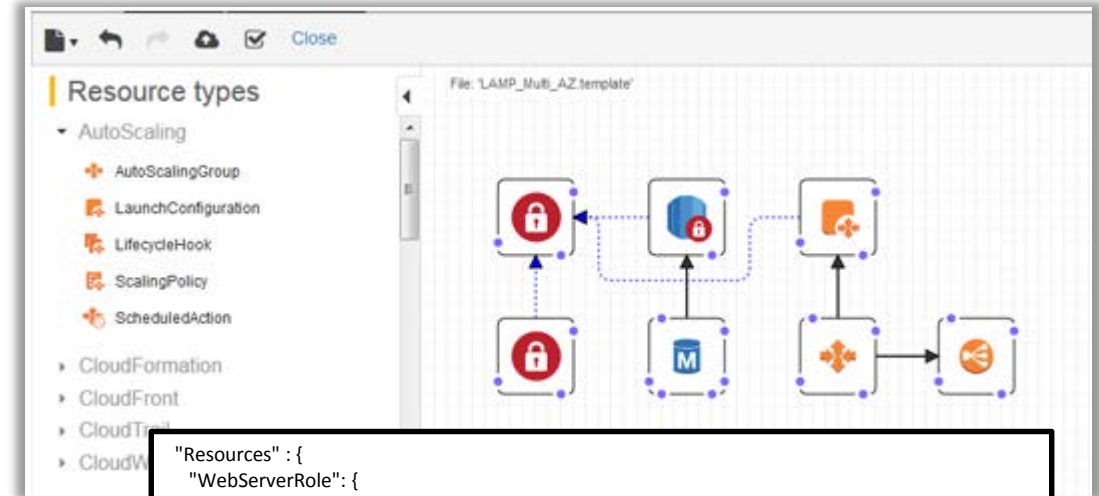
ORGANIZATION HIERARCHIES AND CLOUDS



SIDE NOTE: CLOUDS THRIVE ON AUTOMATION!


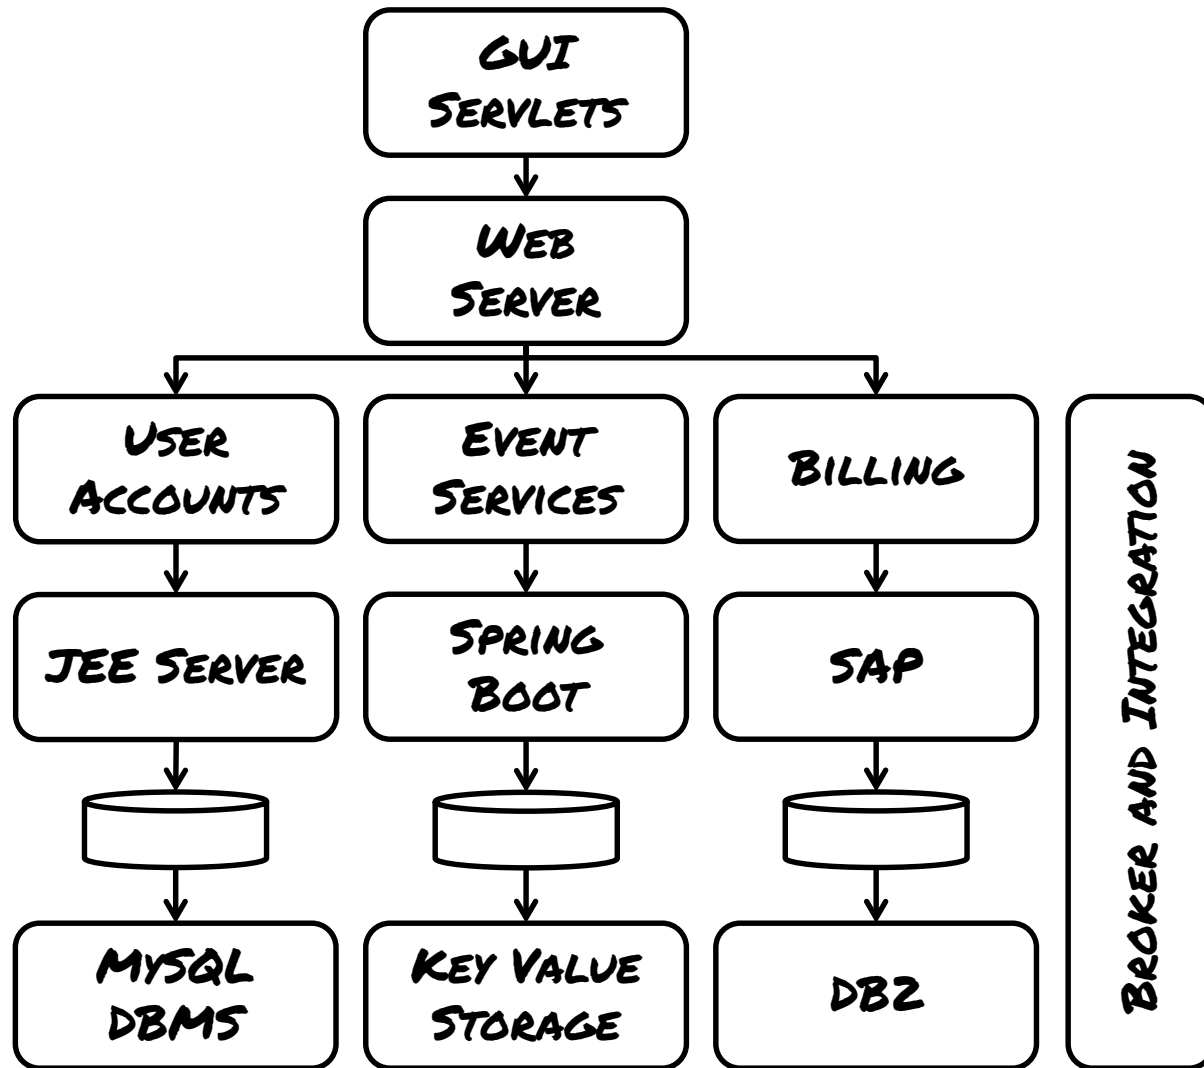


AMAZON CLOUD FORMATION



```
"Resources": {
  "WebServerRole": {
    "Type": "AWS::IAM::Role",
    "Properties": {
      "AssumeRolePolicyDocument": {
        "Statement": [{
          "Effect": "Allow",
          "Principal": { "Service": [{ "Fn::FindInMap": ["Region2Principal", {"Ref": "AWS::Region"}],
            "EC2Principal"] }},
          "Action": [ "sts:AssumeRole" ]
        }
      ],
      "Path": "/"
    }
  },
  "WebServerRolePolicy": {
    "Type": "AWS::IAM::Policy",
    "Properties": {
      "PolicyName": "WebServerRole",
```


SIDE NOTE: CLOUDS THRIVE ON AUTOMATION!



DOCKER COMPOSE




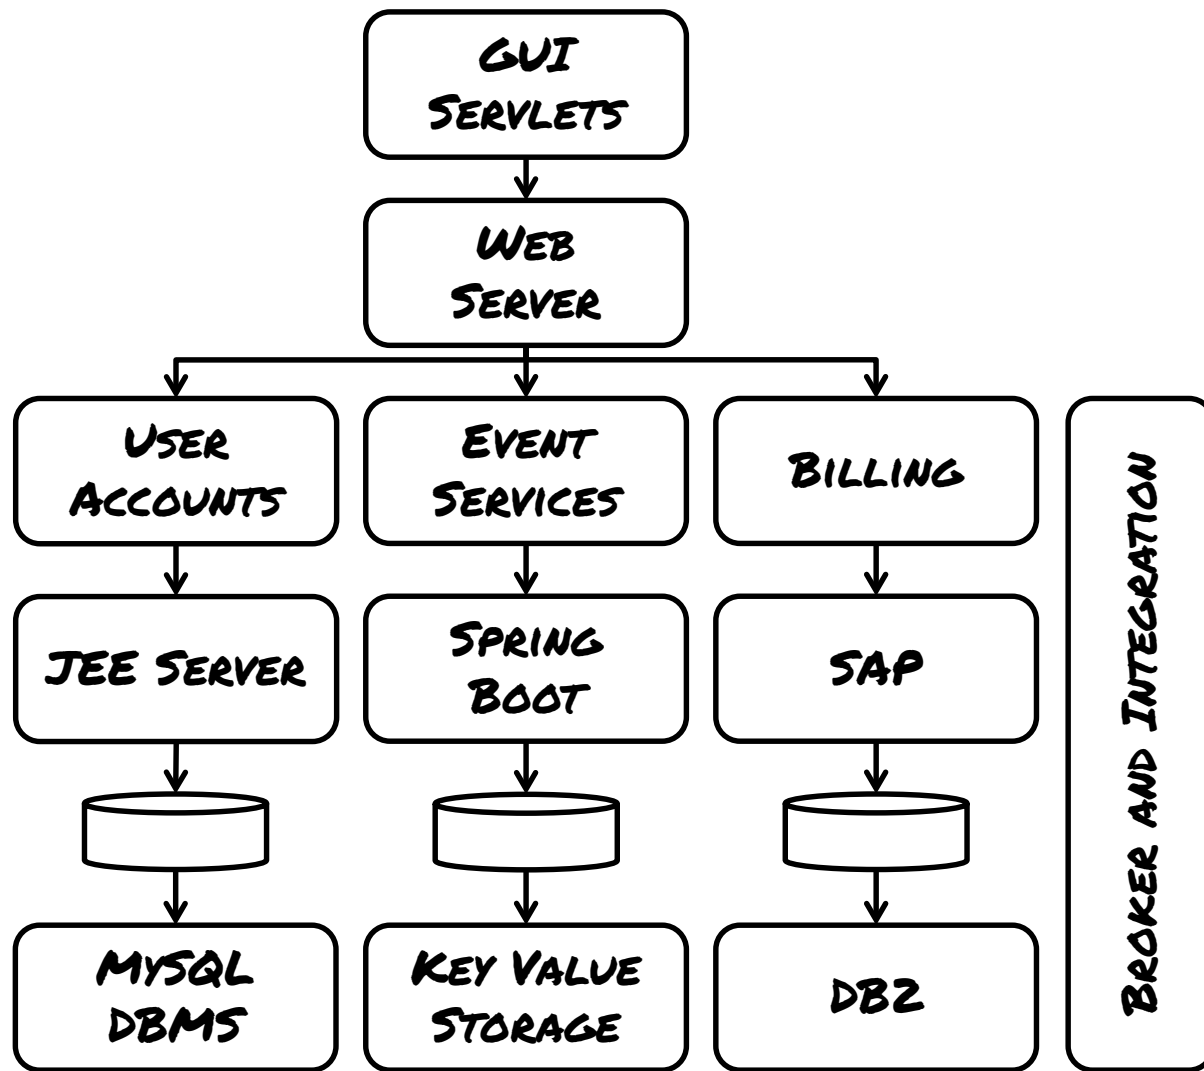
CHEF™



ANSIBLE

->INFRASTRUCTURE AS CODE IS MANDATORY TO
BENEFIT FROM THE CLOUD PROPERTIES!

OK! LET'S AUTOMATE OUR IT STACK!




OK! LET'S AUTOMATE OUR IT STACK!

BUT WAIT...

- WHO CREATES (WHICH PART OF) THE MODEL?
- WHICH SYSTEM INTERPRETS AND EXECUTES THE MODEL?
- WHO CONTROLS THE AUTOMATION SYSTEM?
- WHAT ABOUT EXISTING TICKET SYSTEMS?!
- WHAT ABOUT EXISTING MANUAL TASKS?

- > AUTOMATION REQUIRES CONTROL OVER THE CLOUD TO BE CENTRALIZED - NOT DELEGATED!
- > EMPOWERMENT OF DEVOPS TEAMS IS NEEDED TO USE CLOUDS EFFICIENTLY!
- > EXISTING TEAMS MAY BE AFRAID TO LOOSE CONTROL.



LESSONS LEARNED

REVISE PROCUREMENT PROCESSES, BECAUSE...

... SUPPLIERS USING A CLOUD FOR YOU MAY **CREATE A VENDOR LOCK-IN!**

... MANY **BENEFITS OF THE CLOUD PROPERTIES ARE LOST** IF A SUPPLIER USES A CLOUD FOR YOU!

DEMAND CLOUD-COMPATIBLE LICENSES, BECAUSE...

... **COSTS PER INSTANCE CONFLICTS WITH ARCHITECTURAL GOALS!**

REVISE YOUR ORGANIZATIONAL HIERARCHIES, BECAUSE...


... **CLOUDS THRIVE ON AUTOMATION** AND REQUIRE FEWER DELEGATION OF MANUAL TASKS!

... **SELF-SERVICE INTERFACES ARE MORE AGILE** THAN TICKETING SYSTEMS!


Summary

Part 1: Cloud Computing Patterns @ Mercedes Me


Microservice Template based on Cloud Computing Patterns and Pivotal Cloud Foundry




*Static
Workload*




*Periodic
Workload*




*Once-in-a-lifetime
Workload*




*Unpredictable
Workload*




*Continuously
Changing
Workload*



*Strict
Consistency*




*Eventual
Consistency*




*Data
Abstraction*




*Node-based
Availability*




*Environment-based
Availability*



*User Interface
Component*



*Processing
Component*



Part 1: Cloud Computing Patterns @ Mercedes Me

Microservice Template based on Cloud Computing Patterns and Pivotal Cloud Foundry

Part 2: The Non-technical “Stuff”...

**PROCUREMENT PROCESSES HAVE TO BE
ADJUSTED FOR CLOUD COMPUTING**

- **HOW CAN WE BUY ENVIRONMENTS FOR
AGILE CLOUD DEVELOPMENT?**
- **HOW CAN WE BENEFIT FROM CLOUD
PROPERTIES – NOT OUR SUPPLIERS?**

**ORGANIZATION HIERARCHIES HAVE TO BE
ADJUSTED FOR CLOUD COMPUTING**

- **HOW CAN WE BENEFIT FROM CLOUD
AUTOMATION?**
- **HOW CAN WE ORGANIZE WORK WITHOUT
TICKETING SYSTEMS FOR MANUAL TASKS?**

Part 3: Discussion during OOP ☺

I'm here all week! Contact me: fehling.c@gmail.com +49 170 58 35 456 @ccpatterns

<http://www.cloudcomputingpatterns.org>






Design Steps for Cloud Applications using Patterns



or

to see a Cloud Application Architecture you should go out and have a...



Design Steps for Cloud Applications




 **IDEAL** 
Cloud-Native Application

- ✓ Isolated State
- ✓ Distribution
- ✓ Elasticity
- ✓ Automated Management
- ✓ Loose Coupling




 **How to distribute Application Functionality?**



Distributed Application

A cloud application **divides provided functionality** among multiple application components that can be **scaled out independently**.

Layer-based Decomposition



Components reside on separate functional layers

Often: user interface, processing, storage


Access is only allowed to **same layer and the layer below**

→ Dependencies between layers and interfaces are controlled



Cloud-Native Application


- ☒ Isolated State
- ☒ Distribution
- ☐ Elasticity
- ☐ Automated Management
- ☐ Loose Coupling



Distributed Application

A cloud application **divides provided functionality** among multiple application components that can be **scaled out independently**.

Process-based Decomposition



Business process model determines decomposition

Activities: tasks executed in a specific order (**control flow**)


Data elements: information handled by activities (**data flow**)

Functional application components (**services**) are accessed by process



Cloud-Native Application


- ☒ Isolated State
- ☒ Distribution
- ☐ Elasticity
- ☐ Automated Management
- ☐ Loose Coupling



Distributed Application

A cloud application **divides provided functionality** among multiple application components that can be **scaled out independently**.

Pipes-and-Filters-based Decomposition



Decomposition based on the data processing function

Filter: application component processing data

Pipe: connection between filters (commonly messaging)




Cloud-Native Application

- ☒ Isolated State
- ☒ Distribution
- ☐ Elasticity
- ☐ Automated Management
- ☐ Loose Coupling

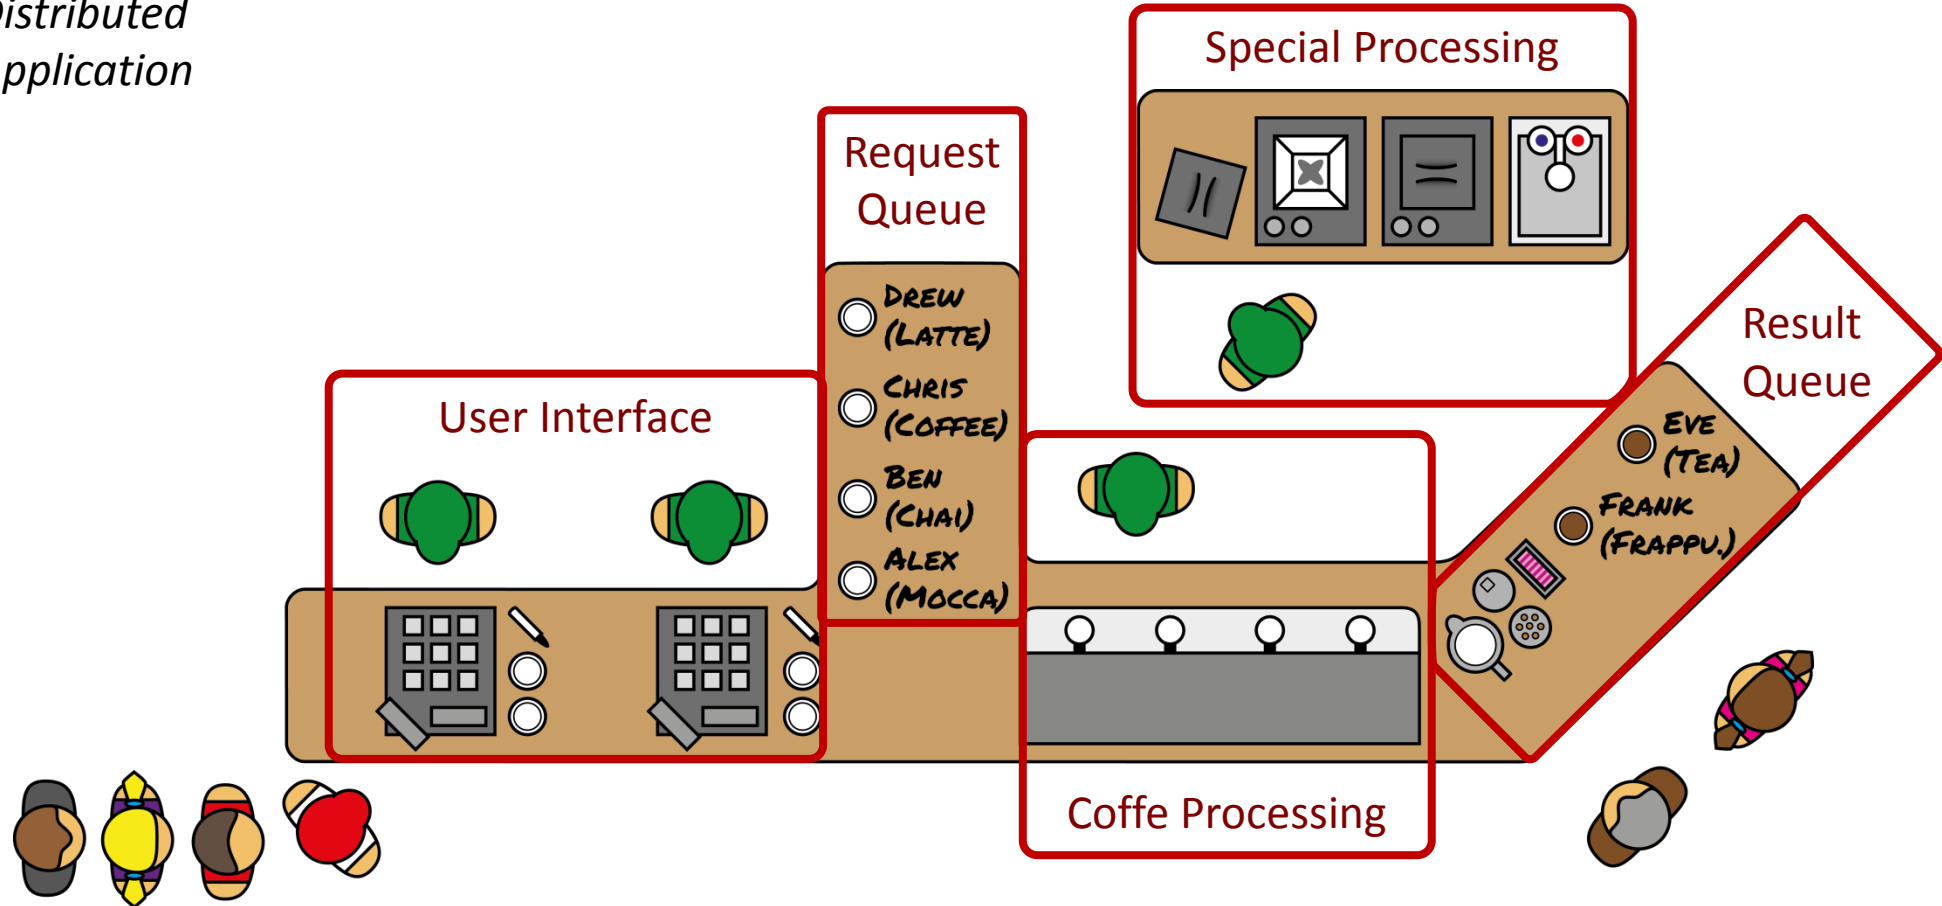


Coffe Shop – Decomposition of Functions

Identify functional components.




*Distributed
Application*






 **What workload do components experience?**




Static Workload

IT resources with an **equal utilization over time** experience static workload.




Periodic Workload

IT resources with a **peaking utilization at reoccurring time intervals** experience periodic workload.




Once-in-a-Lifetime Workload

IT resources with an equal utilization over time disturbed by a **strong peak occurring only once** experience once-in-a-lifetime workload.



Unpredictable Workload

IT resources with a **random and unforeseeable utilization** over time experience unpredictable workload.



Continuously Changing Workload

IT resources with a **utilization that grows or shrinks constantly** over time experience continuously changing workload.

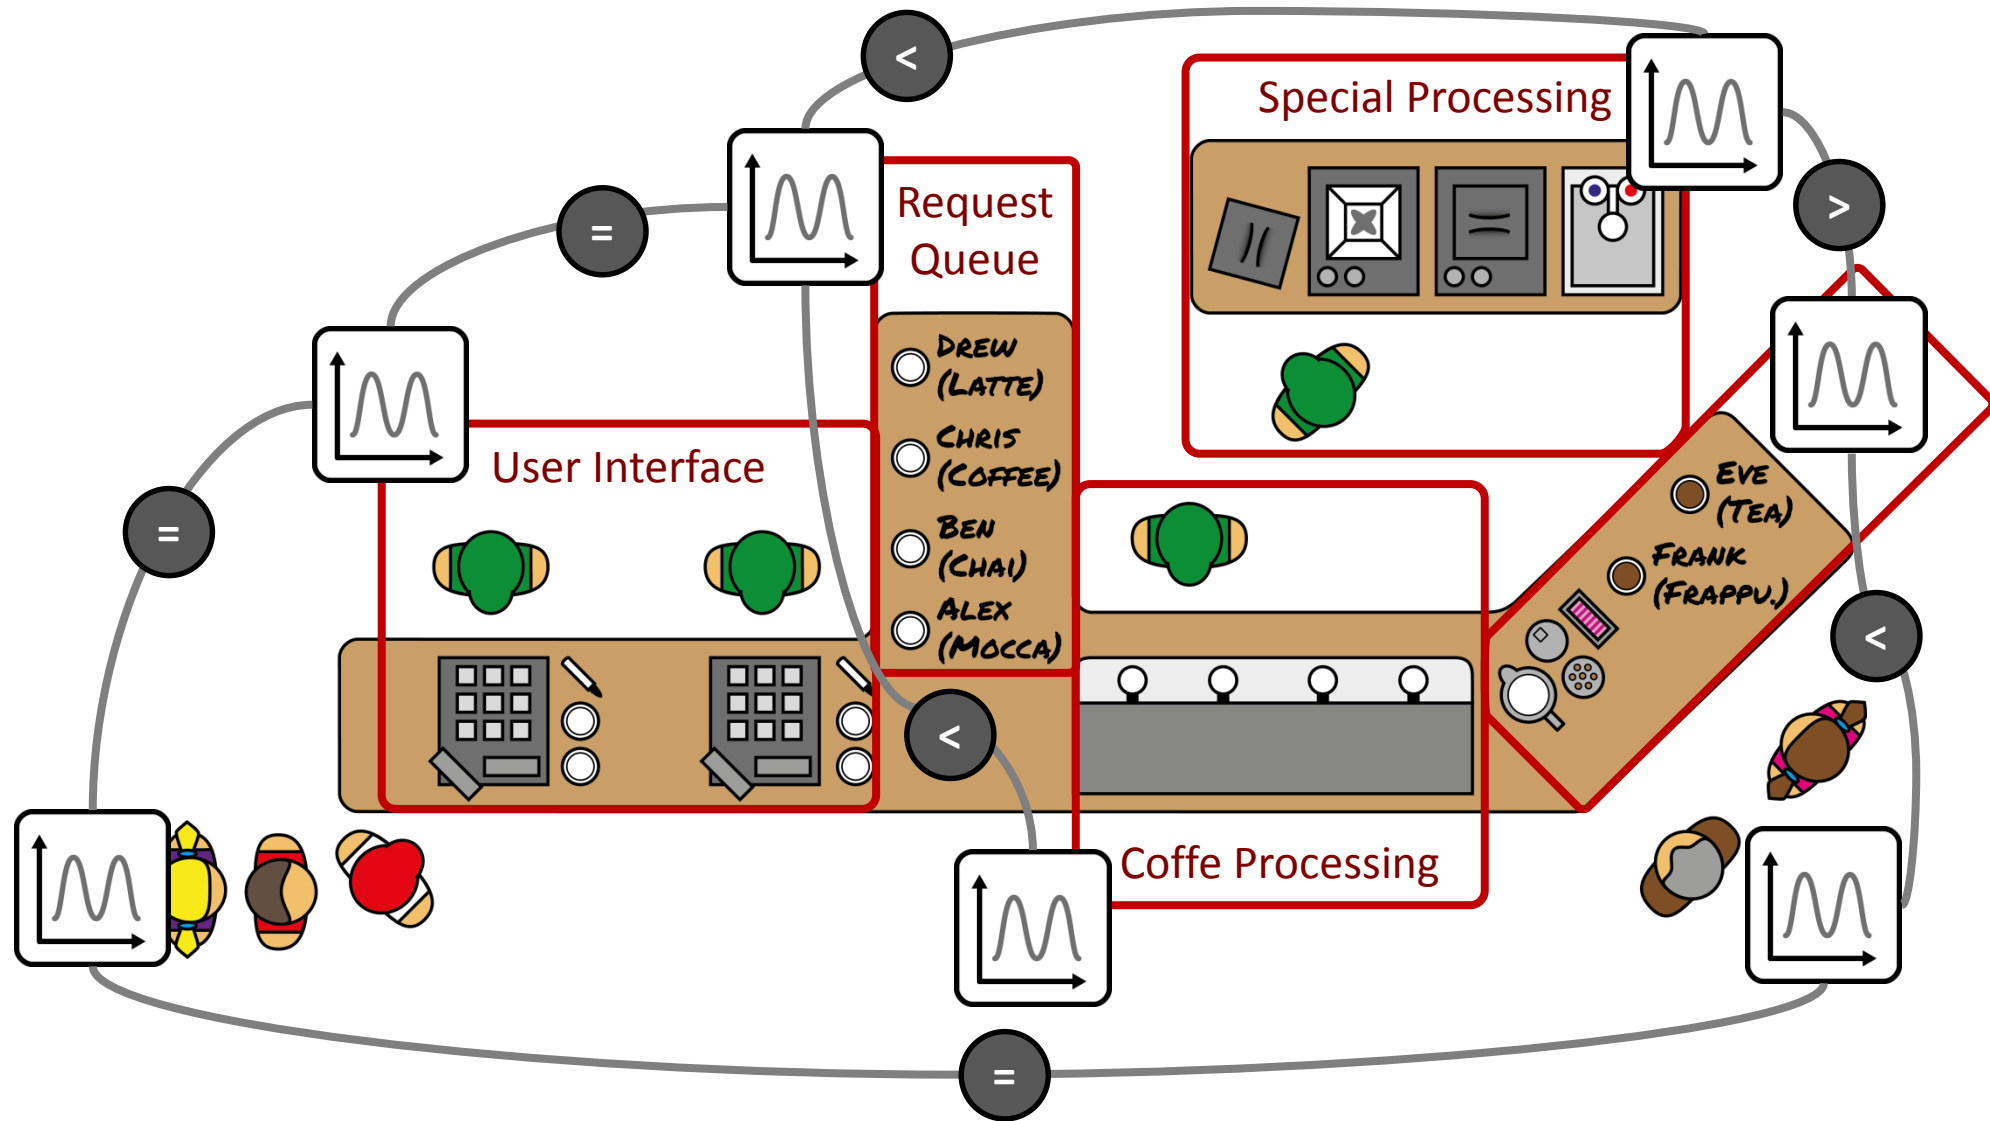


Cloud-Native Application

- ☐ Isolated State
- ☐ Distribution
- ☒ Elasticity
- ☒ Automated Management
- ☐ Loose Coupling


Coffe Shop – Workloads

Identify and compare workload generated by user groups at different components.





 **Where does the application handle state?**




Stateful Component

Multiple instances of a scaled-out application component **synchronize their internal state** to provide a unified behavior.




Stateless Component

State is handled **external of application components** to ease their scaling-out and to make the application more tolerant to component failures.




Strict Consistency

Data is stored at different locations to improve response time and to avoid data loss in case of failures while **consistency of replicas is ensured at all times**.



Eventual Consistency

Performance and availability of data in case of network partitioning are enabled by ensuring **data consistency eventually and not at all times**.



Data Abtractor

Data is **abstracted to inherently support eventually consistent data** storage through the use of abstractions and approximations.

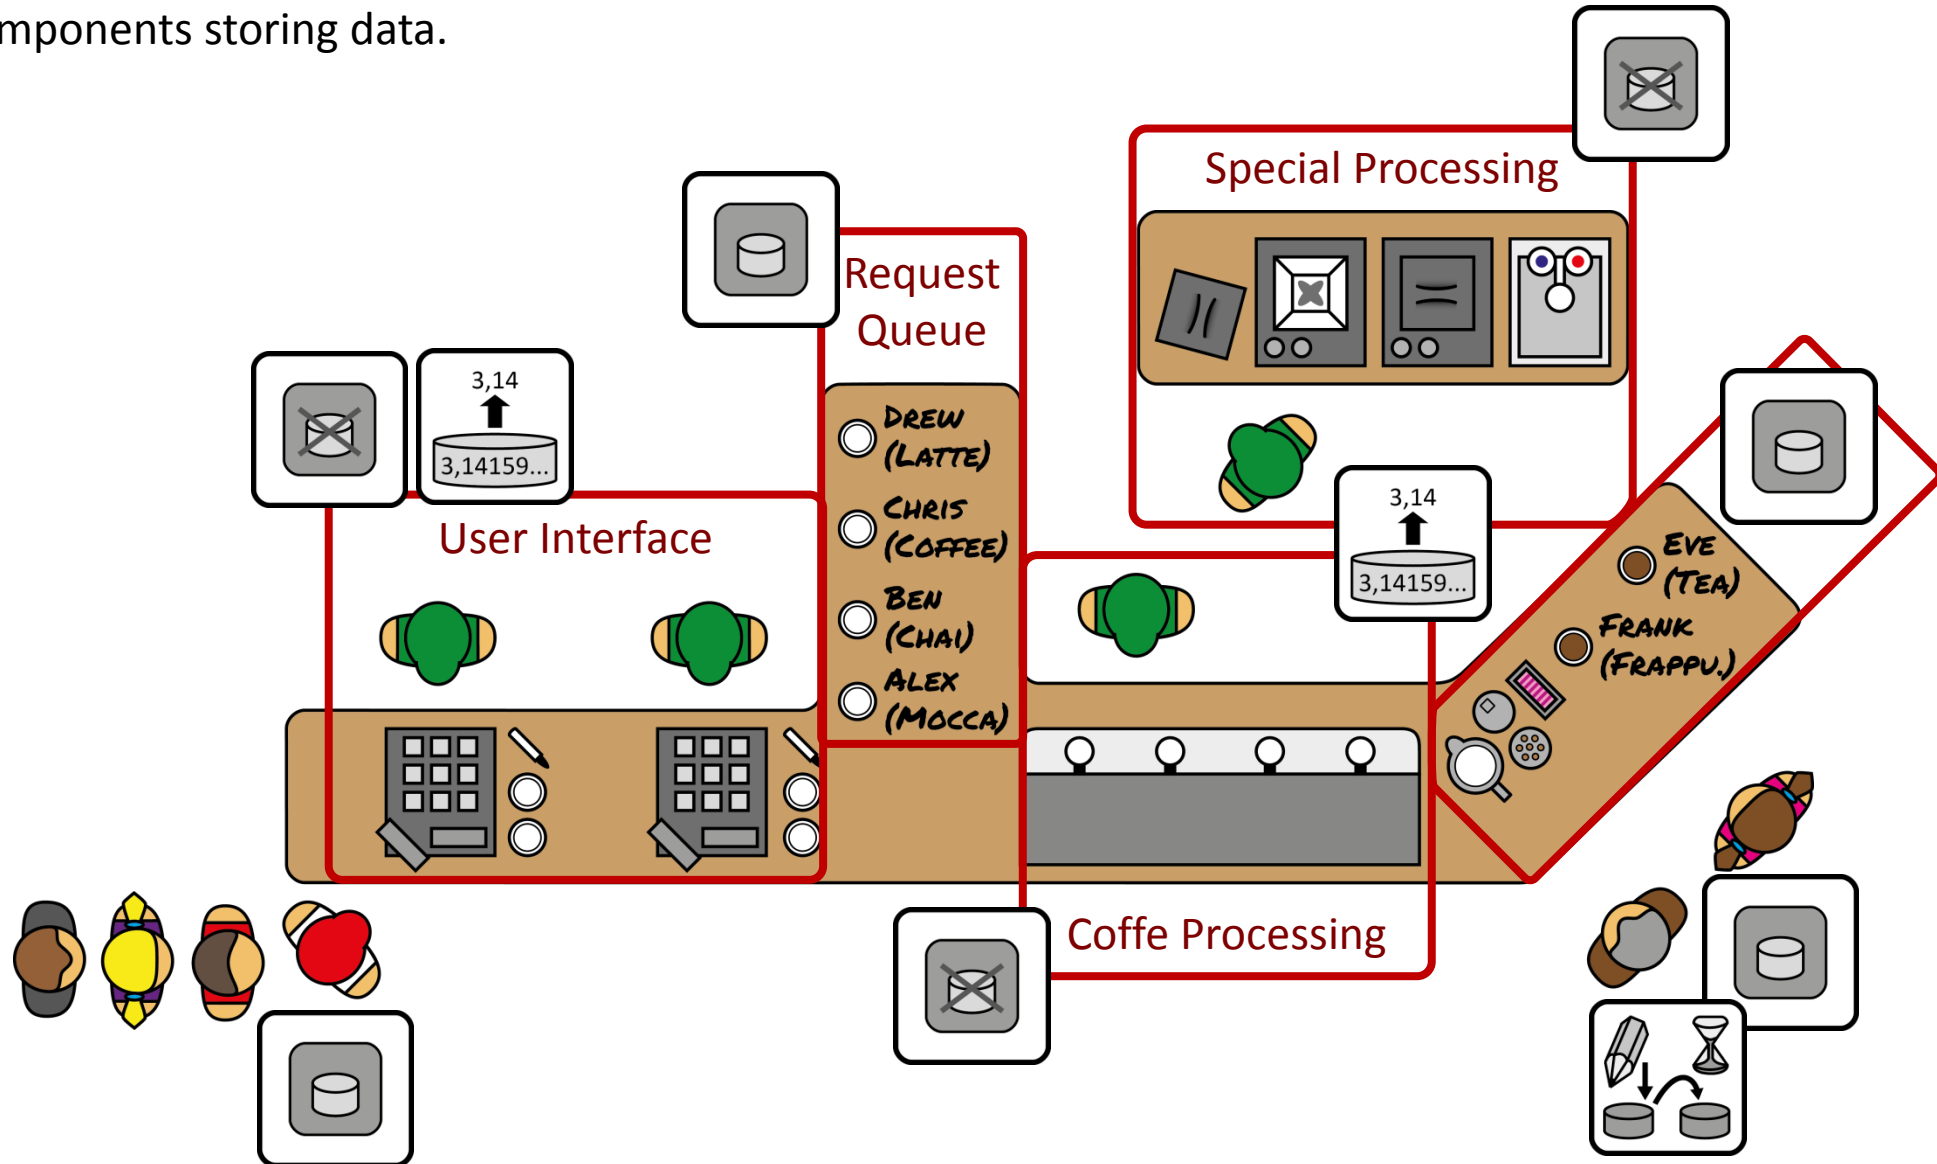


Cloud-Native Application

- ☒ Isolated State
- ☐ Distribution
- ☐ Elasticity
- ☒ Automated Management
- ☐ Loose Coupling


Coffee Shop – Data

Identify components storing data.






 **How are components implemented?**




Message-oriented Middleware

Asynchronous communication is provided while hiding complexity of addressing, routing, or data formats to make interaction robust and flexible.




User Interface Component

Synchronous user interfaces are accessed by humans, while application-internal interaction is realized asynchronously to ensure loose coupling.




Processing Component

Processing functionality is handled by elastically scaled components.



Batch Processing Component

Requests are delayed until environmental conditions make their processing feasible.



Multi-component Image

Virtual servers host multiple application components that may not be active at all times to reduce provisioning and decommissioning operations.




Cloud-Native Application

- ☒ Isolated State
- ☐ Distribution
- ☐ Elasticity
- ☐ Automated Management
- ☒ Loose Coupling


Decide how to implement components.

Decide how to implement components.






Elasticity and Resiliency




Elastic Load Balancer

The number of **synchronous accesses** to an elastically scaled-out application is used to determine the **number of** required application component **instances**.




Elastic Queue

The number of **accesses via messaging** is used to **adjust the number of** required application component **instances**.



Node-based Availability

A cloud provider guarantees the **availability of nodes**, such as individual virtual **servers**, **middleware** components or hosted **application components**.



Environment-based Availability

A cloud provider guarantees the **availability of the environment hosting** individual nodes, such as virtual **servers** or hosted **application components**.



Watchdog

Applications **cope with failures by monitoring and replacing** application component instances if the provider-assured availability is insufficient.

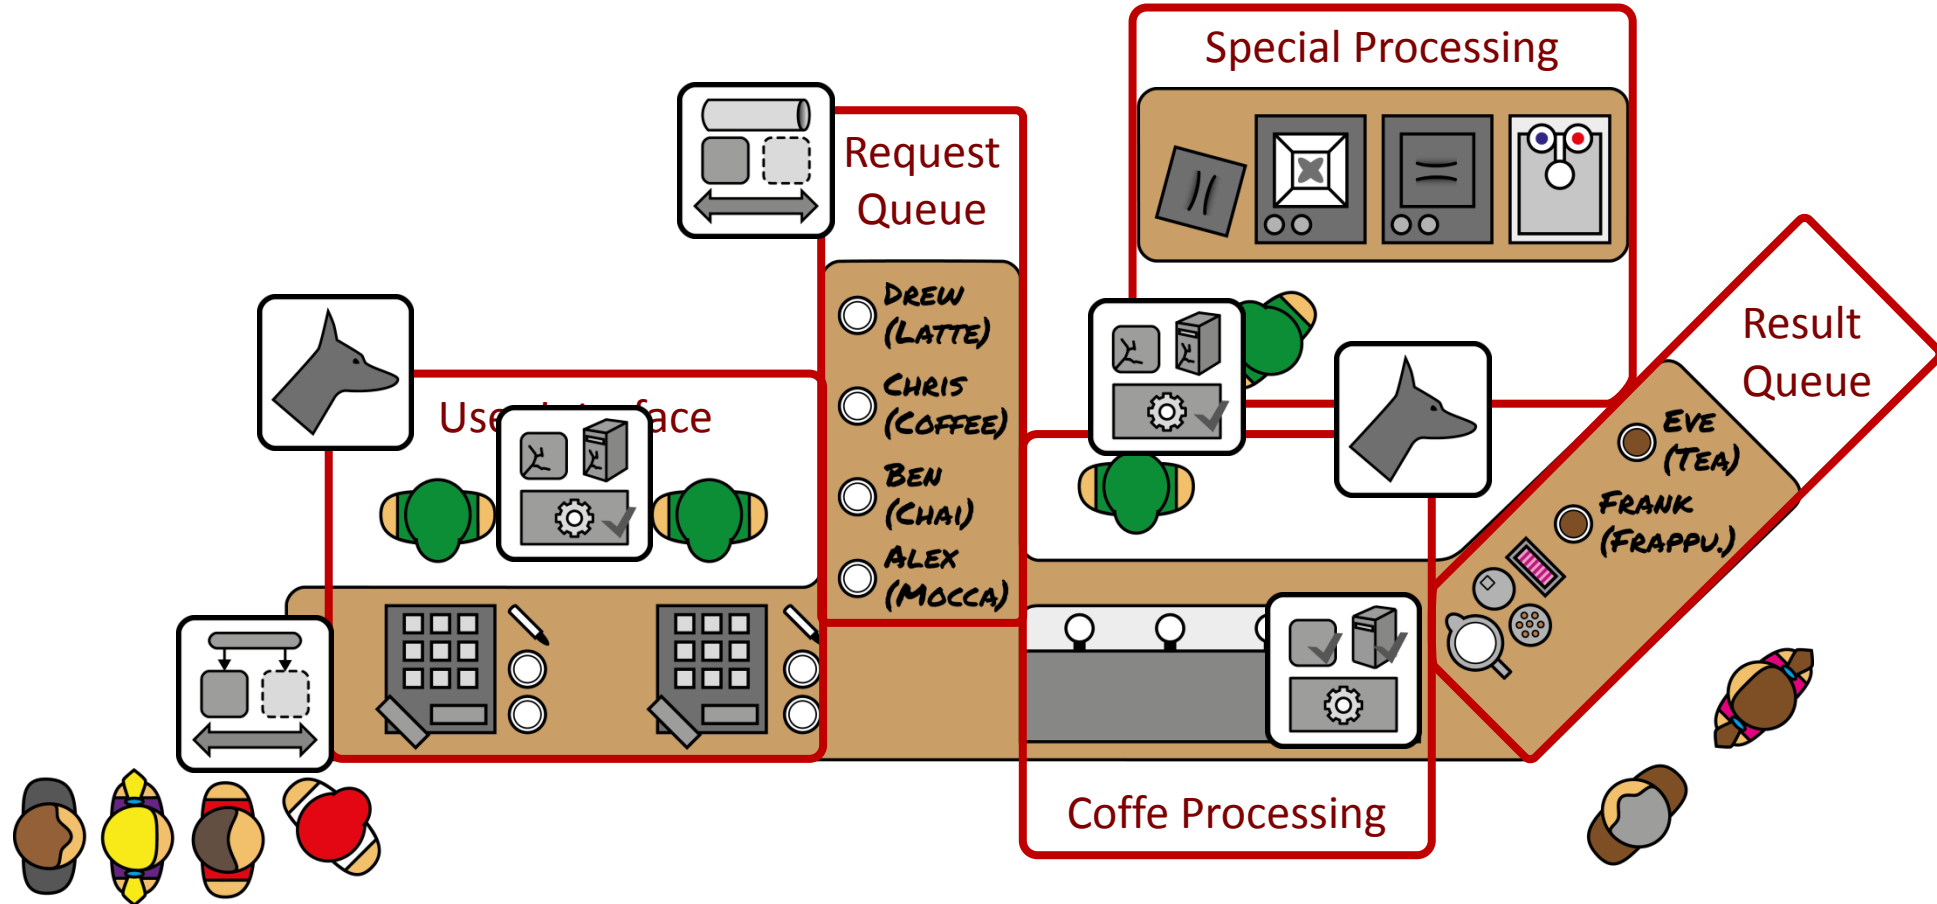


Cloud-Native Application

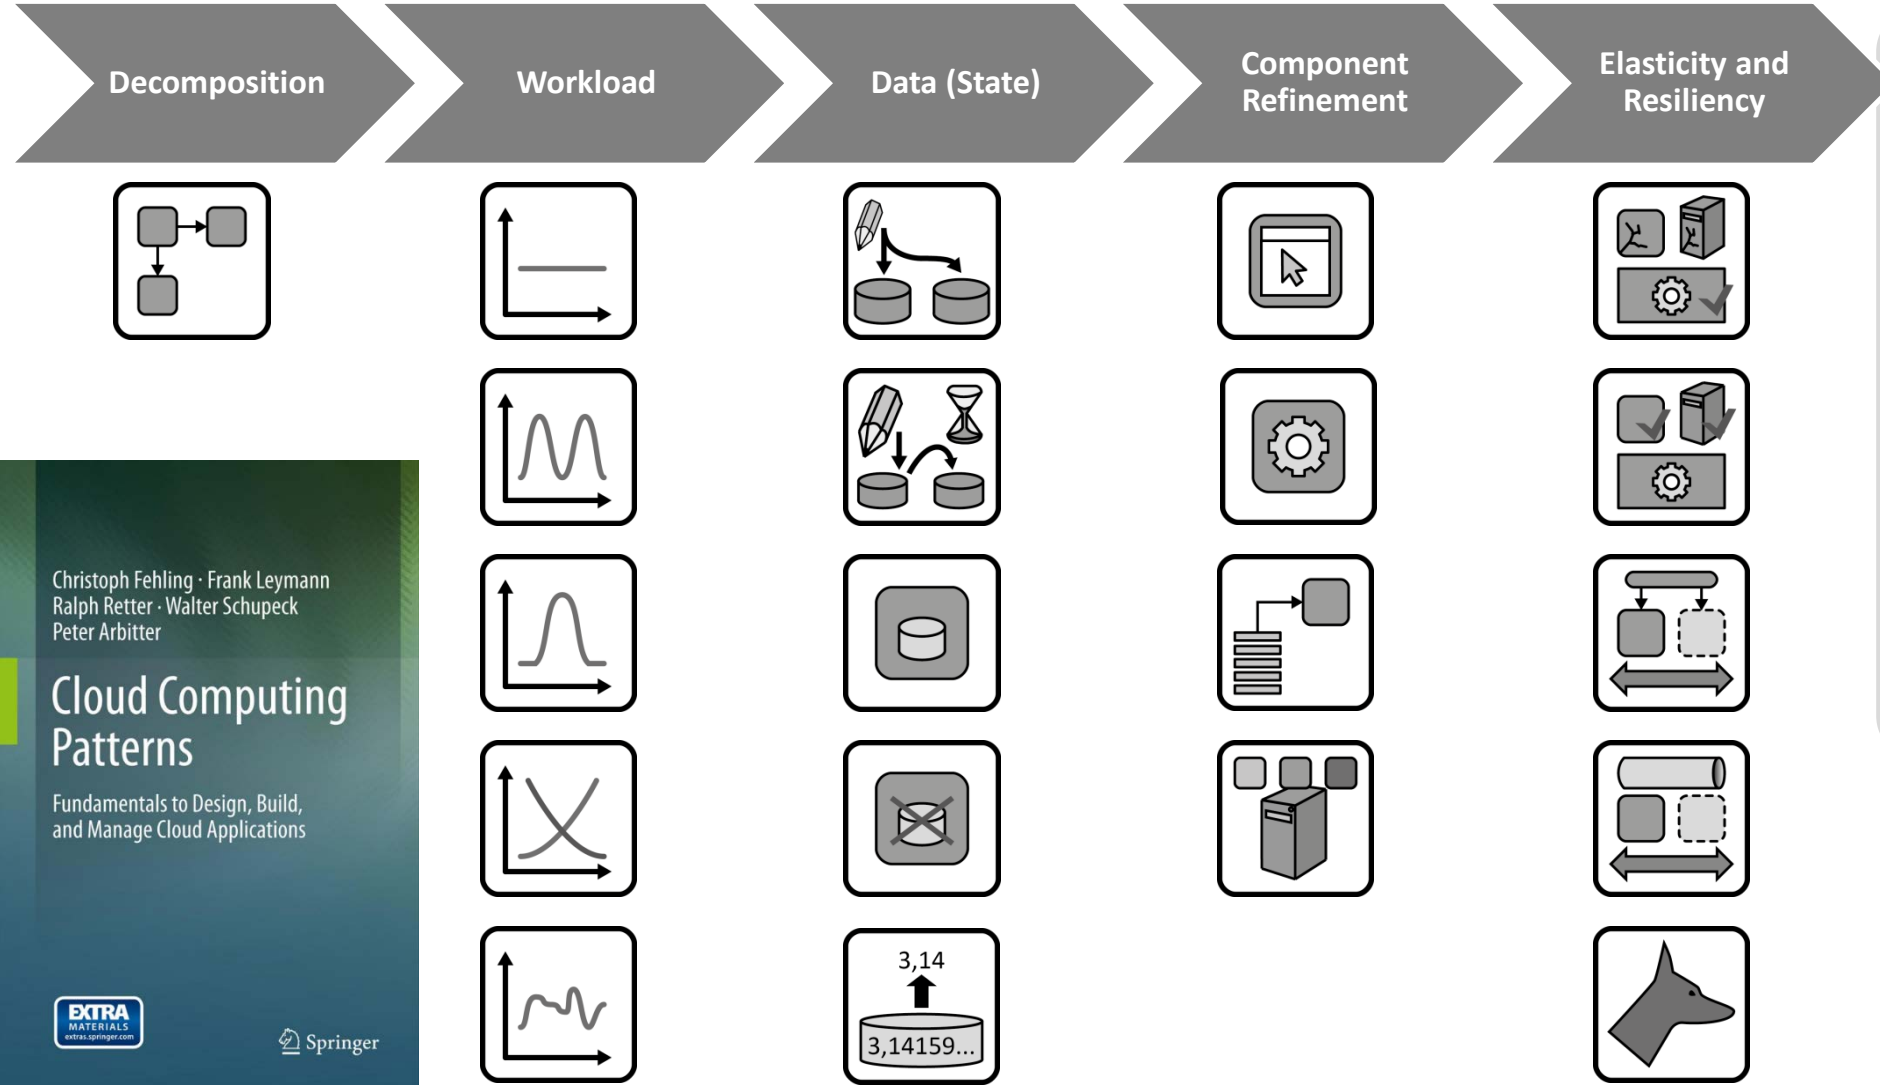
- ☐ Isolated State
- ☐ Distribution
- ☒ Elasticity
- ☒ Automated Management
- ☐ Loose Coupling

Elasticity and Resiliency

What shall happen if workload changes or something fails?



Design Steps for Cloud Applications using Patterns



IDEAL

Cloud-Native Application

- ✓ Isolated State
- ✓ Distribution
- ✓ Elasticity
- ✓ Automated Management
- ✓ Loose Coupling

