

UK Systems 24 - Encrypting data-at-rest in Apache Kafka[®]

Keith Wall
Senior Principal Software Engineer,
Red Hat

Why

Some consequences of a Data Breach

Legal Implications

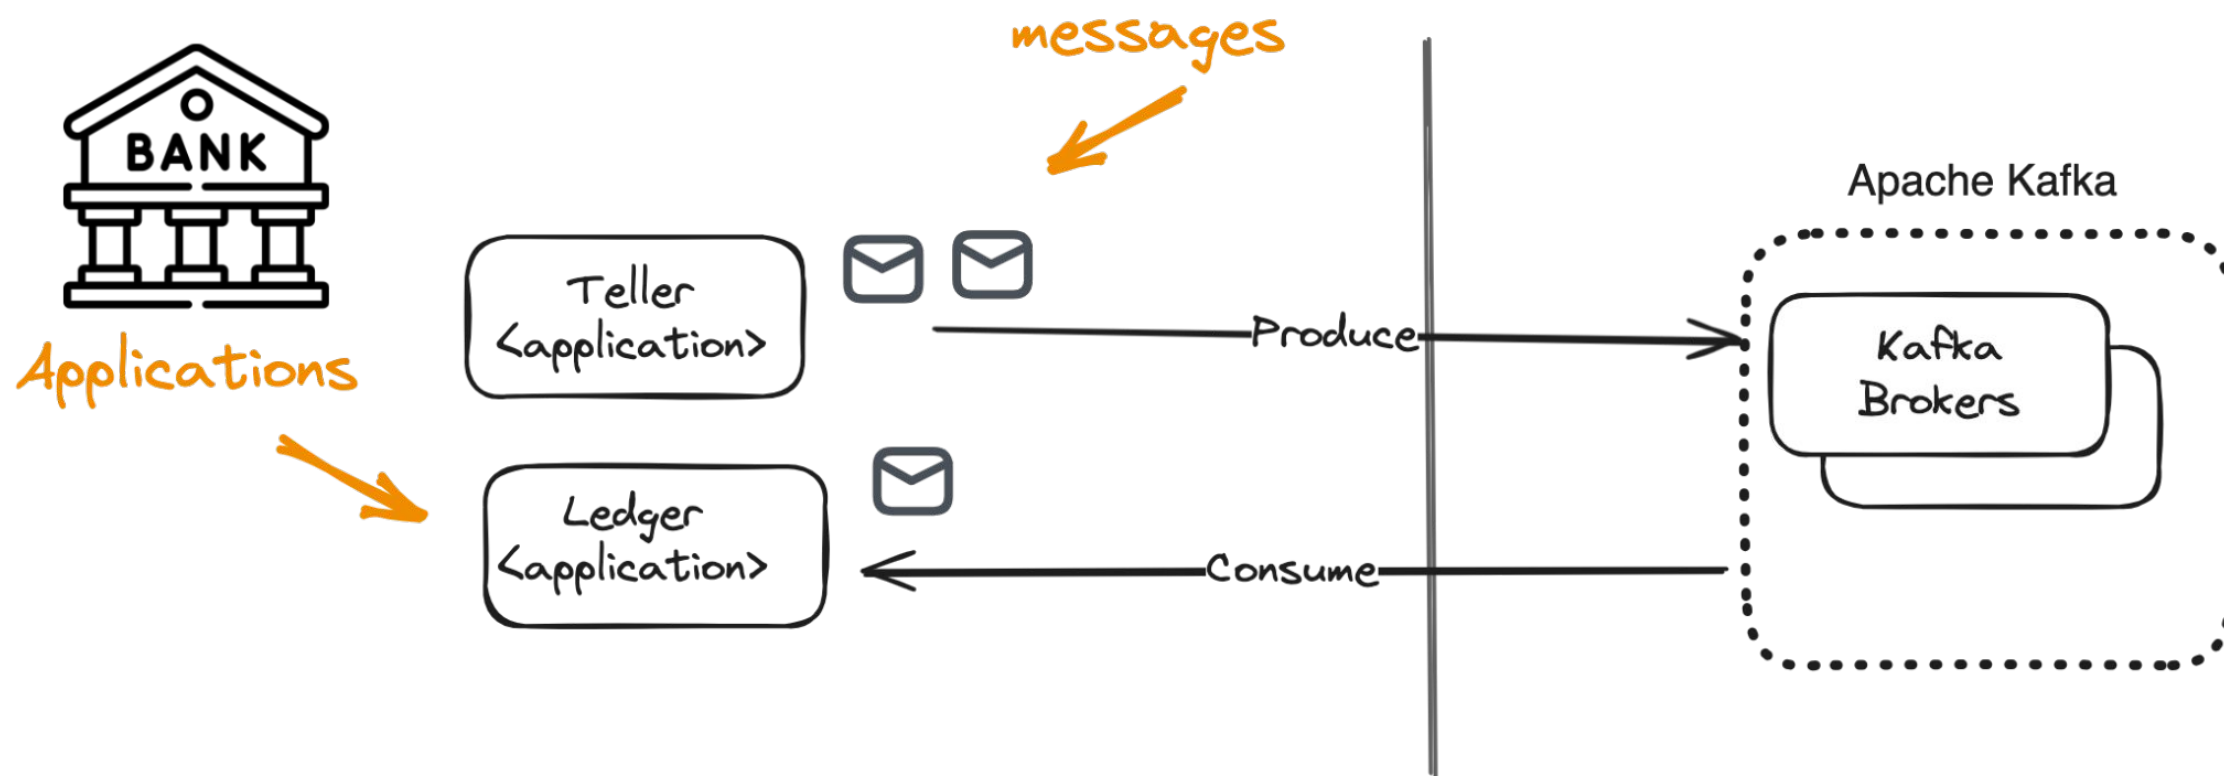
Reputational Damage

Harm to Individuals

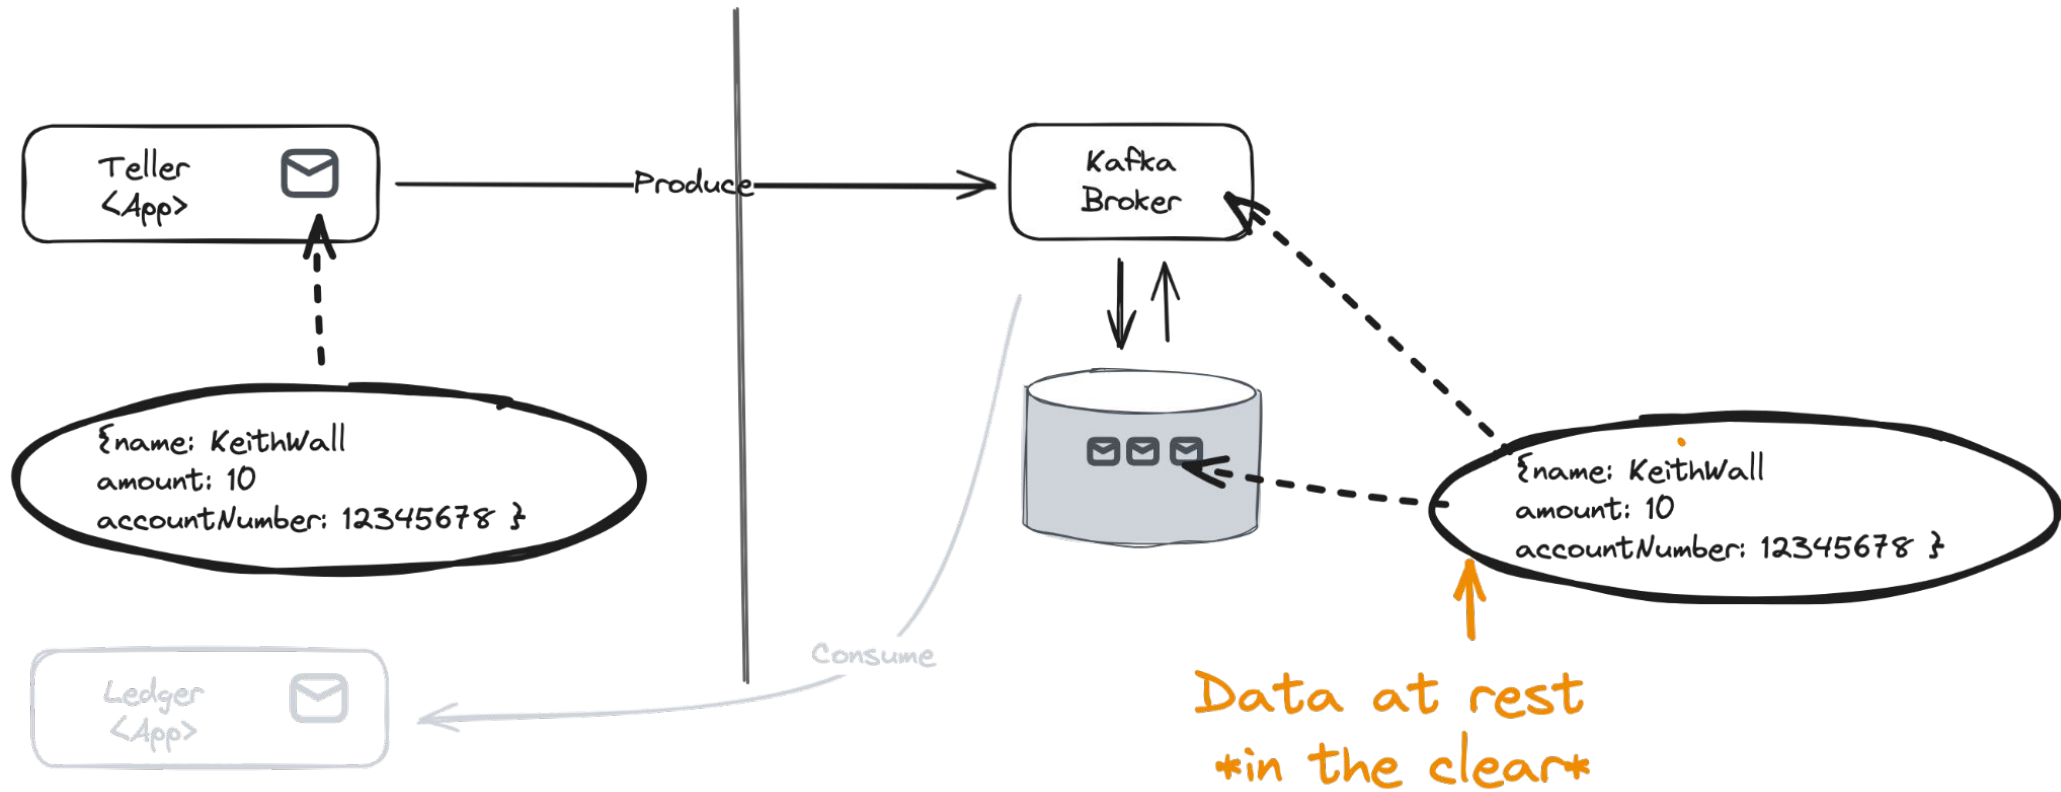
.....

.....

What is Apache Kafka[®]?



Data at rest stored in the clear



So, you trust your Kafka Admin, right?

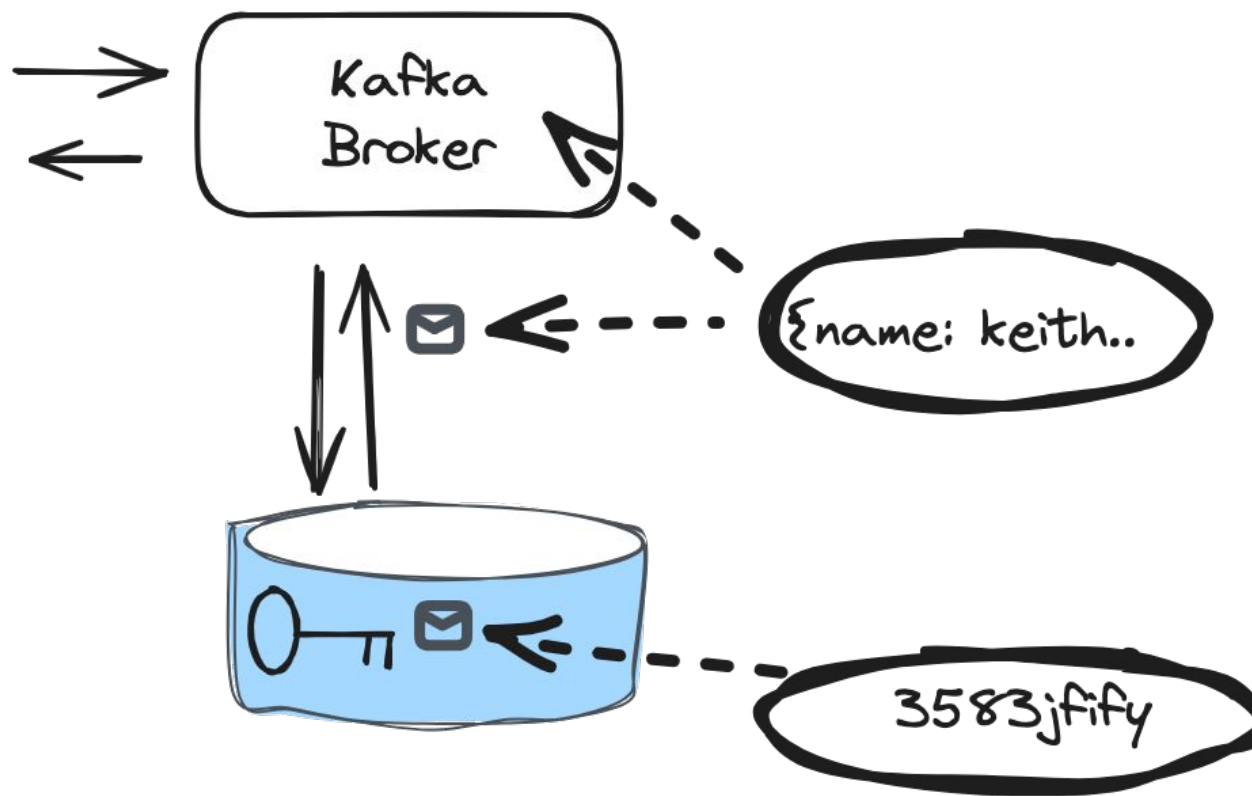


Not all bad guys looks like bad guys.

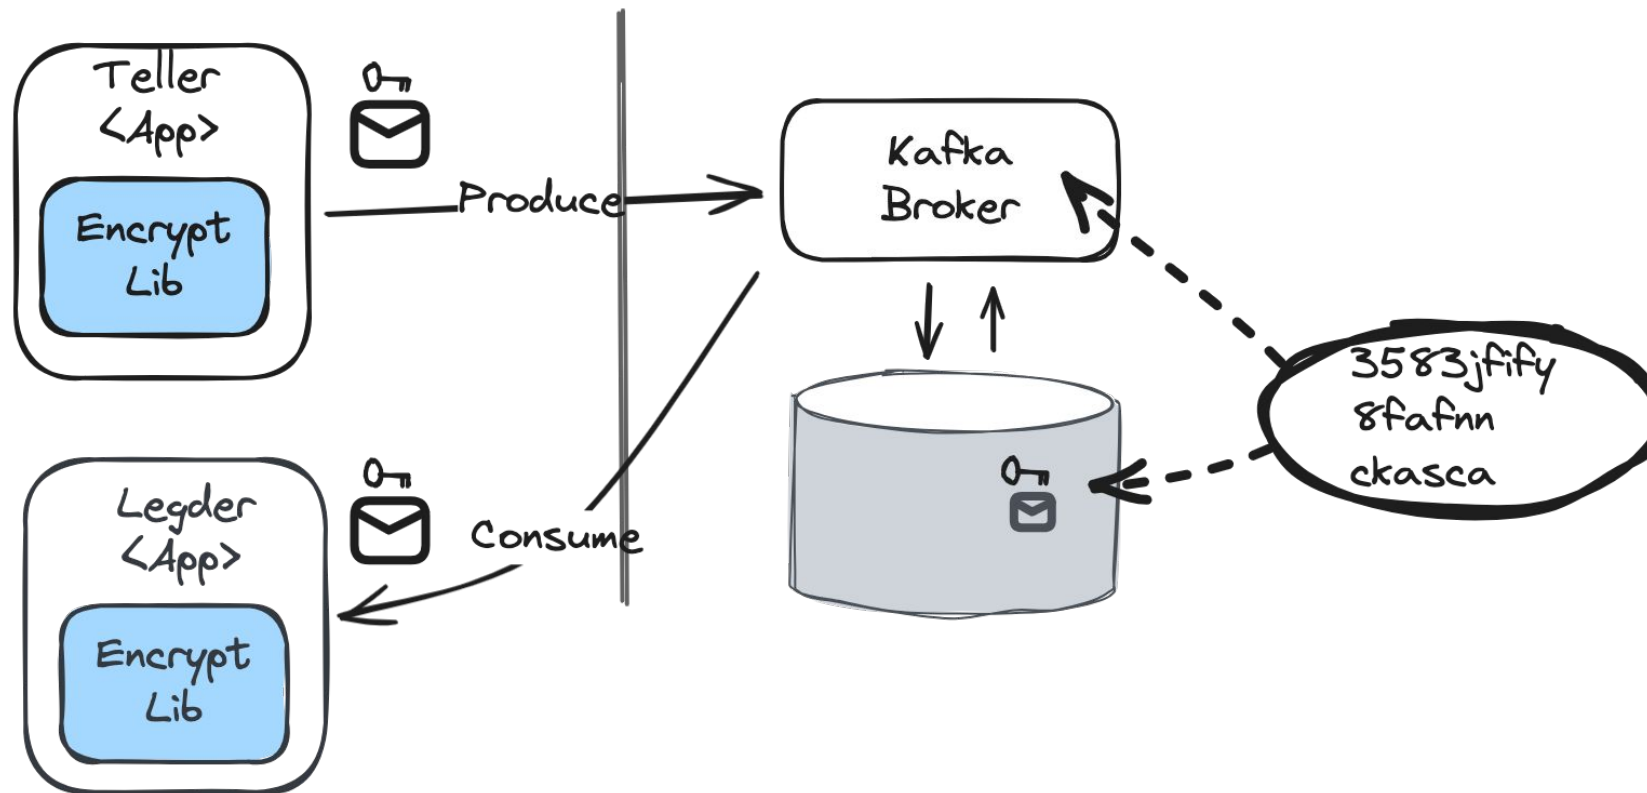


Use Encryption.. but where?

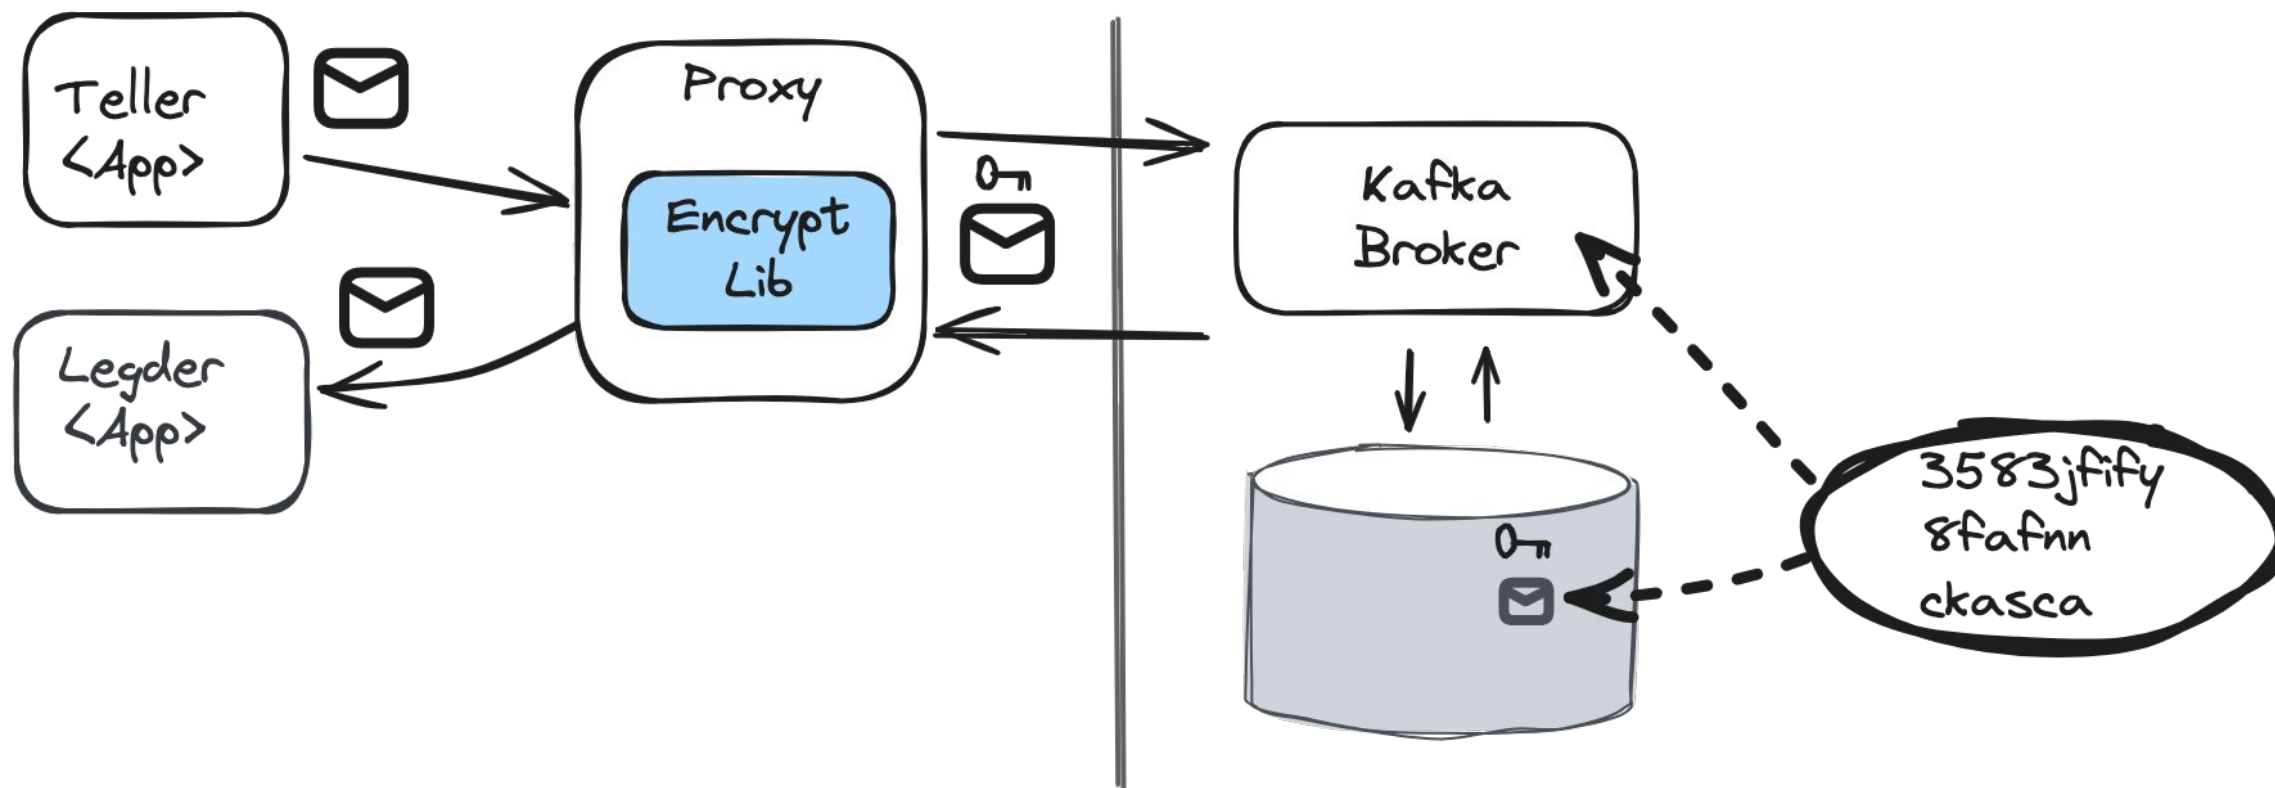
Disk Encryption



On the application side



With a proxy




What to encrypt?

Introducing the Kafka Record

```
{  
key: <bytes>  
value: <bytes>  
headers: [ { key: <bytes> , value: <bytes> } ]  
timestamp: long  
}
```

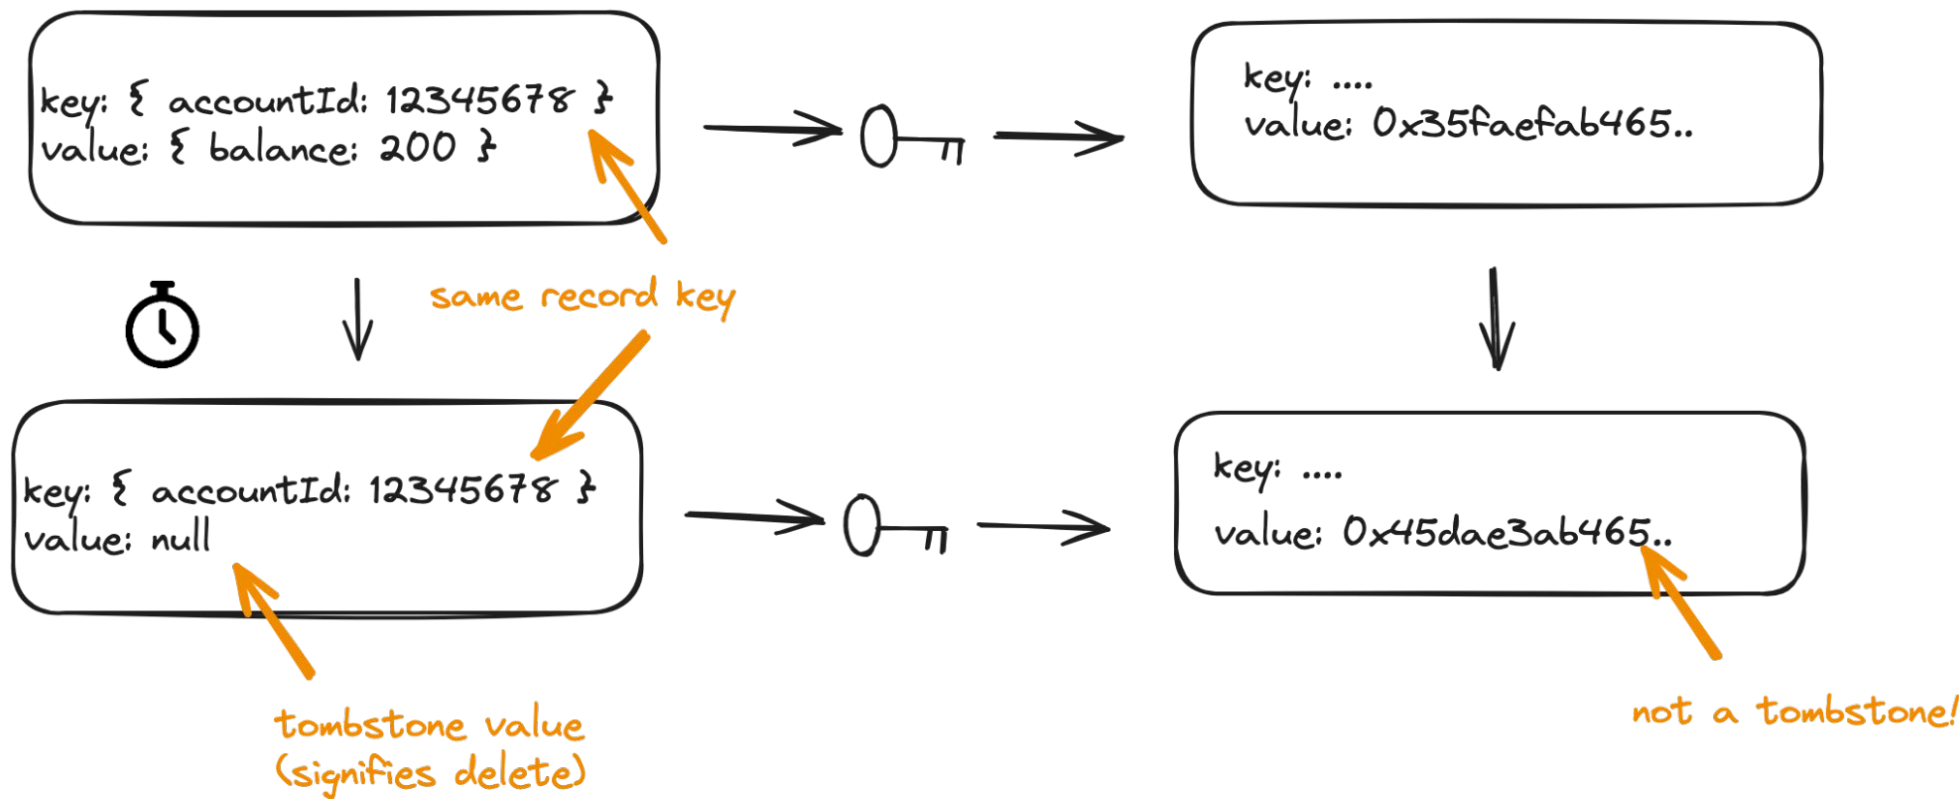
Encrypting the Kafka Record (v0)

```
key: { txnId: 462742 }  
value: { accountId: 12345678  
        name: Keith  
        value: 200  
      }  
headers: [  
  {key: "auth", value: "PASSPORT"},  
  {key: "authid", value: "89482893:"}],  
timestamp: 123456789
```

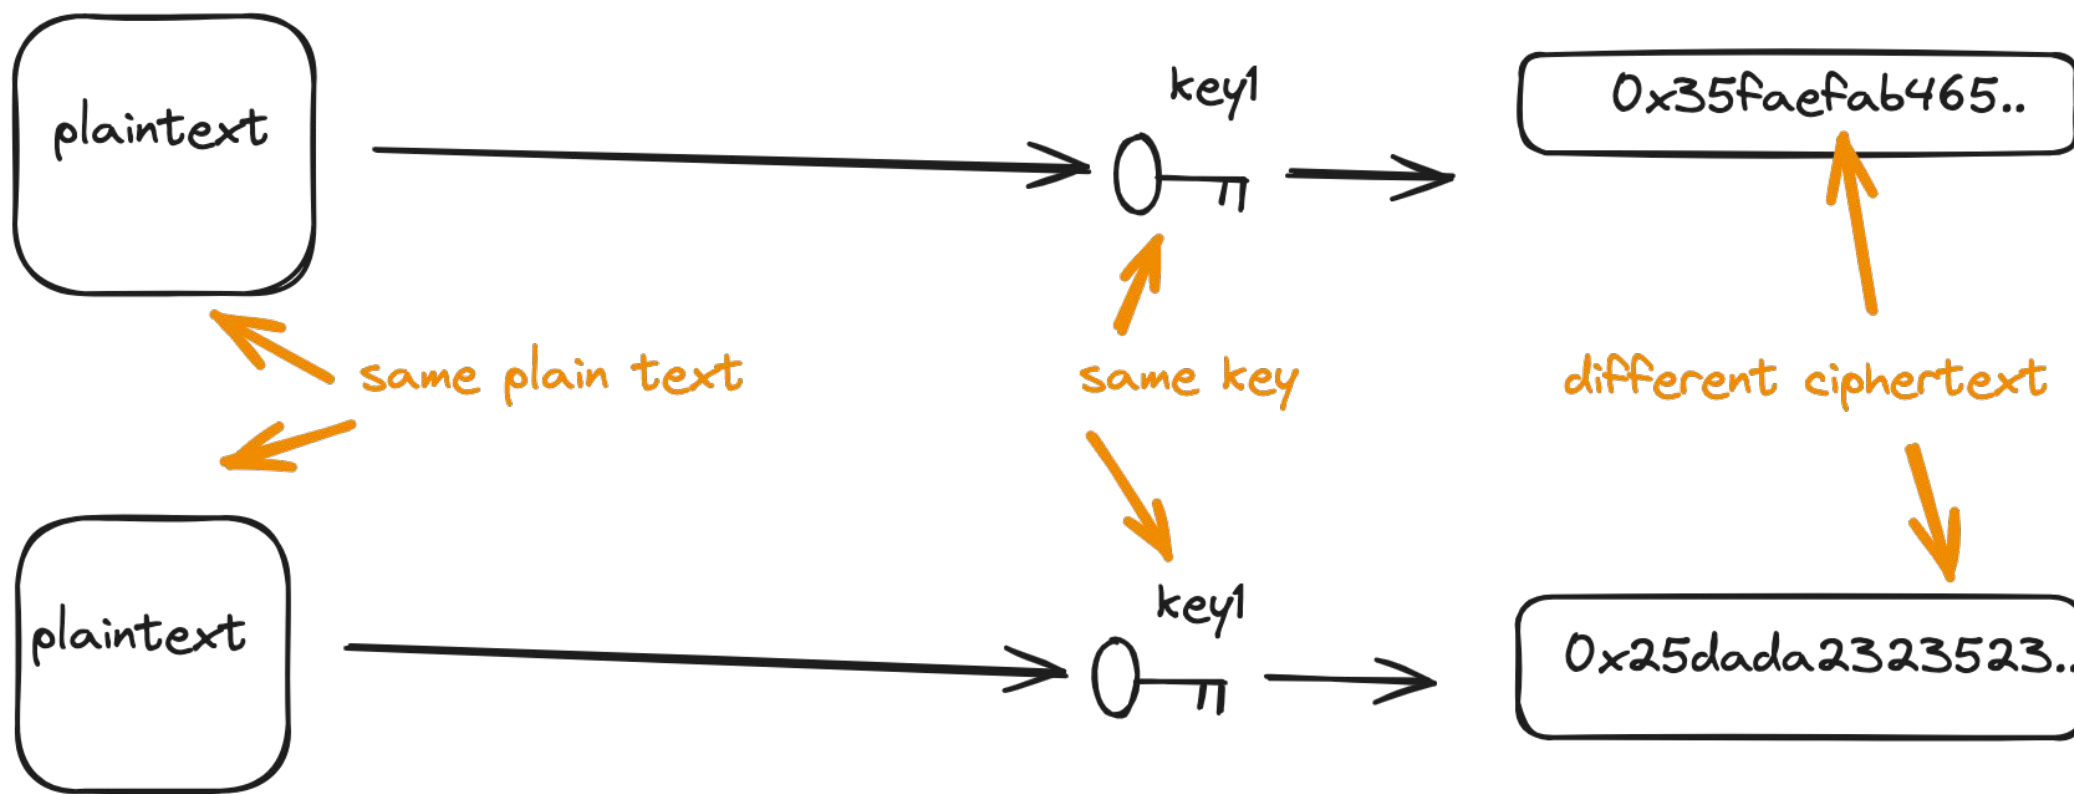


```
key: 0x25dada2323523..  
value: 0x35faefab465..  
headers: [  
  {key: 0x34748.., value: 0xbcd43232.},  
  {key: 0x5238d.., value: 0x523abcd.}],  
timestamp: 123456789
```

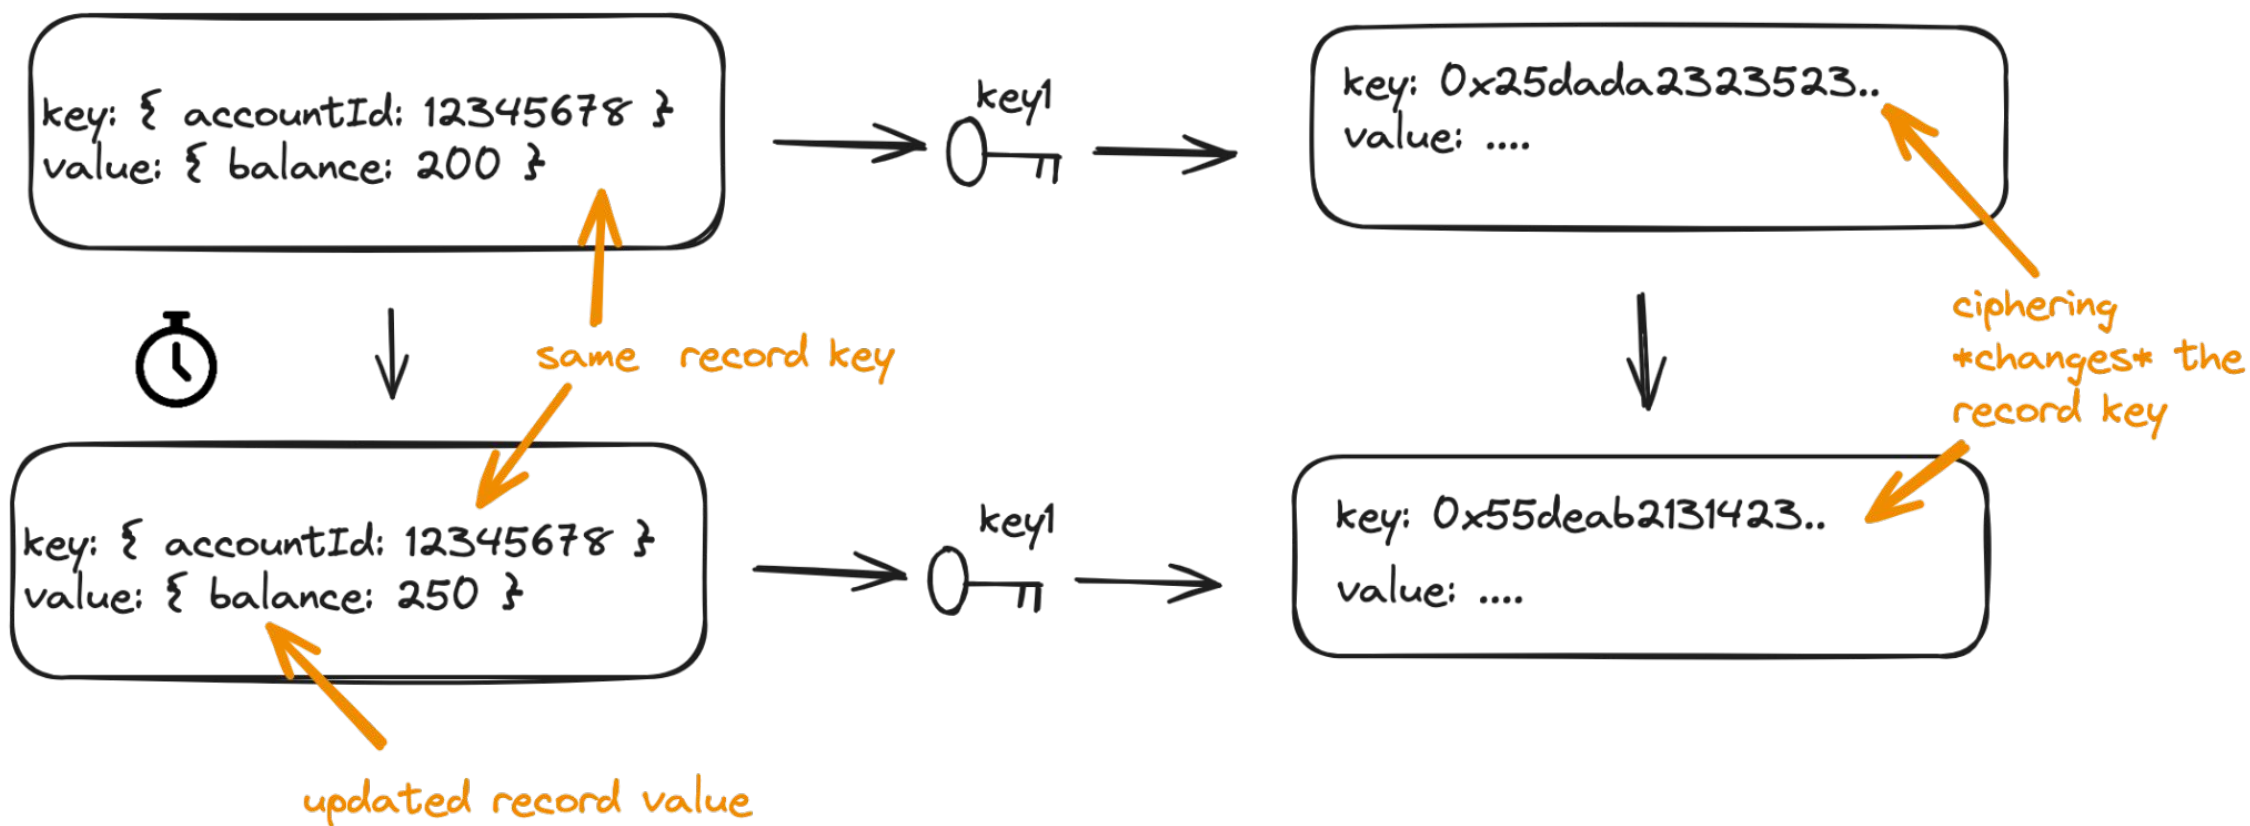
Compacted Topics: tombstone problem




Semantic Encryption



Compacted Topics: unstable record keys problematic




Compacted topic: stabilizing the record keys



Summary

We have covered:

- Some motivations for encrypting data at rest in Kafka
- How Kafka stores data in the clear
- Options for introducing encryption into Kafka
- Some of the challenges when encrypting Kafka Records



Kroxylicious

Kroxylicious, the snappy open source proxy for Apache Kafka®

Is a pluggable open-source proxy for Apache Kafka where we are exploring these encryption ideas and much more.

Come and find out more.

<https://kroxylicious.io>

<https://kroxylicious.slack.com/>