Red Hat Summit

Connect

Managed Kafka jak zacząć?

Andrzej Kowalczyk

Red Hat | intel

Cloud Services Managed by Red Hat

Managed OpenShift + Application Services + Data Services

Full stack management and unified experience

Maximize full value of Red Hat OpenShift

Hybrid cloud flexibility

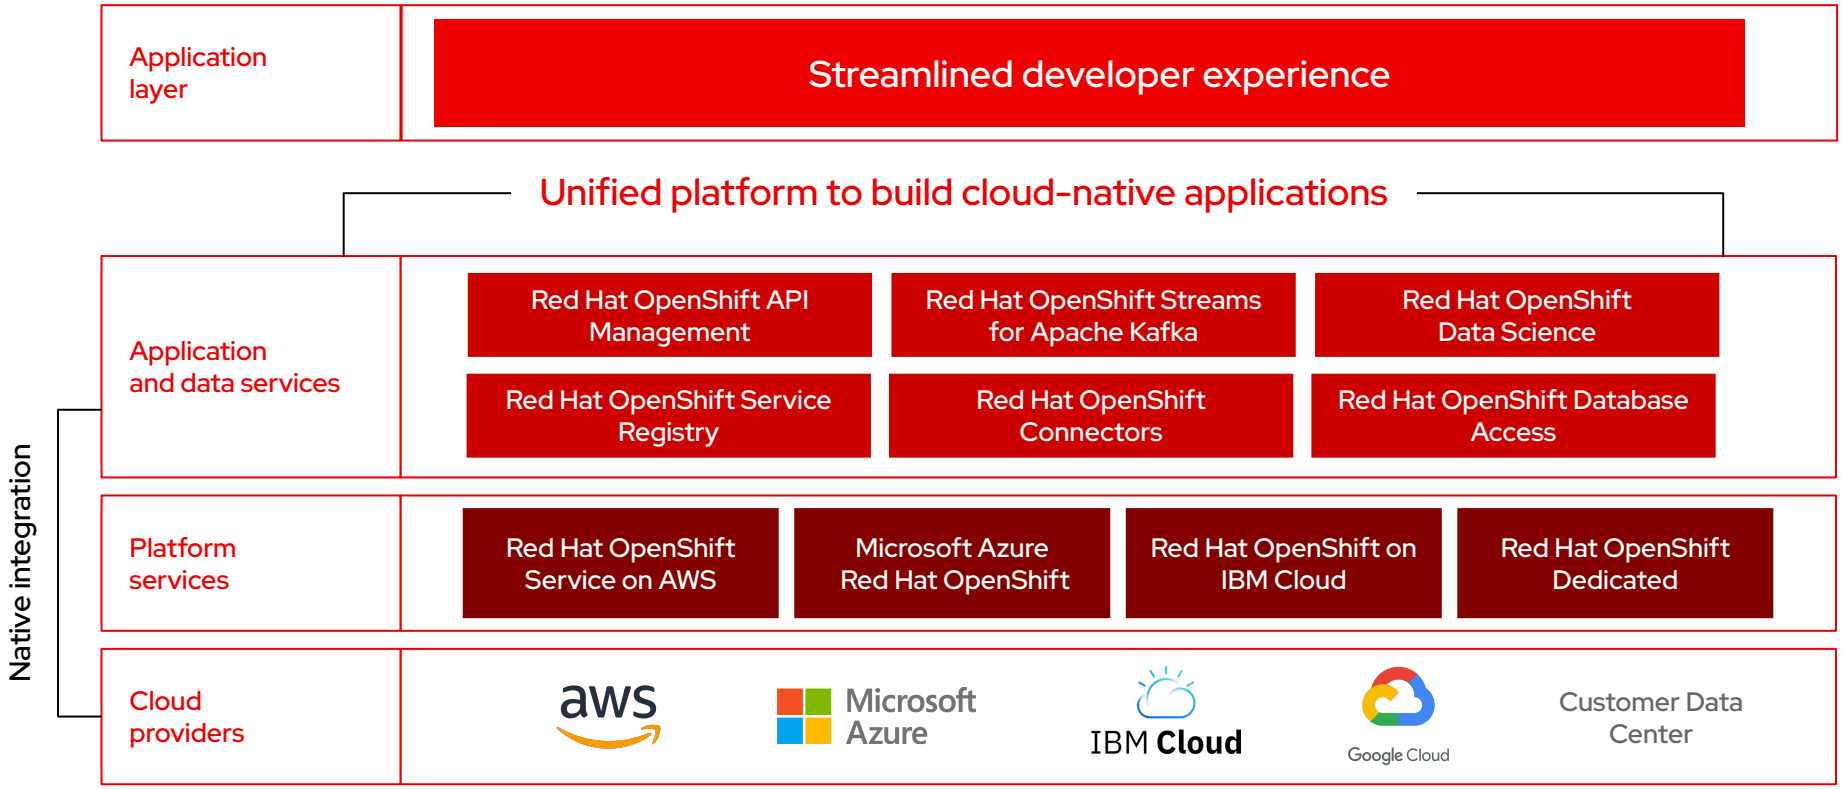

Red Hat | intel

Red Hat OpenShift Streams for Apache Kafka

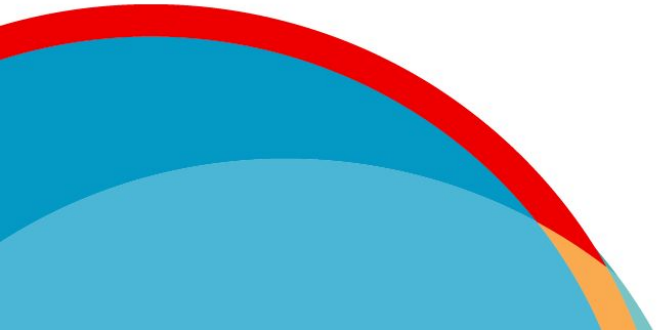

Red Hat | intel

Choosing your Streams

Self Managed

Red Hat Integration (AMQ Streams)

Physical

Virtual

Private cloud

Public cloud

Managed cloud
(Azure, AWS, IBM, Red Hat)

Fully Managed by Red Hat

Red Hat OpenShift Streams for Apache Kafka

Managed cloud
(AWS)

Red Hat

The value of Red Hat OpenShift Streams for Apache Kafka

Faster application velocity

Begin developing immediately and continuously respond to change

Unified experience across all clouds

Seamlessly connects applications across public and private clouds

Kafka ecosystem for streams-based applications

Delivers a curated set of cloud services to simplify delivery of stream-based applications

Red Hat OpenShift Streams for Apache Kafka - Key features

Complete solution for stream-based applications

Real-time, streaming data broker - Dedicated Apache Kafka cluster delivered as a service in the cloud and location of choice

Schema registry - OpenShift Service Registry is included, making it easy for dev teams to publish, communicate and discover any streaming data topics.

Delivered as a service, managed by Red Hat - Reduce operational overhead and complexity with pre-built connectors managed by Red Hat.

Streamlined developer experience - developer-first, consistent experience that shields the user from administrative tasks

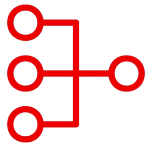

Connectors - securely connect to distributed services to consume and share streaming data between apps, enterprise systems, and cloud provider services.

Red Hat OpenShift Streams for Apache Kafka

Red Hat managed solution for stream-based applications

Streamlined developer experience: a curated solution with a developer-first, consistent experience

Delivered as a service, managed by Red Hat SRE - 24x7 global support and a 99.95% service-level agreement (SLA)

Real-time, streaming data broker - Dedicated Apache Kafka cluster deliver as a service in the cloud and location of choice

- Access to Kafka brokers, topics, and partitions
- Managed ZooKeeper
- Metrics and monitoring
- Integrated identity & access management

Hosted & managed service offering

Red Hat cloud services are managed and operated by Red Hat's Site Reliability Engineers

- SREs serve as the cloud provider account owner and cluster administrator owning the **SLA (99.95%) and uptime**
- Responsible for the **24x7 support** for all managed and hosted environments
  - Including for building, installing, upgrading, **managing and maintaining every cluster**
- SRE teams are distributed **across 3 regions**: APAC, EMEA and Americas
- The team ensures **open communications channels** centralized around the dedicated customer portal

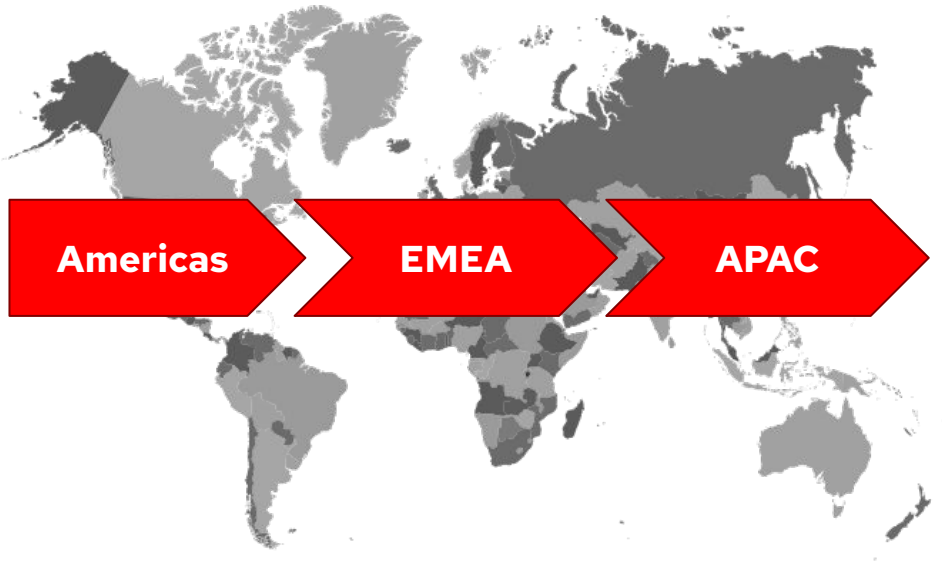

Red Hat Management

We deliver premium support and 99.95% uptime

Planned & Emergency Maintenance Events

Included and performed by Red Hat Service Reliability Engineering (RH SRE)

Monitoring, Logging and Alerting

Available for RH SRE teams to track the performance of the technology stack installed by Red Hat

CVEs

Performed by RH-SRE where Red Hat Product Security dictates

Backups and DR

Performed by RH-SRE where Red Hat Product Security dictates

Red Hat

Red Hat OpenShift Streams for Apache Kafka

Components pre-configured to optimize for high availability and reliability

- OpenShift Streams for Apache Kafka provides **dedicated Kafka instances**
- Service was designed following **experts recommended configurations**
  - Each Kafka cluster consists of **three brokers,** with each broker in a different AZ of the target region
  - **All topics have a replication factor of 3**, resulting in a partition replica placed in each AZ
  - All topics have an **in-sync replica count of 2**
- The Zookeeper environment supporting Kafka instances is completely managed by Red Hat and not accessible to end users

Red Hat OpenShift Streams for Apache Kafka

How to get started?

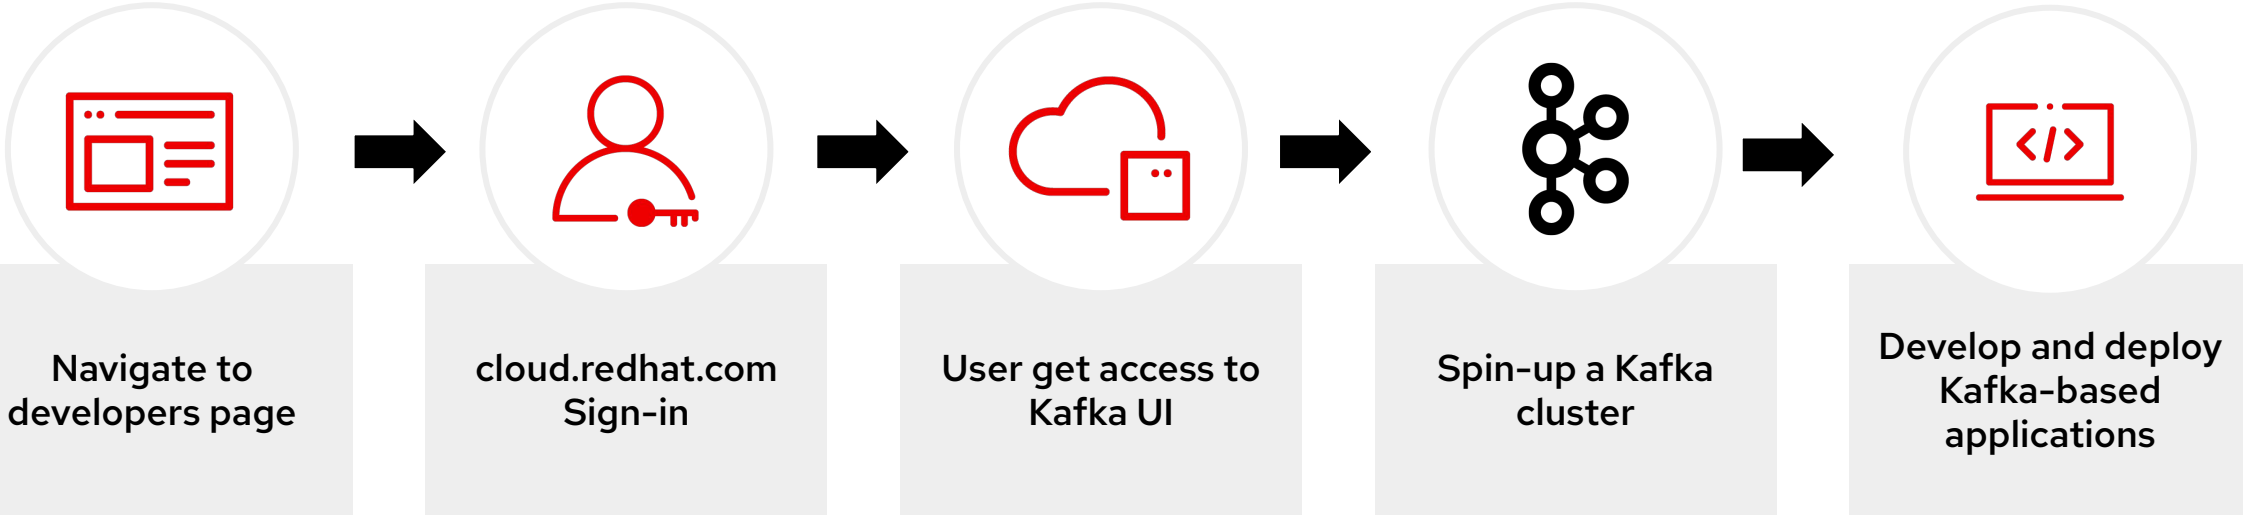

Moje doświadczenia

console.redhat.com - tu się zaczyna

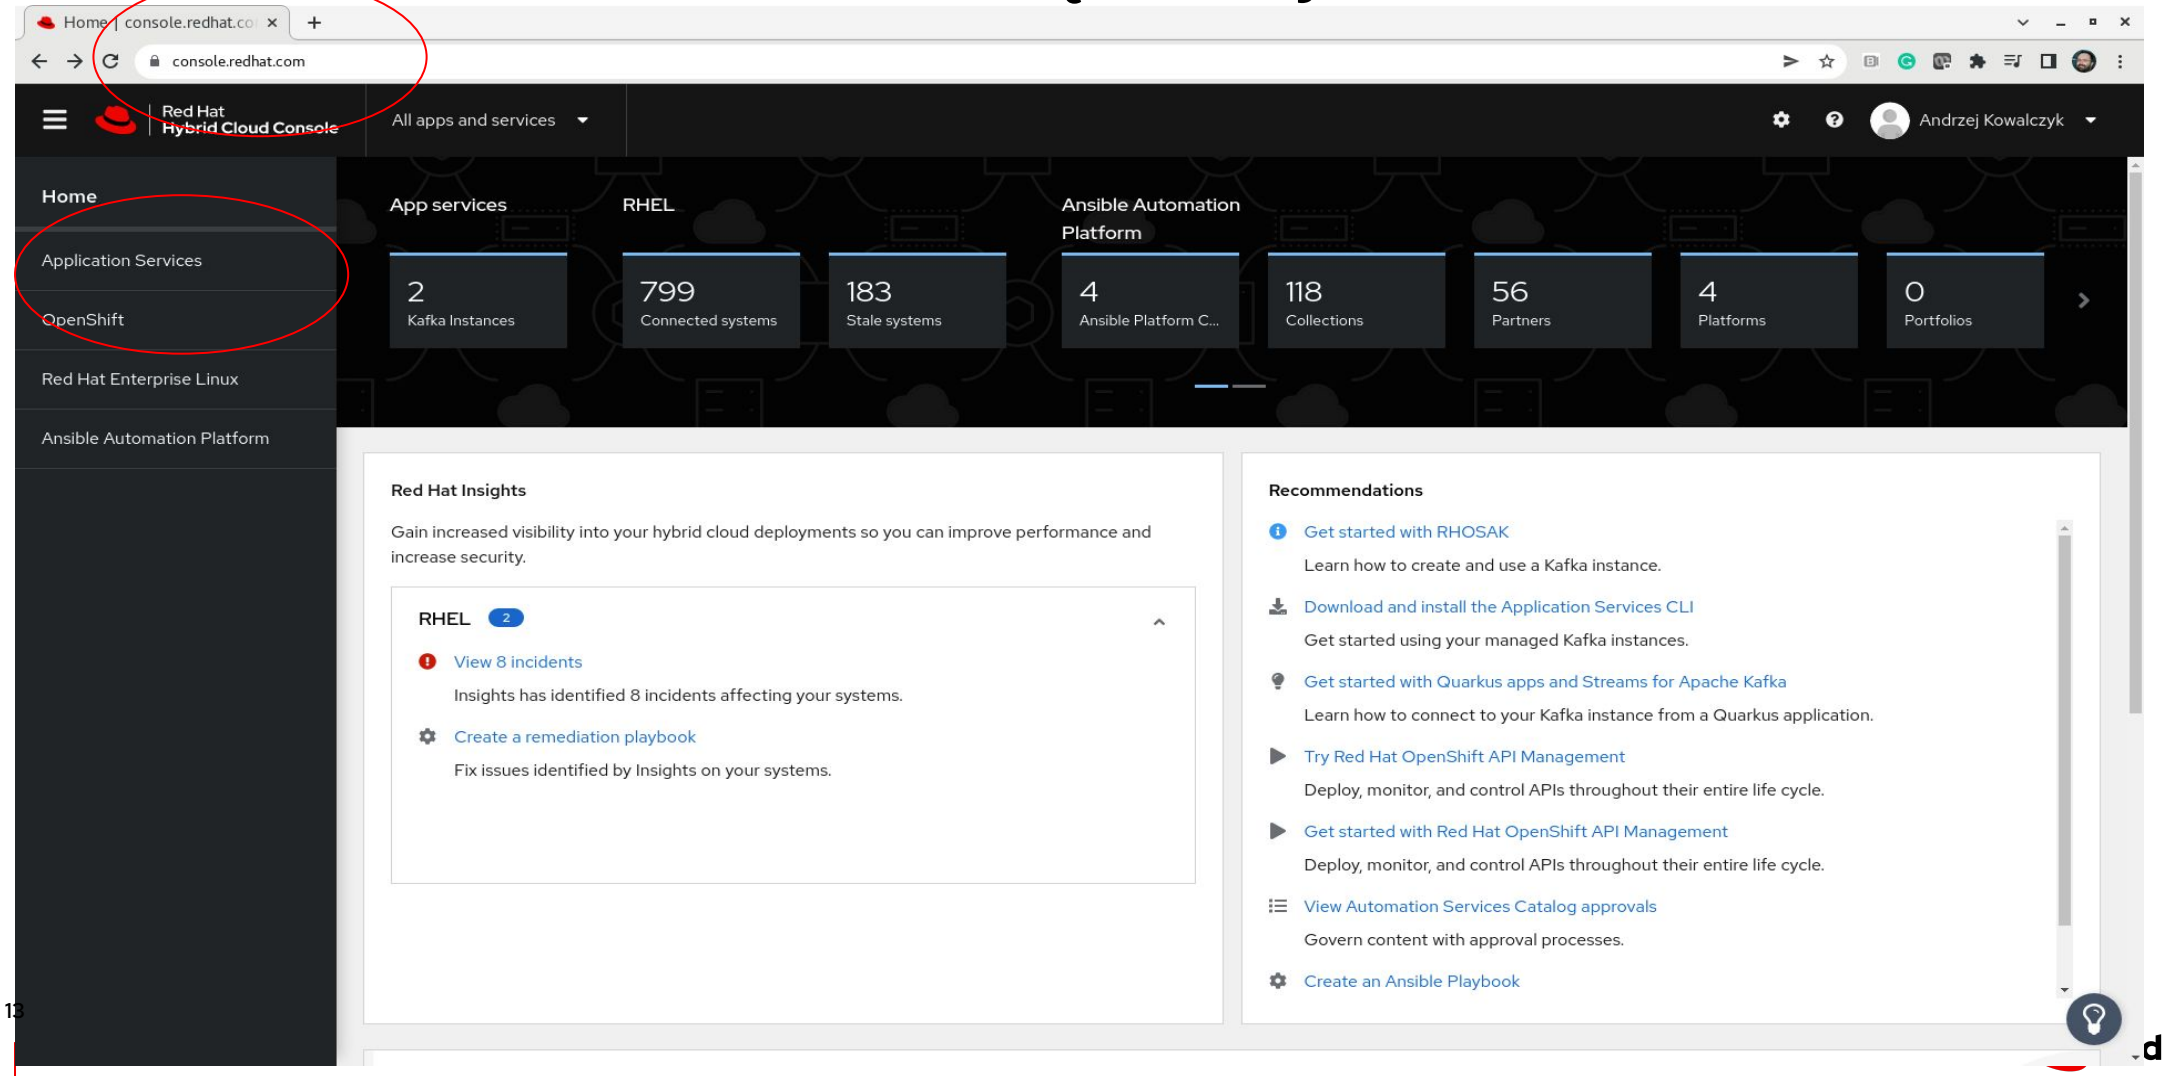

Pokażę jak działa wersja trial (instancja będzie usunięta po 48h)

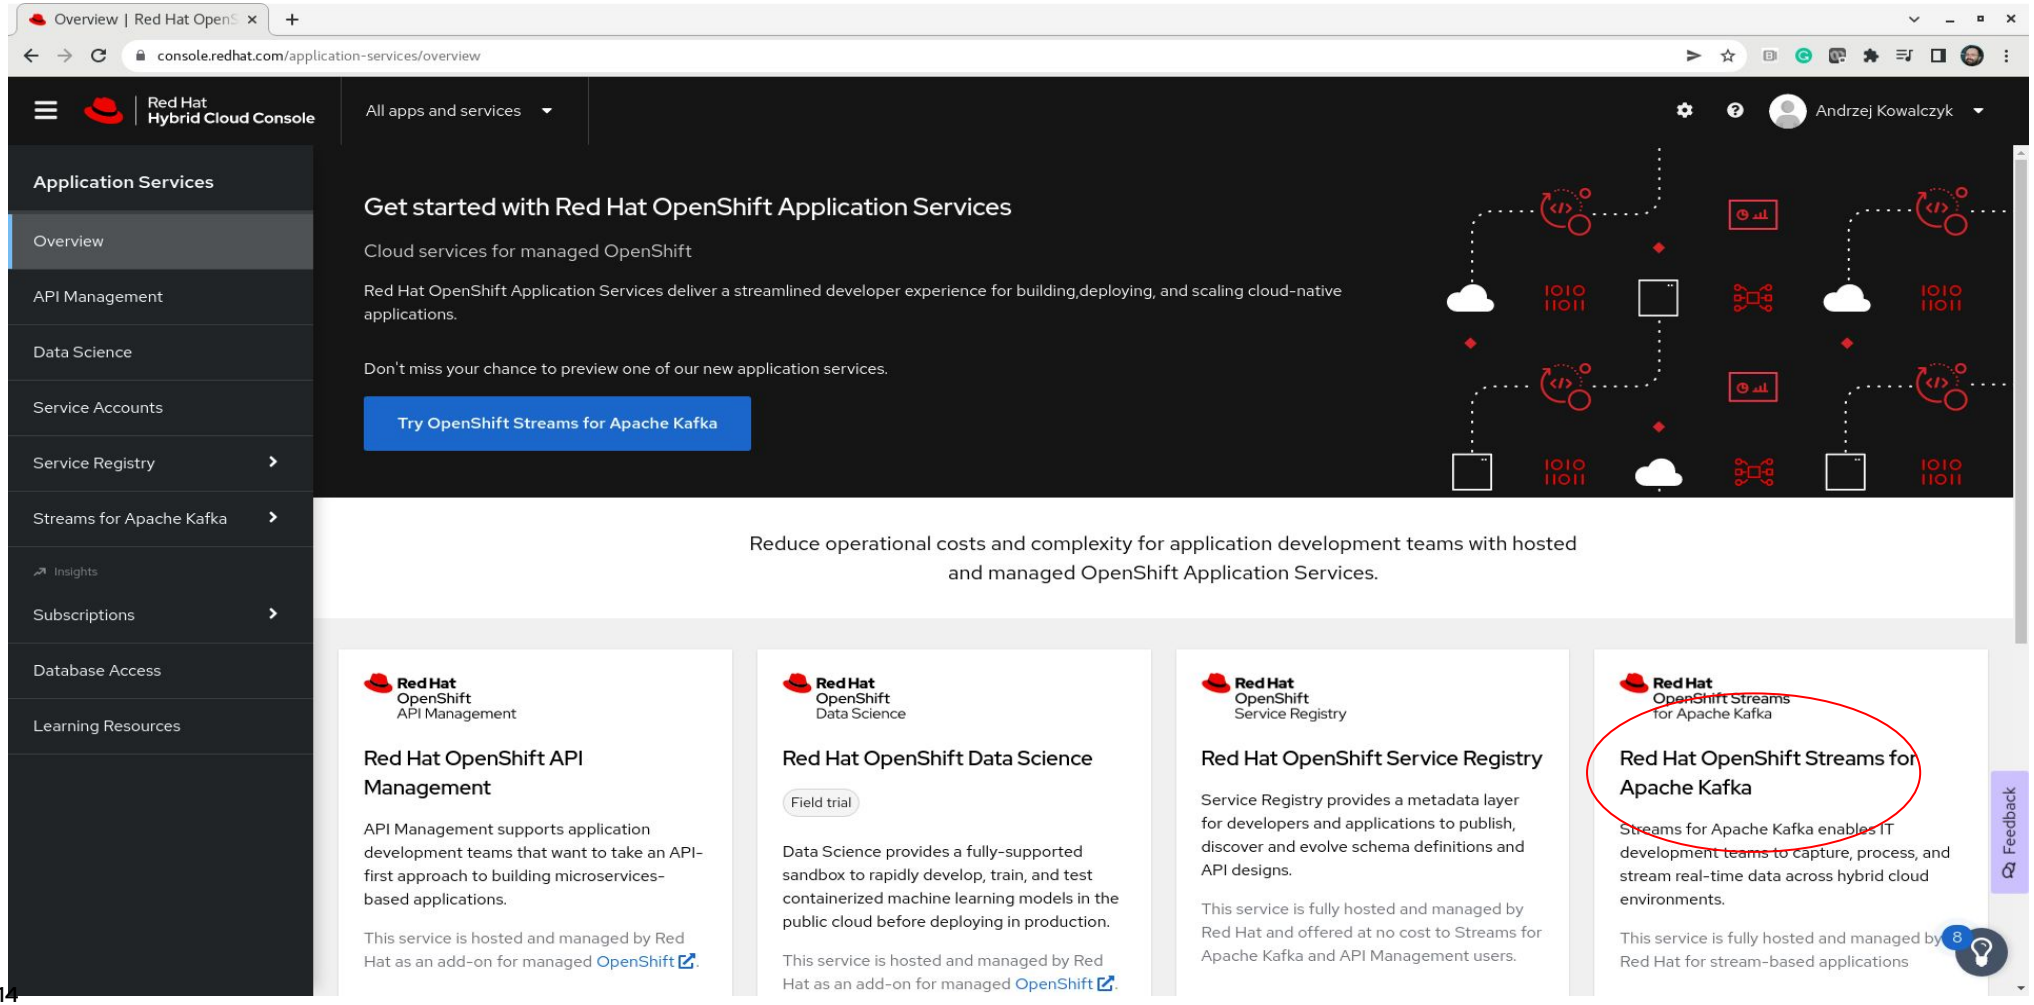

Red Hat | intel.

Można użyć już istniejący broker (twój) lub stworzyć nowy

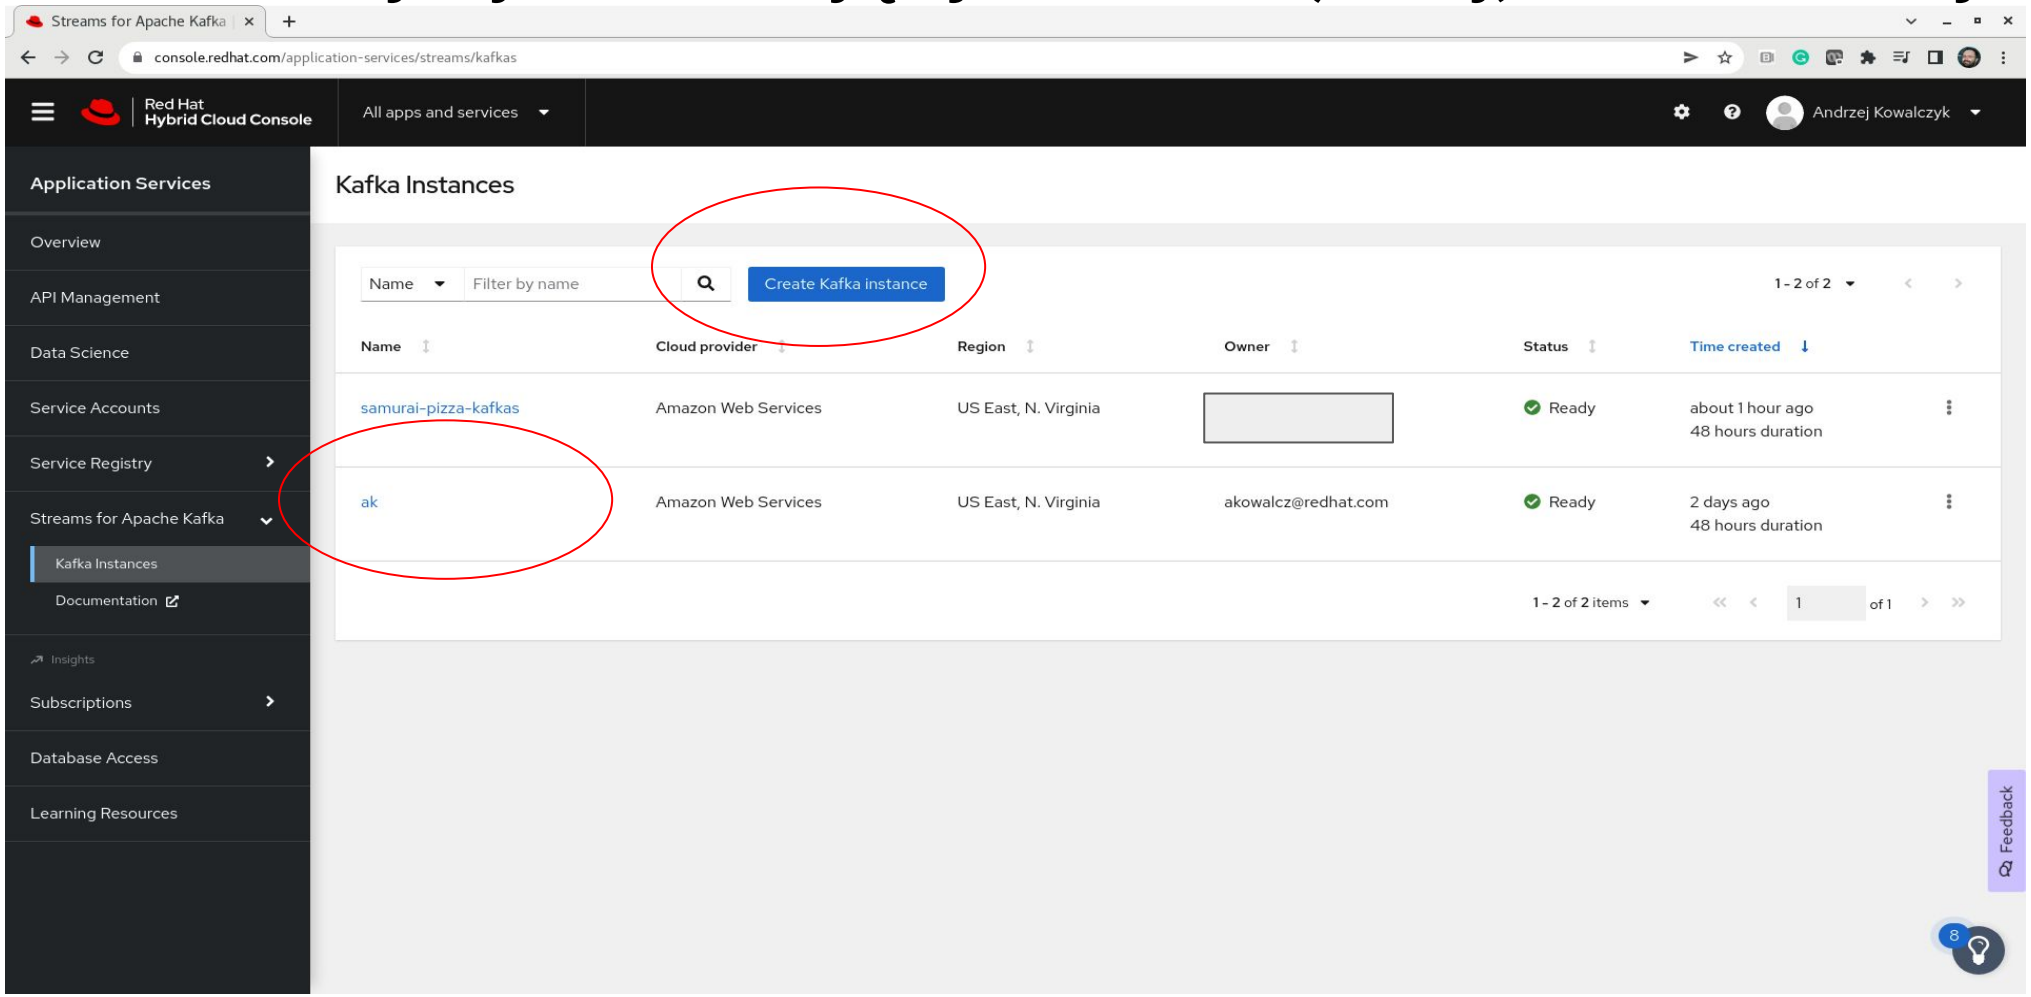

Kopiuję szczegóły połączenia i tworzę service account

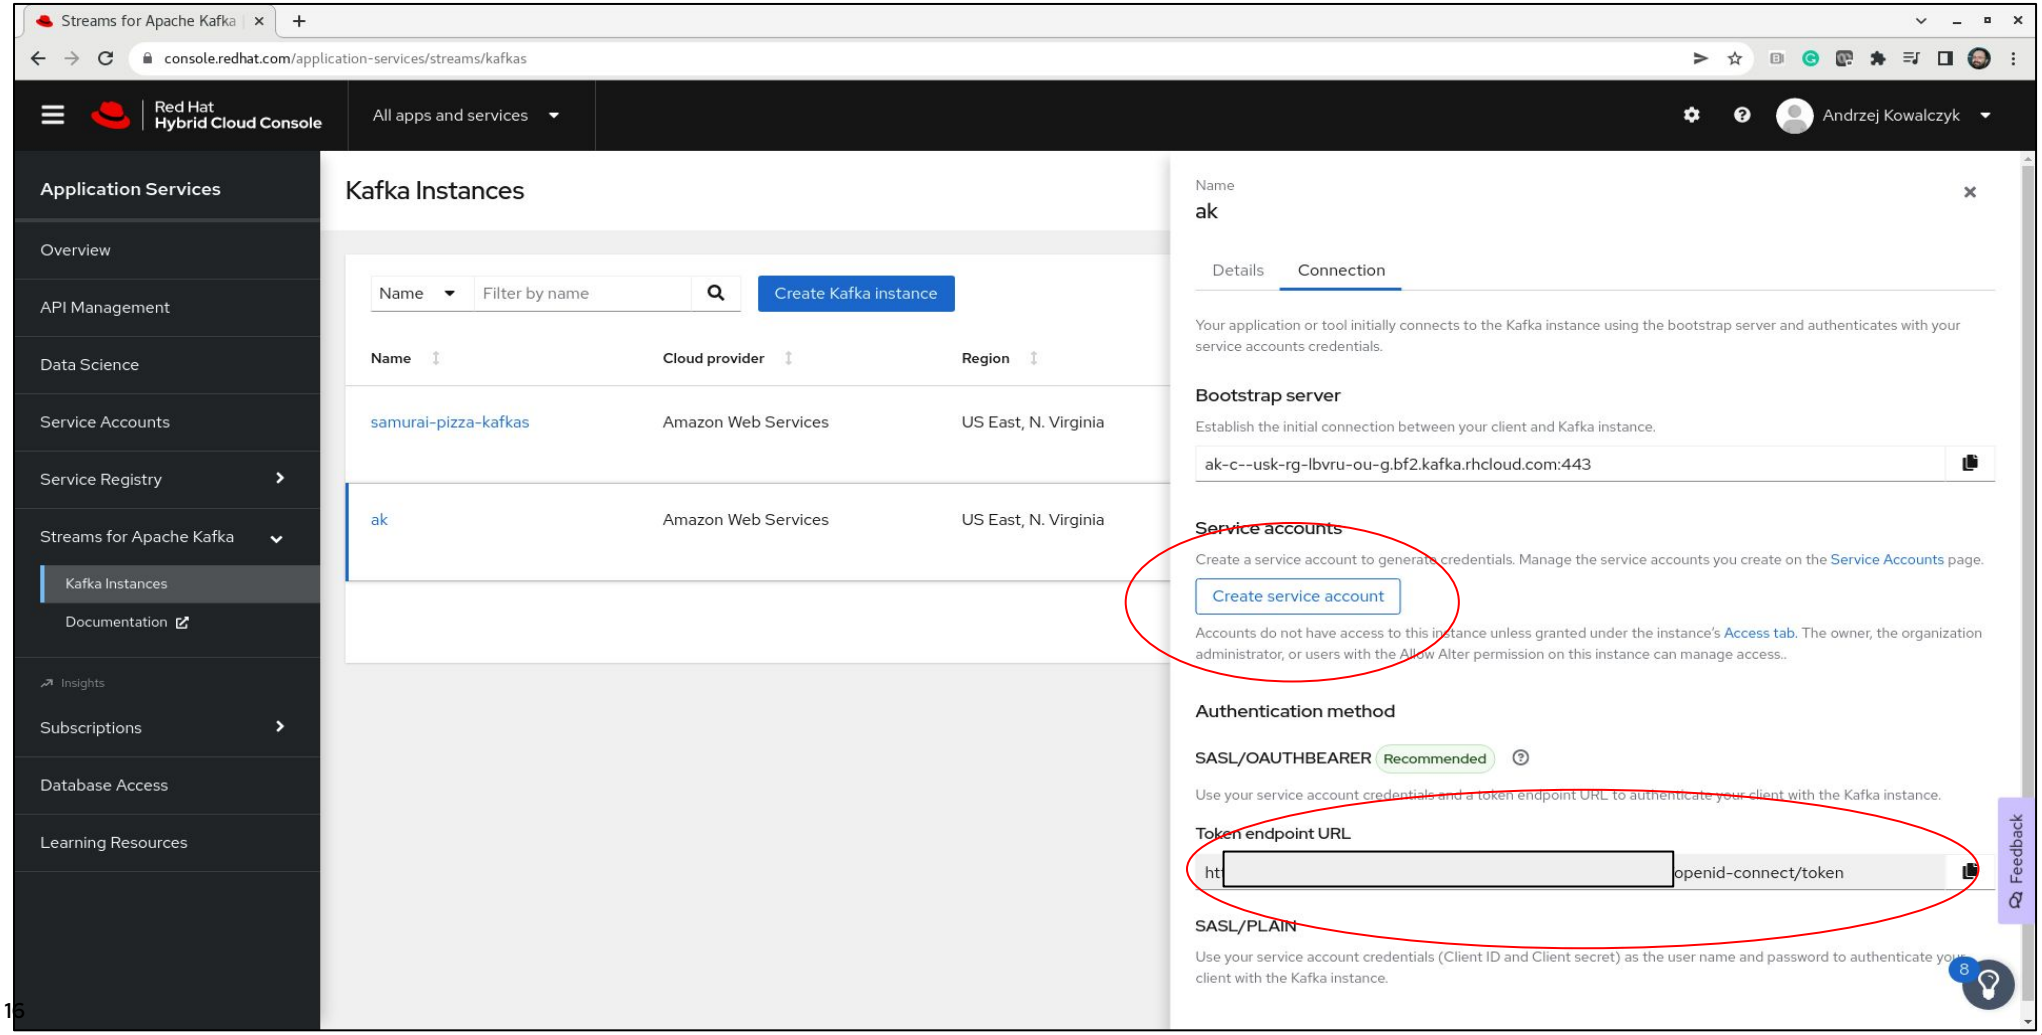

Red Hat | intel.

Tworzę tematy (topics)

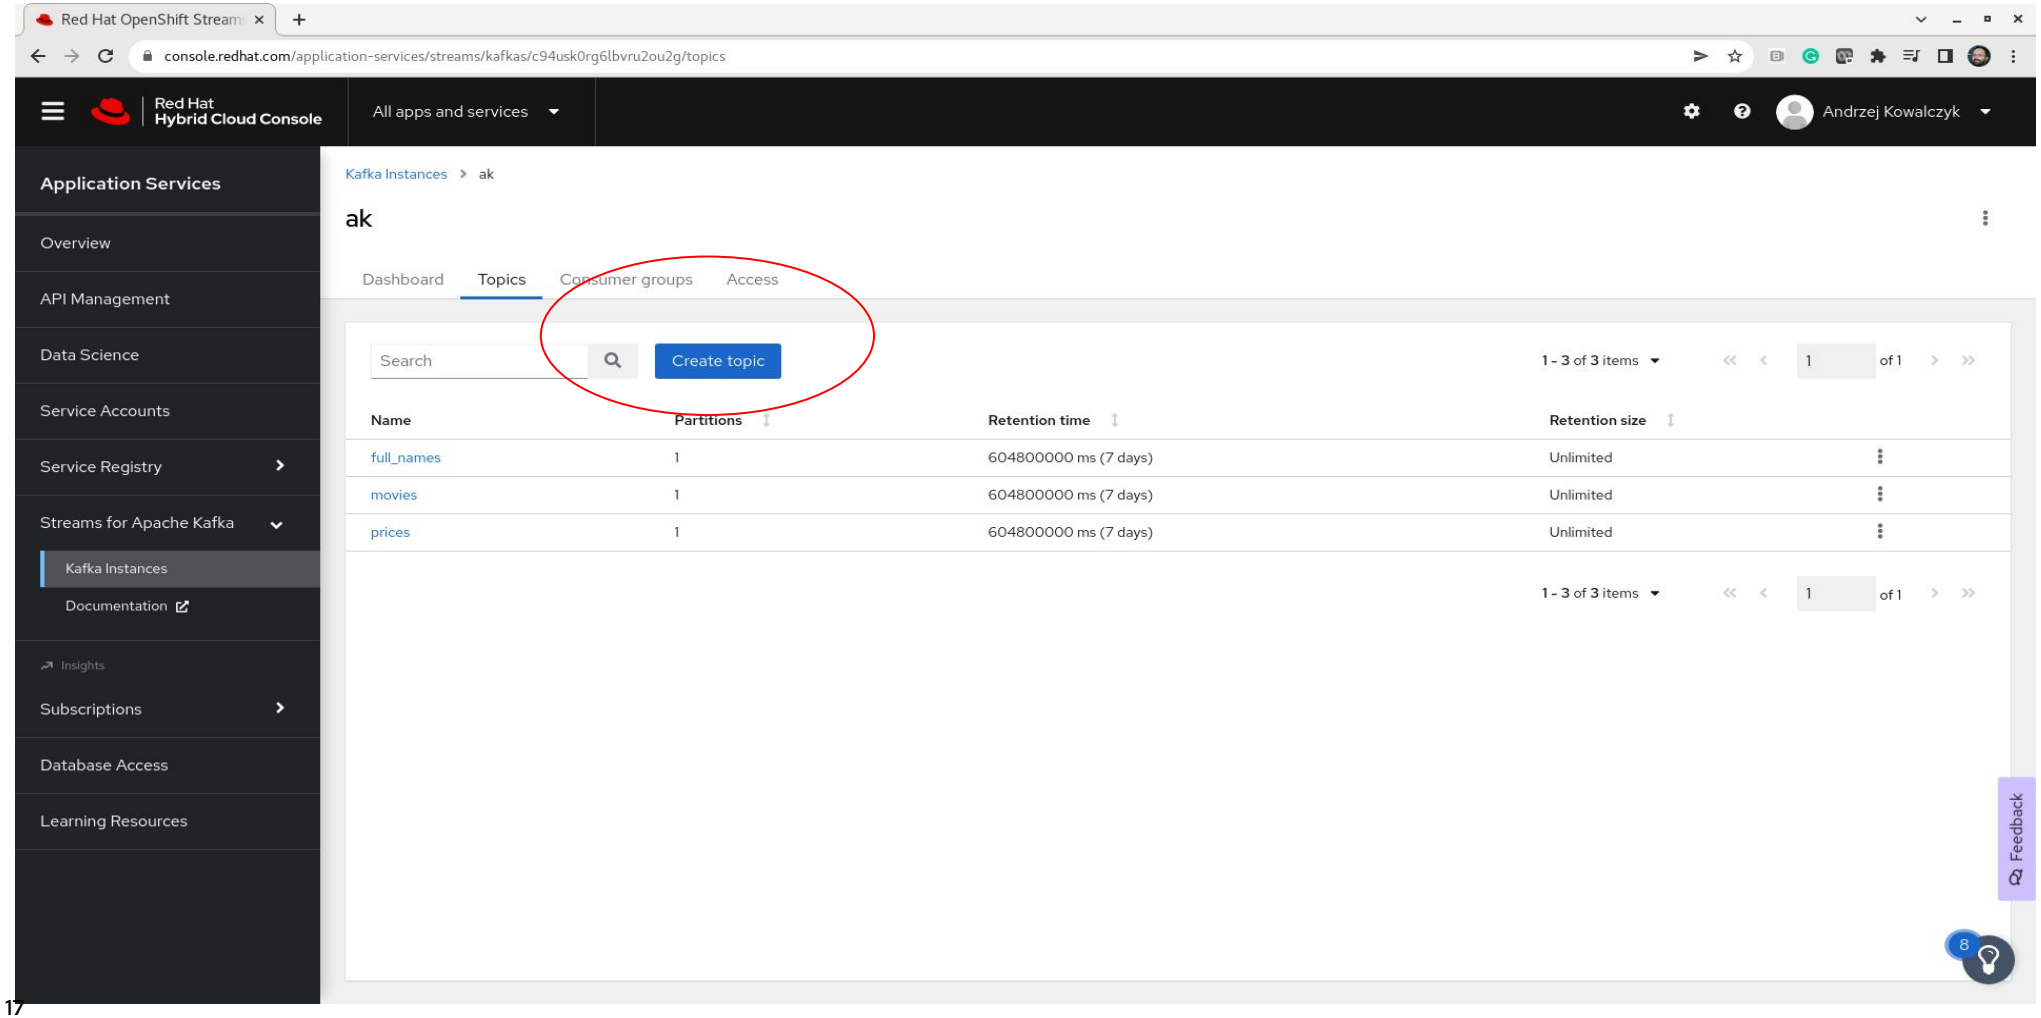

Red Hat | intel.

Dodaję uprawnienia (connect/read/write)

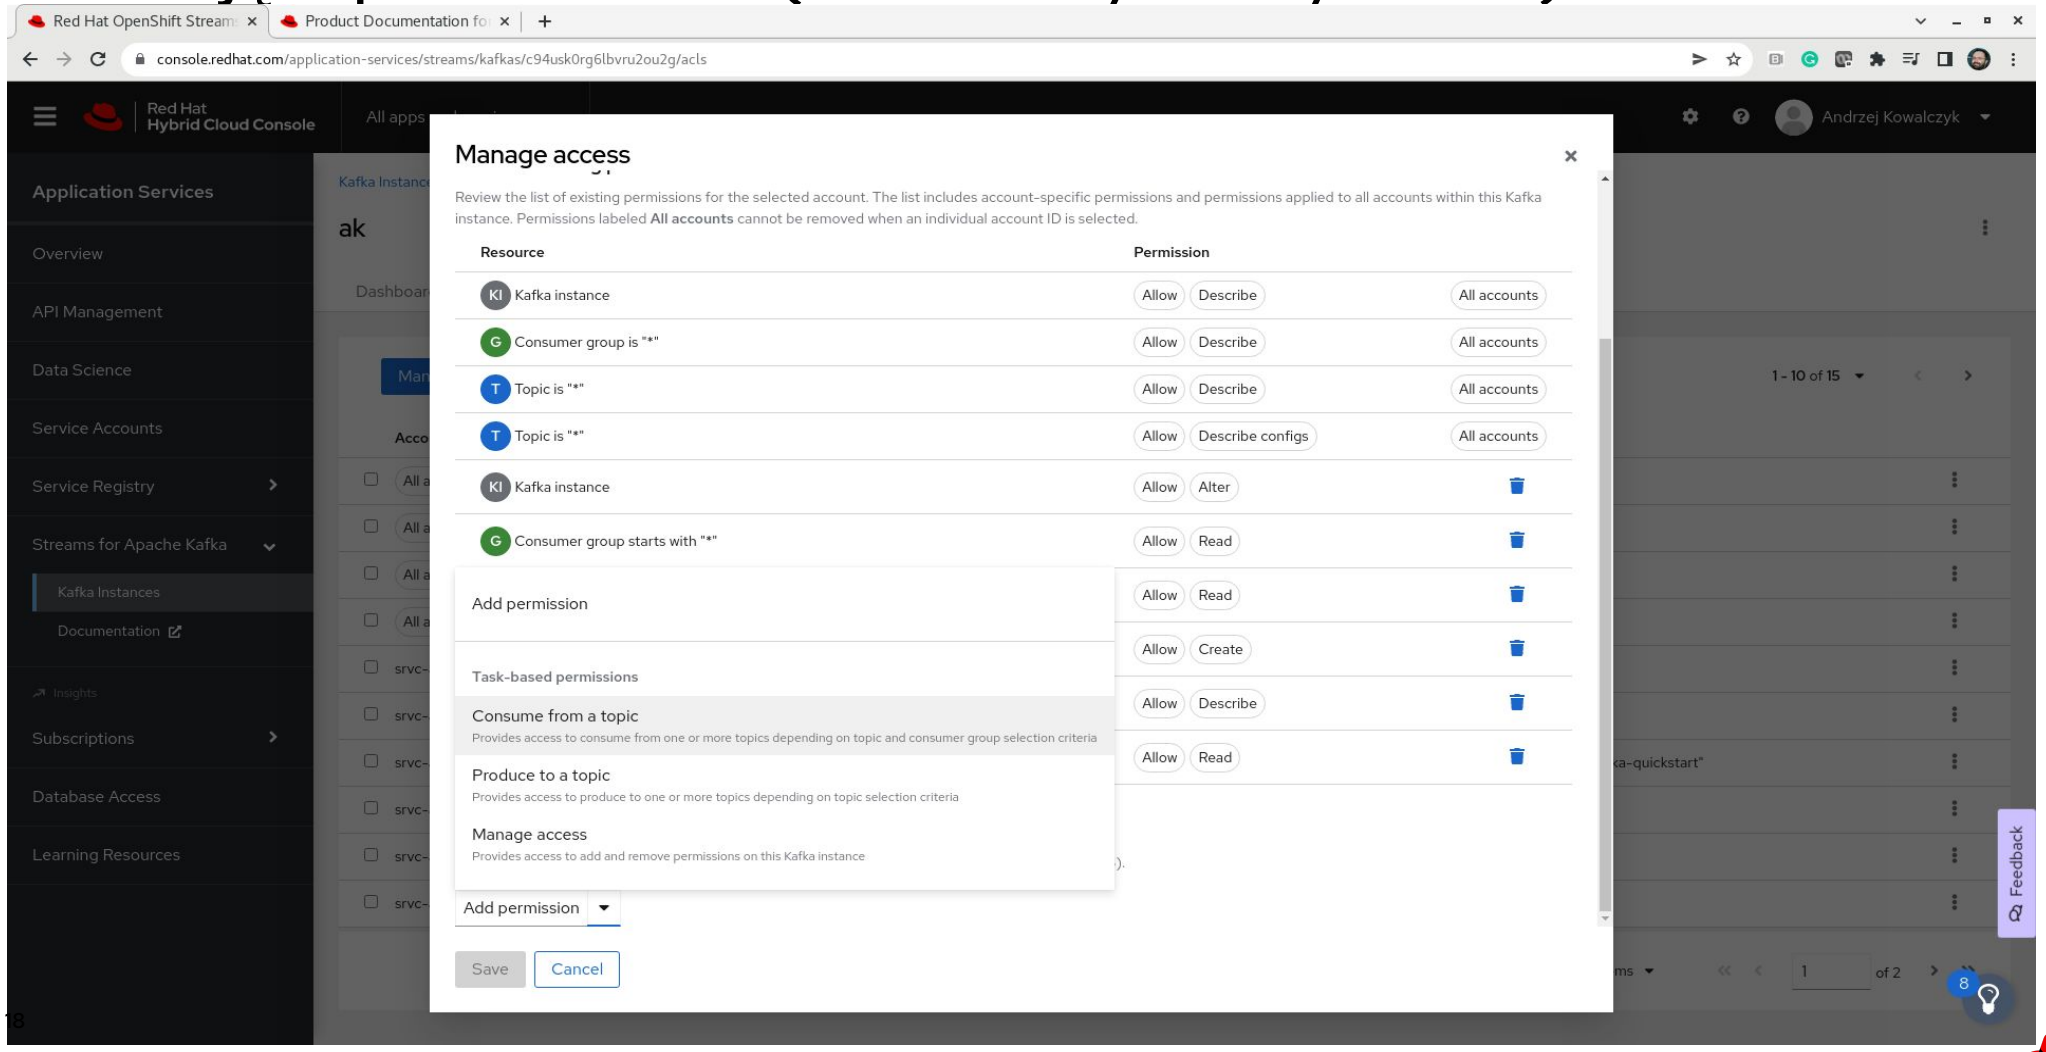

Red Hat | intel.

Jak przetestować czy działa?

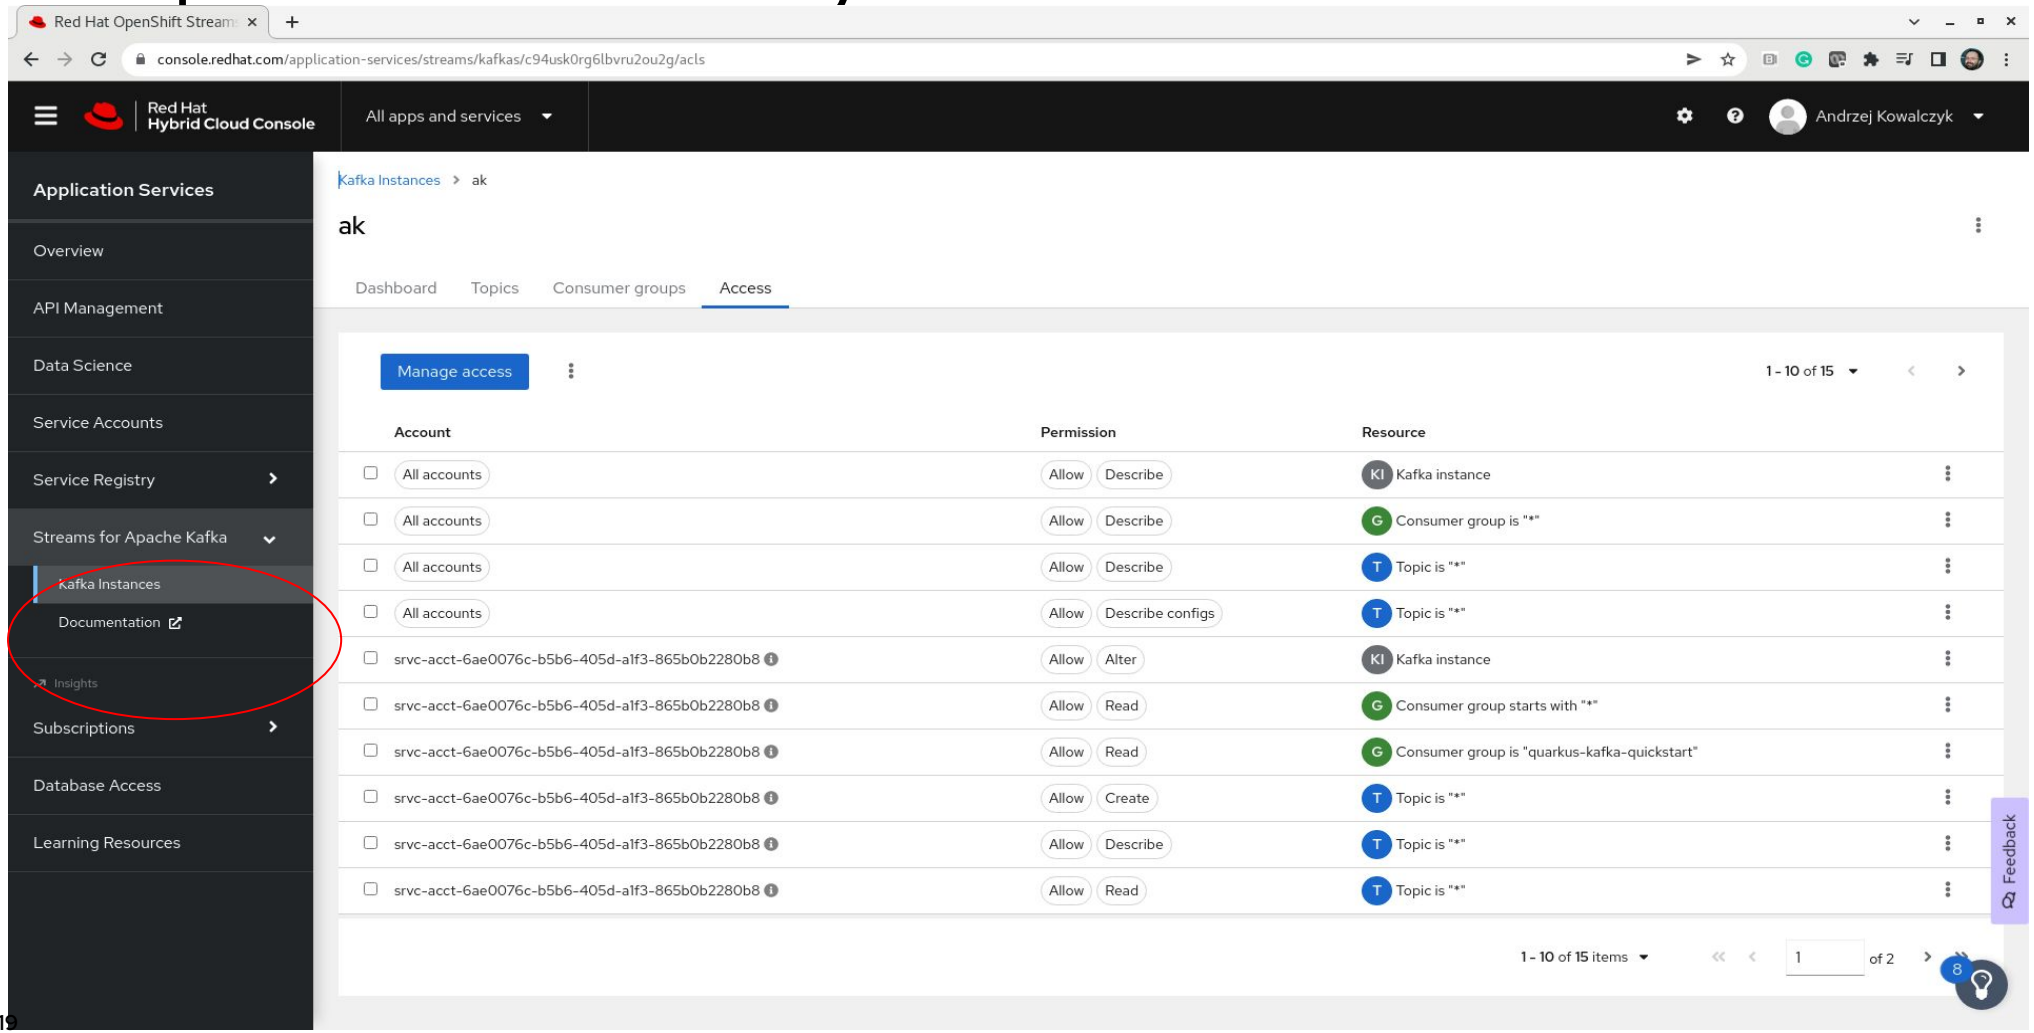

Red Hat | intel.

Najprościej uruchomić przykład z dokumentacji

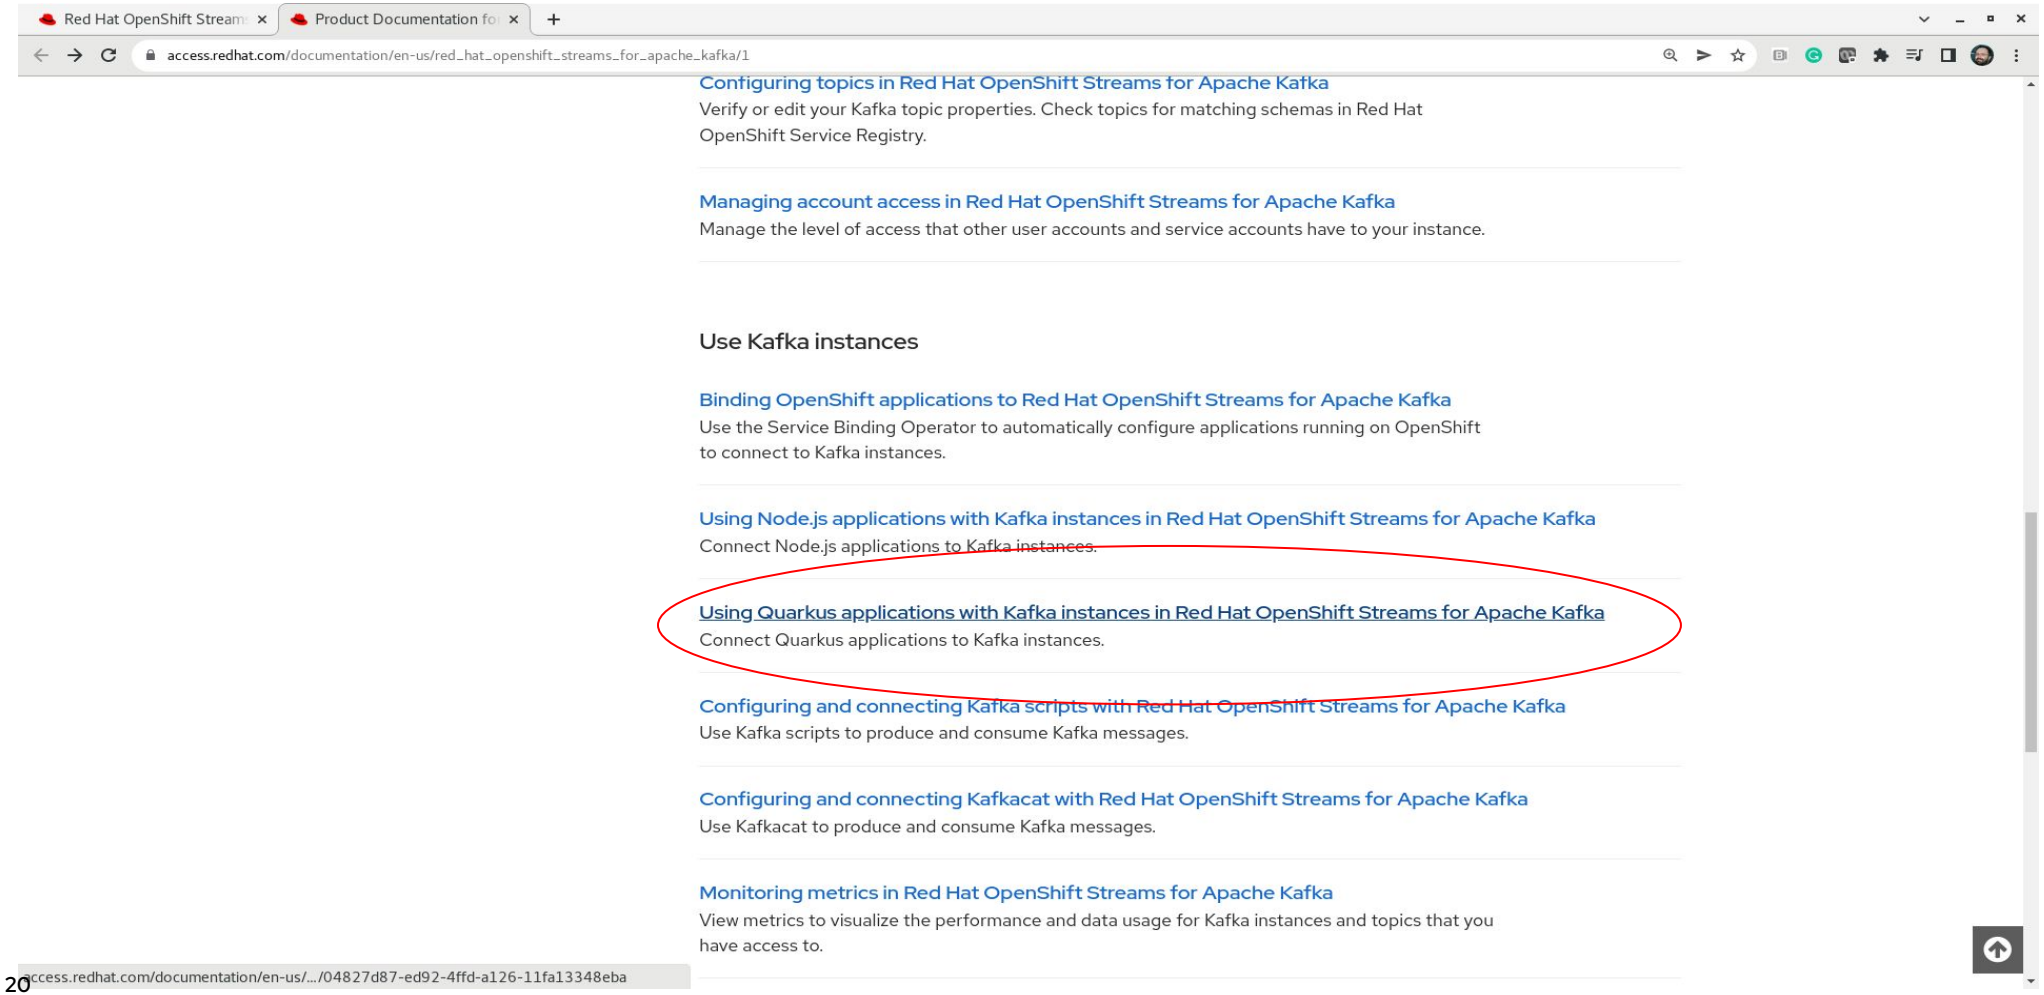

Red Hat | intel

Podążam za dokumentacją...

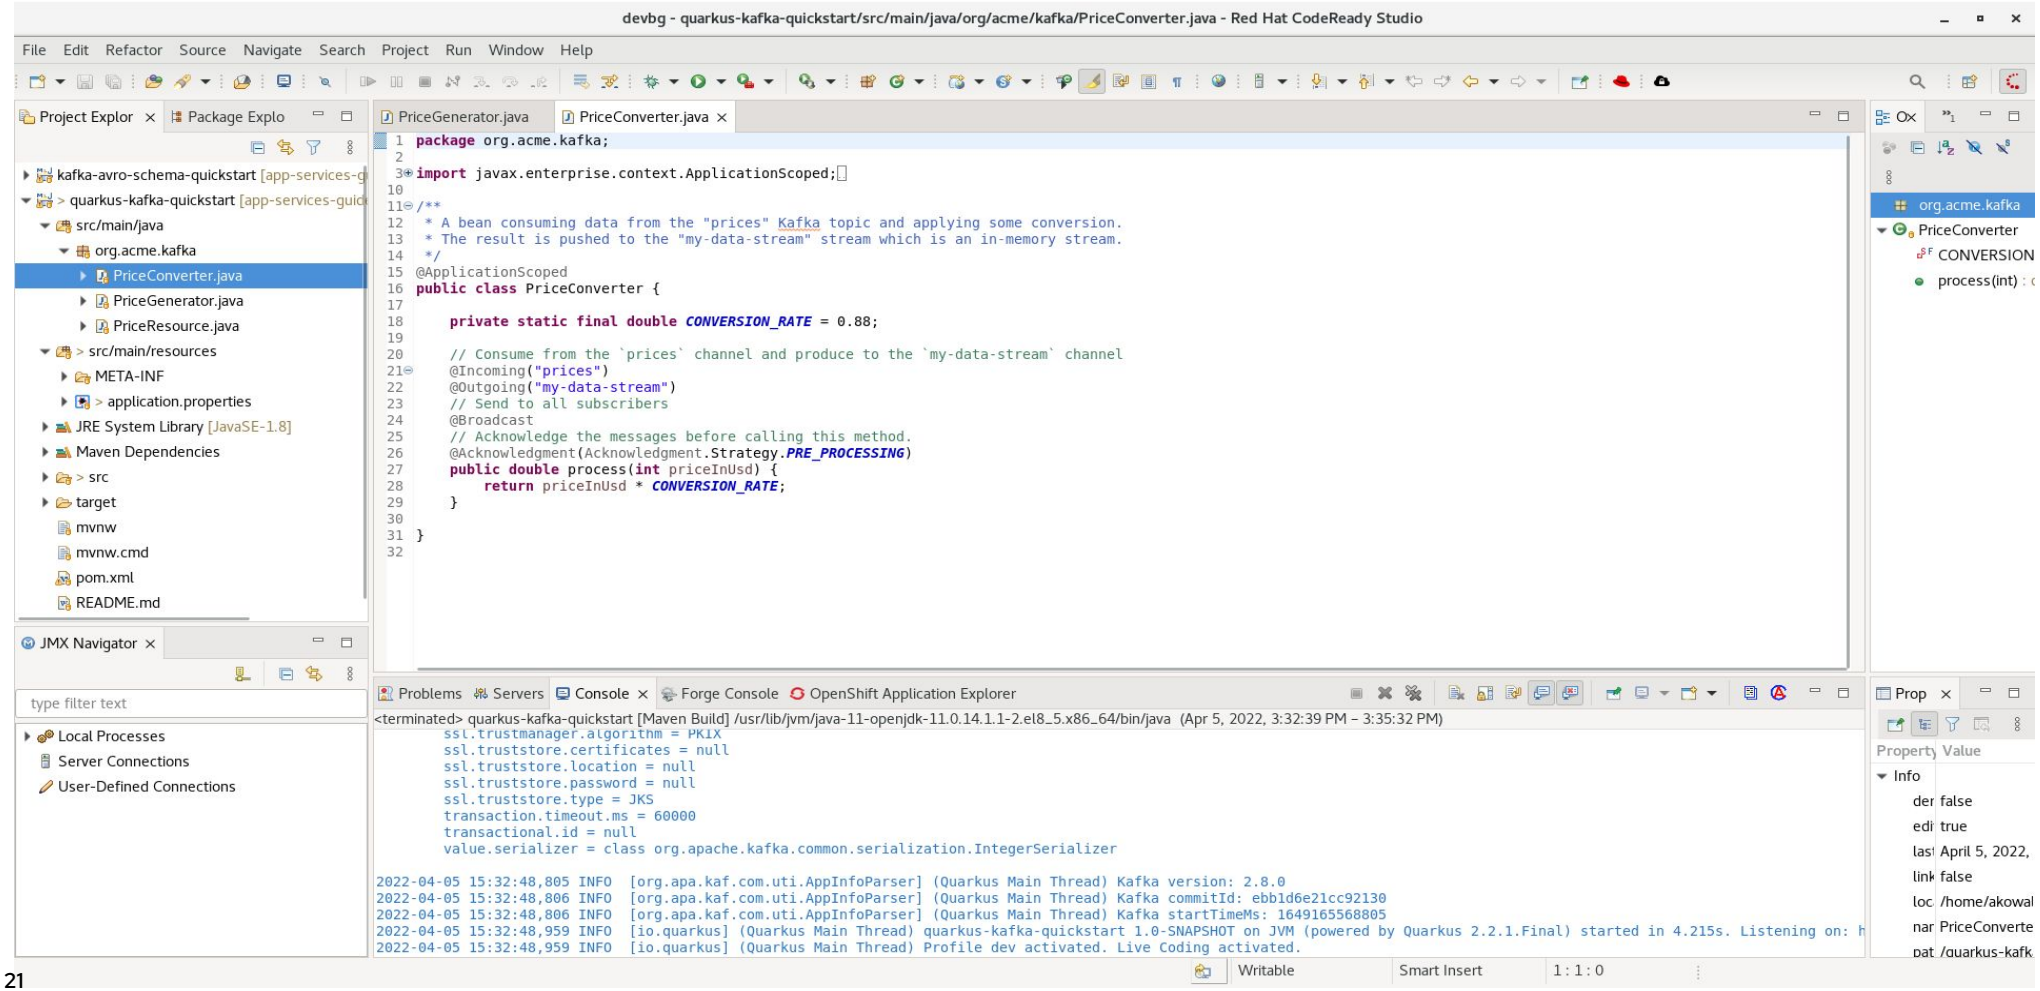

Red Hat | intel.

Hmm... jest błąd

Red Hat | intel.

Tak, zapomniałem utworzyć consumer group i dodać uprawnienia

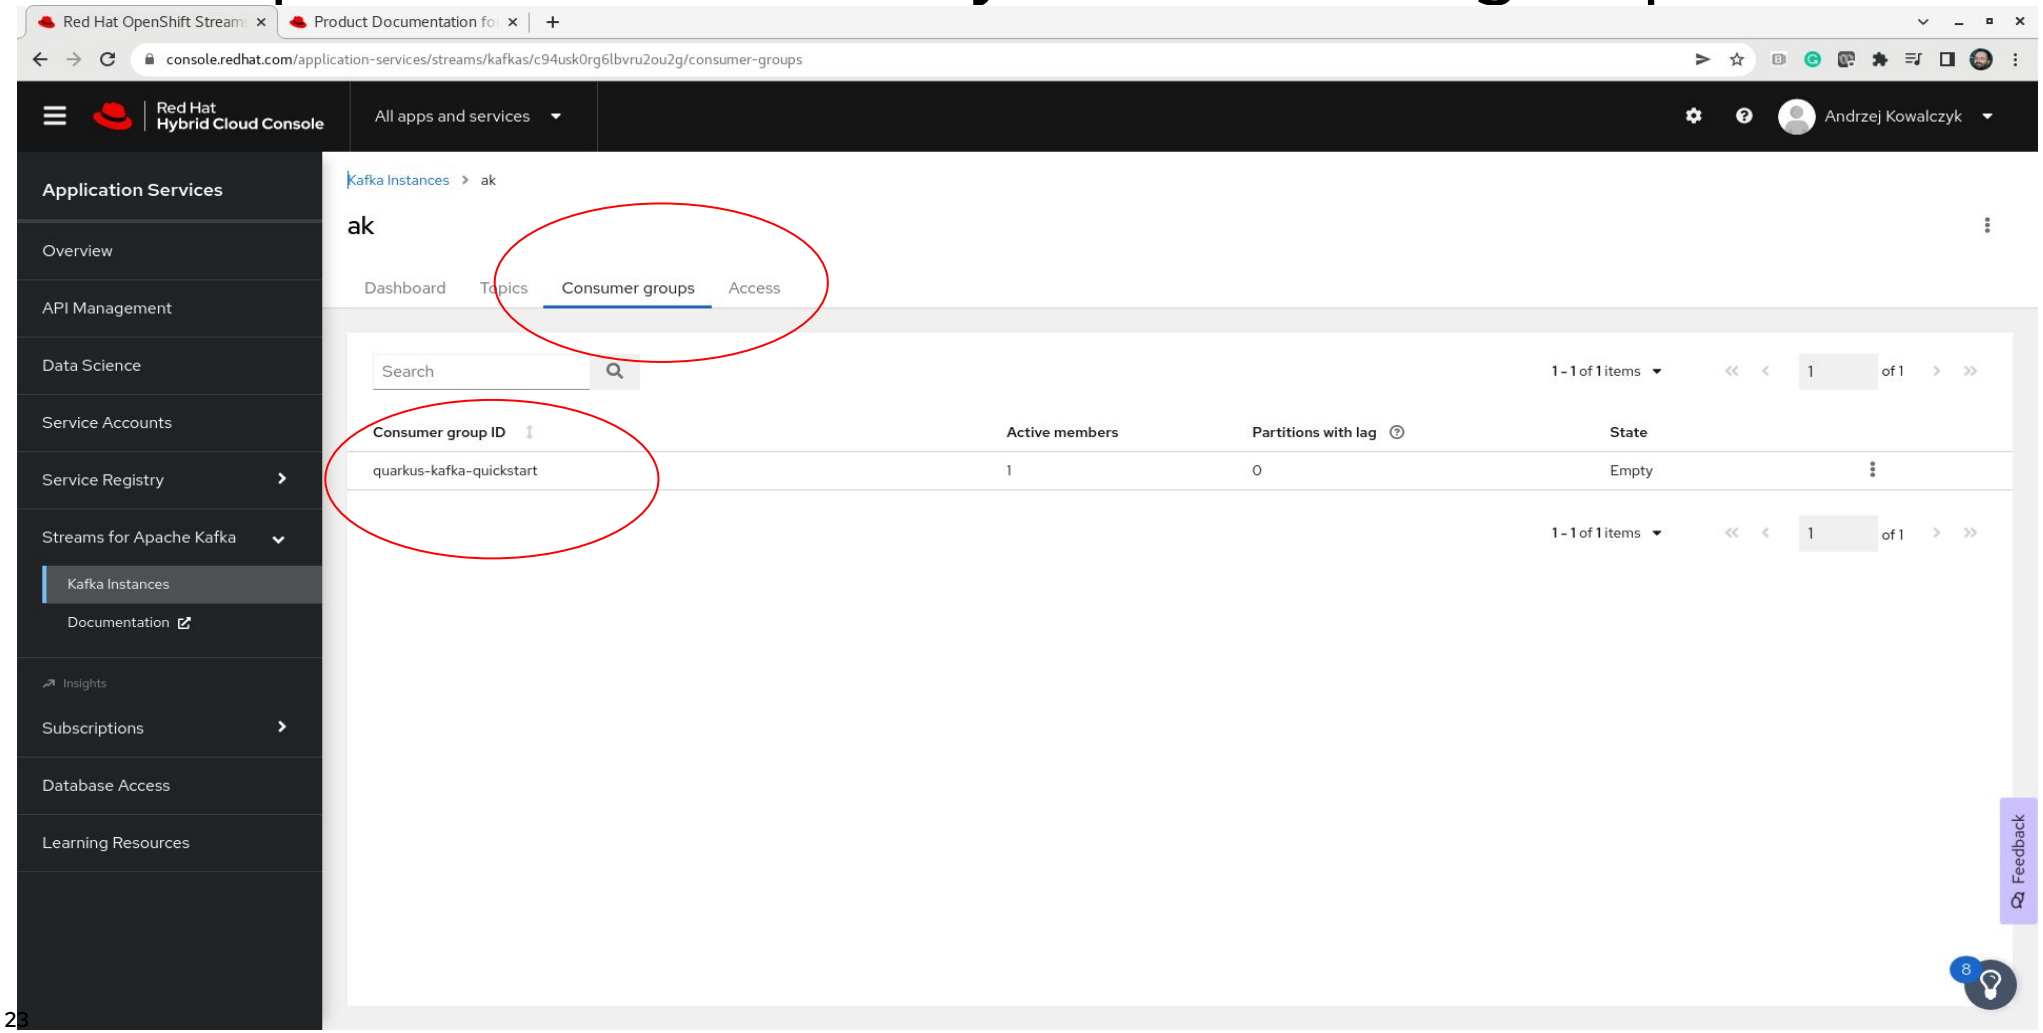

Red Hat | intel.

I w końcu po 20 minutach ... :)

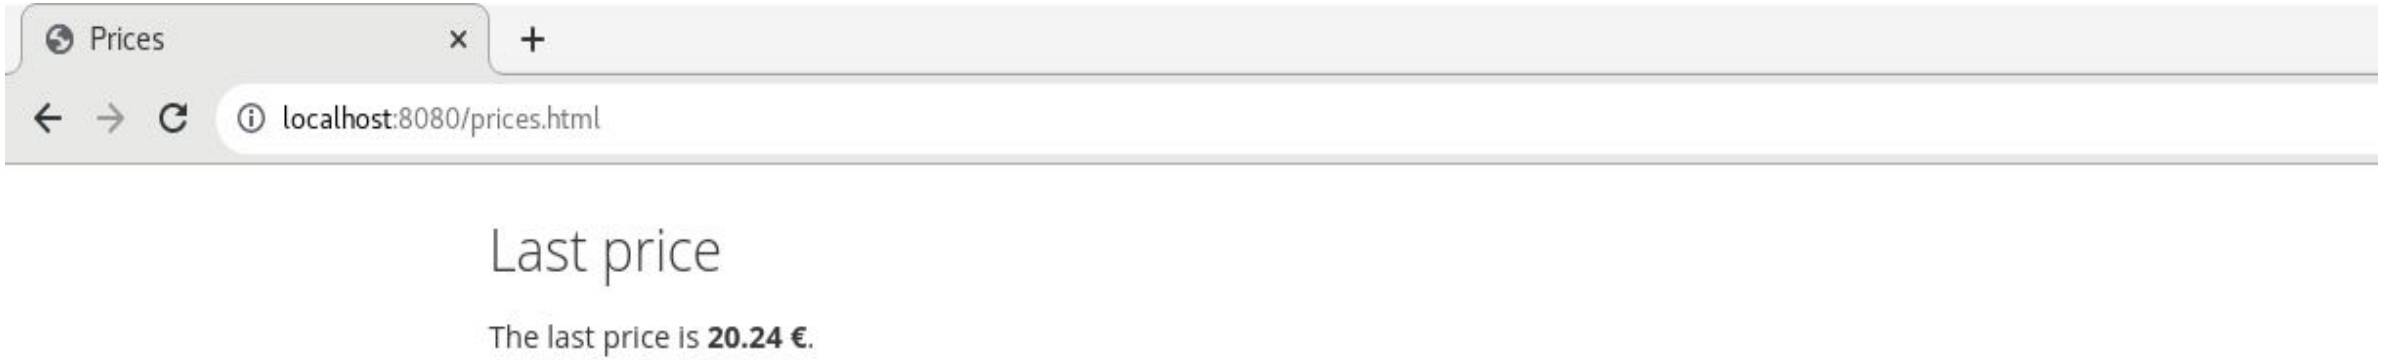

Red Hat | intel

Wszystko trwało 20 minut do uruchomienia kodu.

I ponad godzinę zrobienie zrzutów z ekranu i przygotowanie slajdów :)

Cloud Services Managed by Red Hat

Managed OpenShift + Application Services + Data Services

Full stack management and unified experience

Maximize full value of Red Hat OpenShift

Hybrid cloud flexibility

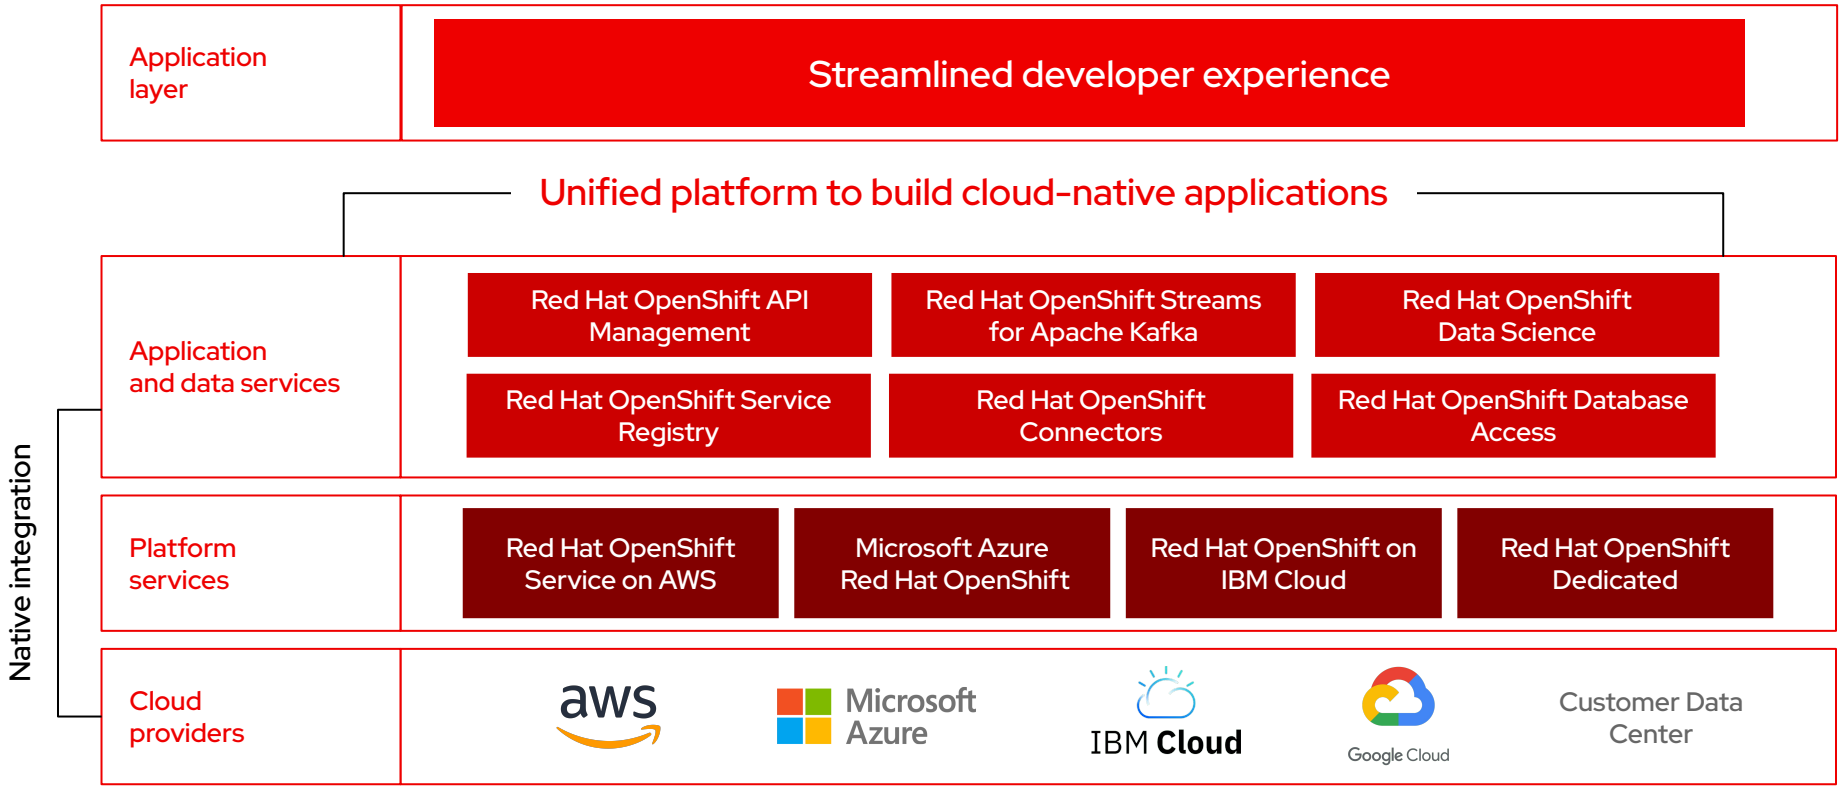

Red Hat | intel

Red Hat OpenShift Streams for Apache Kafka

Try Kafka!

No cost - no strings attached

red.ht/TryKafka

Managed Kafka cluster

- Spin up your own Kafka cluster
- Create your topics and its partitions
- Connect your producers and consumers
- Get started with the quick starts
- Integrate your apps to the service

Time and resource limited

- Access for 48 hours
- Limited number of topics & brokers

Sign-up

- Go to: red.ht/TryKafka
- Create your own Red Hat account
- Sign-in to try the service

Red Hat Summit

Connect

Thank you

 linkedin.com/company/red-hat

 facebook.com/redhatinc

 youtube.com/user/RedHatVideos

 twitter.com/RedHat

Red Hat | intel