Accelerating End-to-End AI/ML with Red Hat

How Data Scientists can increase AI/ML scalability, agility, and deliver intelligent applications

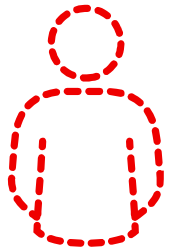Steve Meldrum

Senior Solutions Architect

September 1, 2020

Red Hat

All I want to do is get my model into production!

Talent shortage

Lack of key skills make it difficult to find and secure talent to maintain operations.

Readily usable data lacking

Lots of data is collected, but finding and preparing the right data is difficult.

Lack of collaboration across teams

Unable to implement quickly due to slow, manual and siloed operations.

Unavailability of infrastructure & software

No rapid availability of infrastructure and software tools slows data scientists and developers

Red Hat Influence to help solve challenges

Creating better technology the open source way
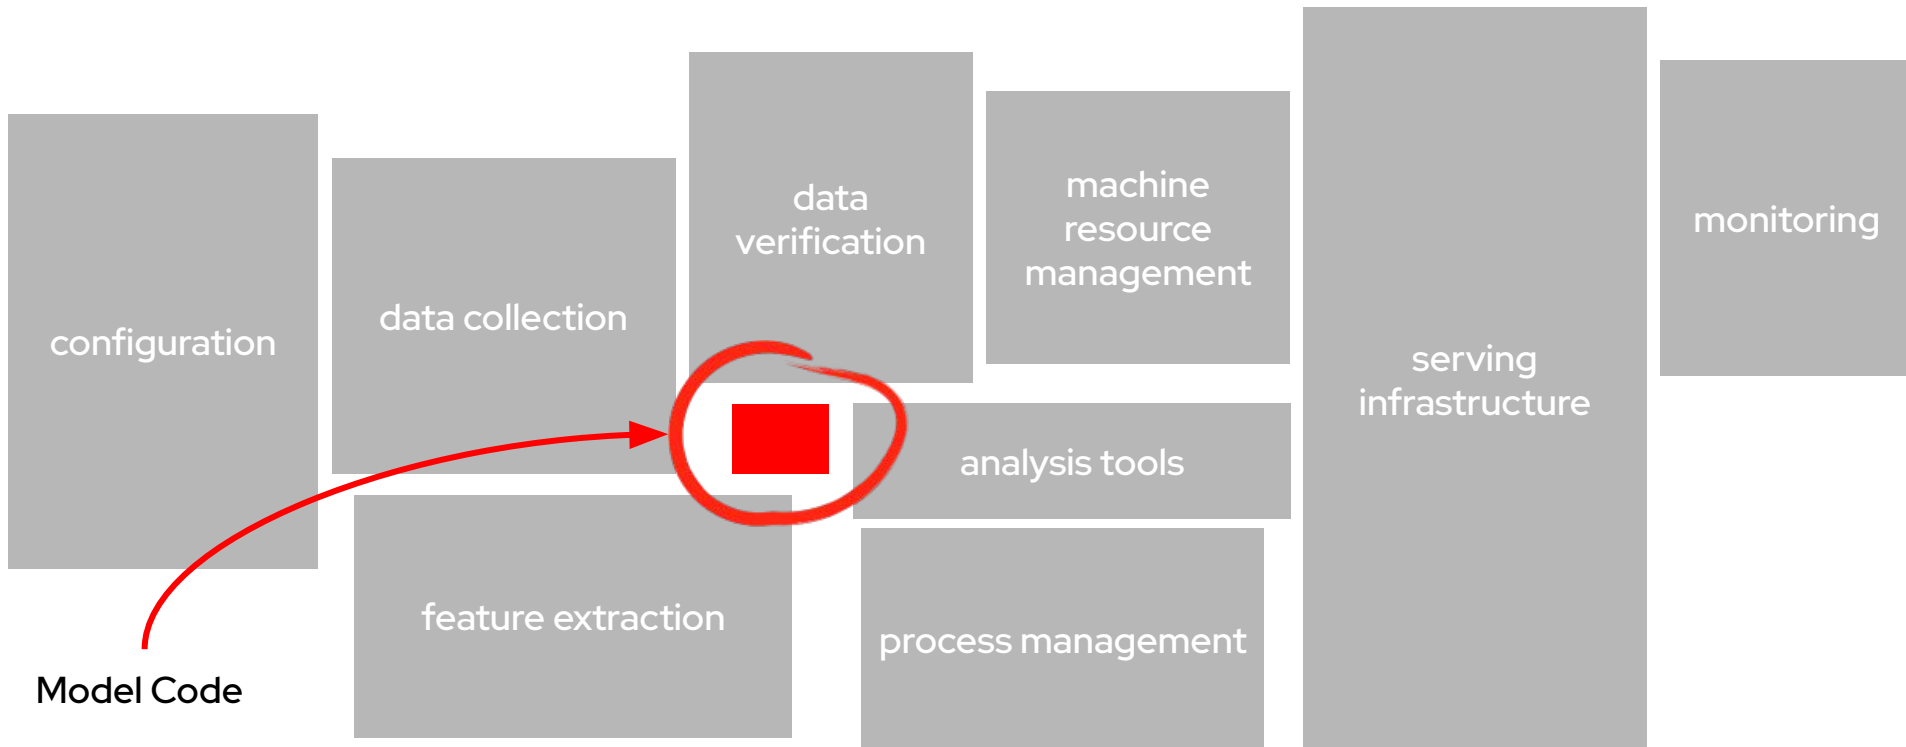

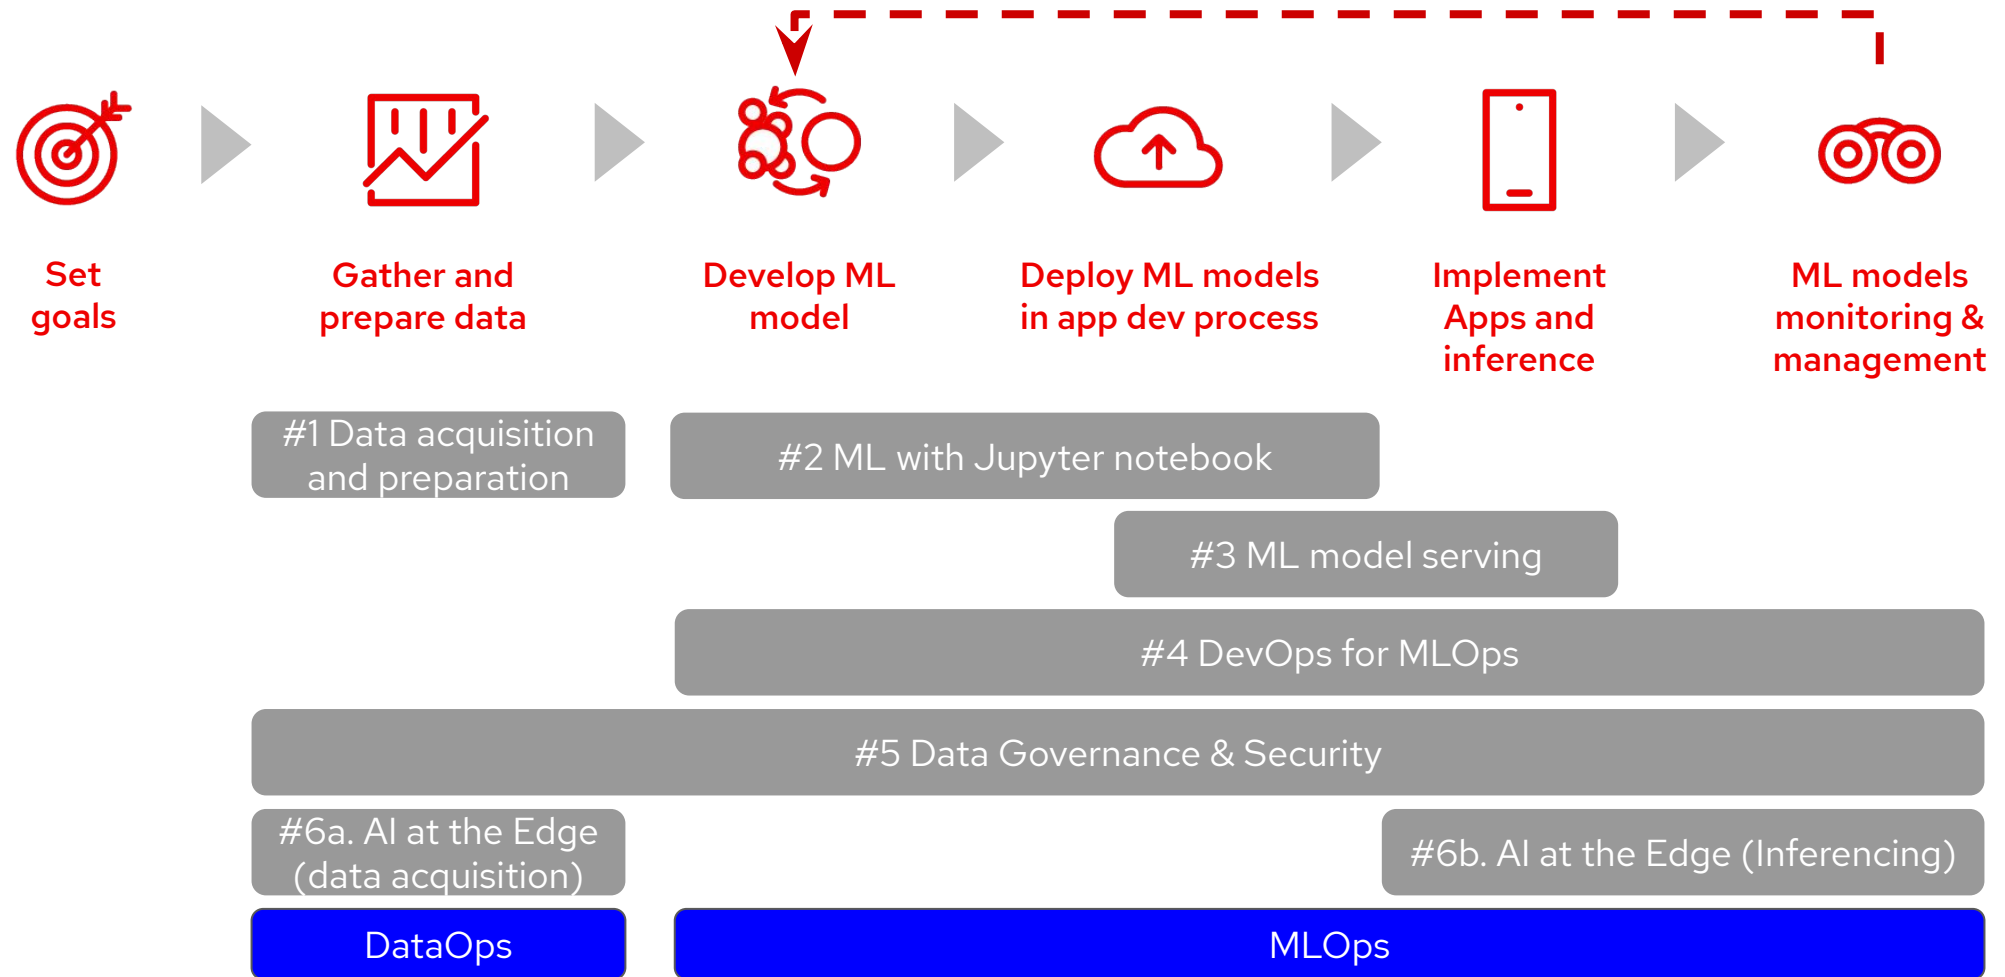| IT optimization | Hybrid cloud infrastructure | Cloud-native development | Agile integration | Management and automation |

Intelligent applications are distributed systems

The model is a tiny part, but let's focus on it

configuration

data collection

data verification

machine resource management

monitoring

analysis tools

serving infrastructure
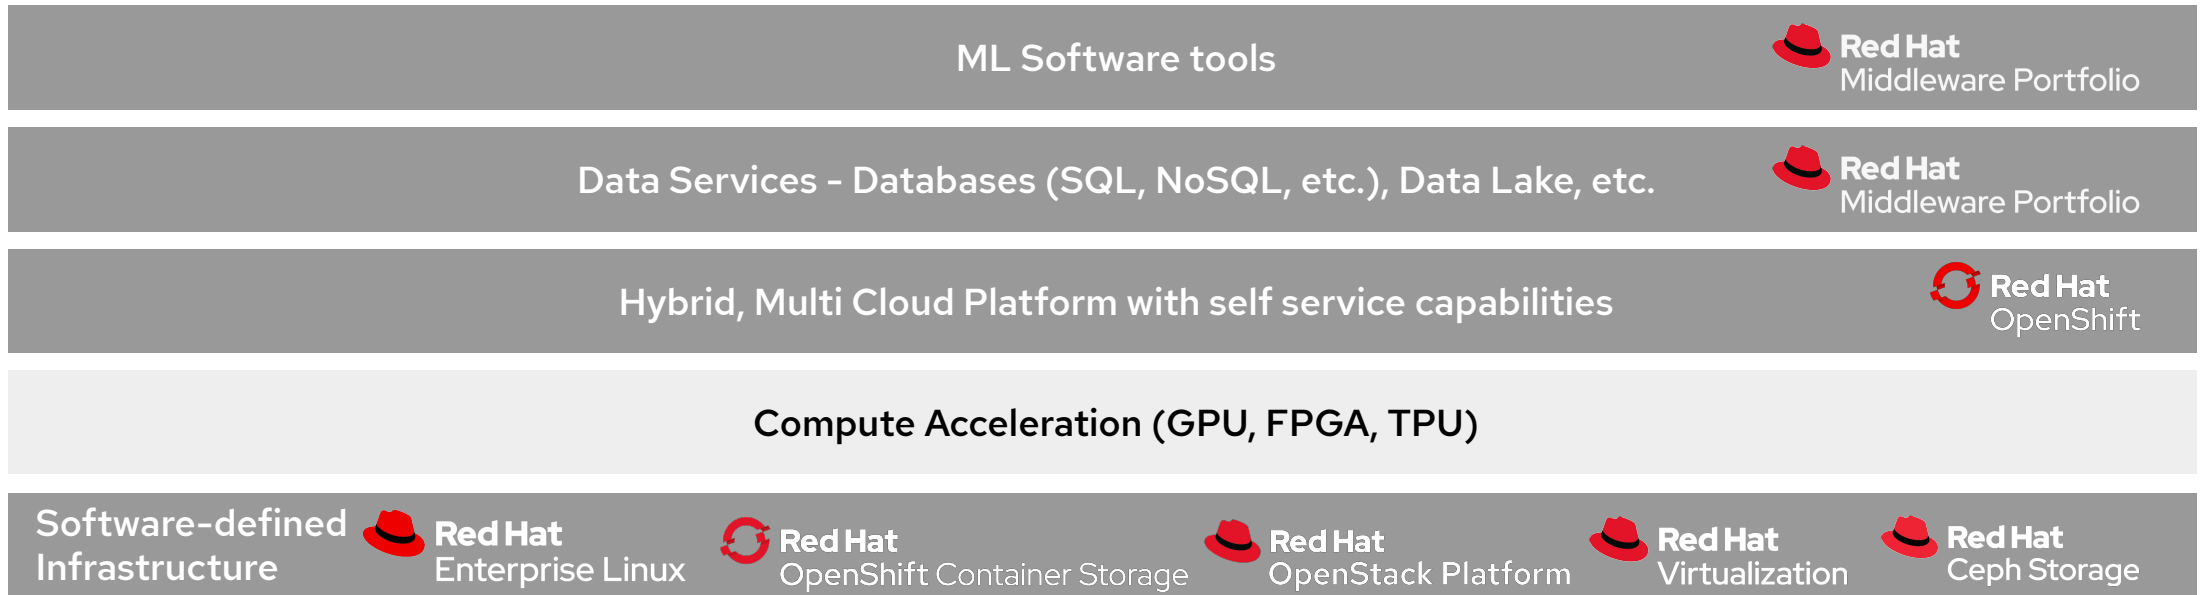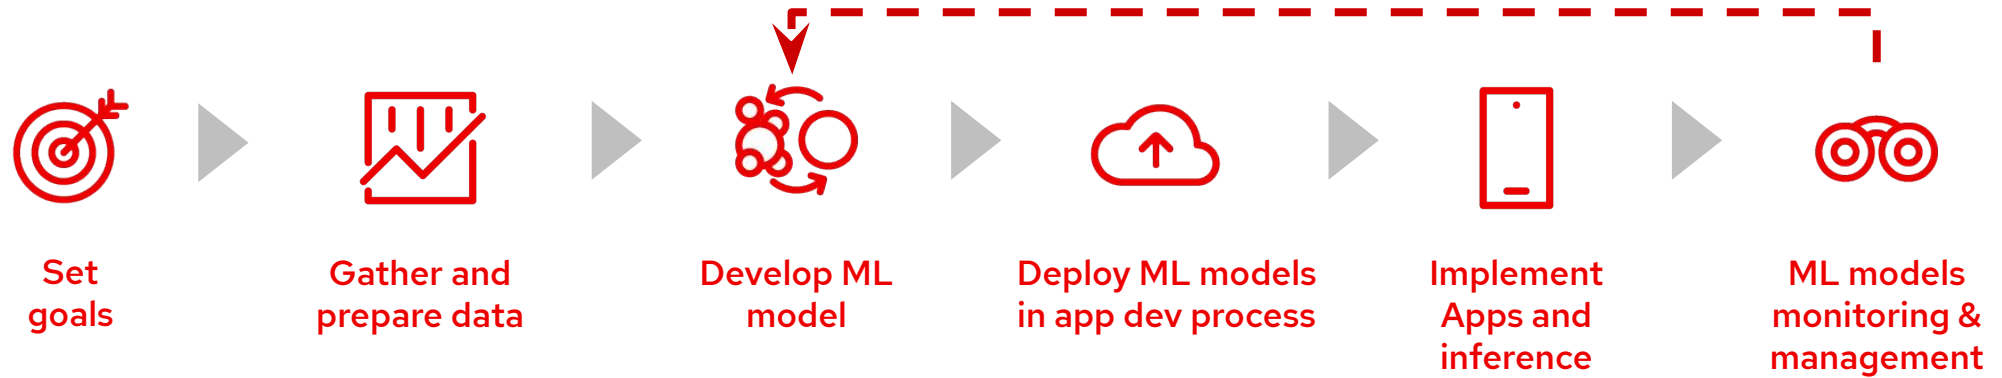
feature extraction

process management

Model Code

(Adapted from Sculley et al., "Hidden Technical Debt in Machine Learning Systems." NIPS 2015)

Red Hat

Accelerate AI/ML initiatives and rapidly move models into production

Set goals

Gather and prepare data

Develop ML model

Deploy ML models in app dev process

Implement Apps and inference

ML models monitoring & management

#1 Data acquisition and preparation

#2 ML with Jupyter notebook

#3 ML model serving

#4 DevOps for MLOps

#5 Data Governance & Security

#6a. AI at the Edge (data acquisition)

#6b. AI at the Edge (Inferencing)

DataOps

MLOps

Red Hat

Red Hat portfolio helps complete the ML architecture

| Set goals | | Gather and prepare data | | Develop ML model | | Deploy ML models in app dev process | | Implement Apps and inference | | ML models monitoring & management |

ML Software tools — Red Hat Middleware Portfolio

Data Services - Databases (SQL, NoSQL, etc.), Data Lake, etc. — Red Hat Middleware Portfolio

Hybrid, Multi Cloud Platform with self service capabilities — Red Hat OpenShift

Compute Acceleration (GPU, FPGA, TPU)

Software-defined Infrastructure — Red Hat Enterprise Linux · Red Hat OpenShift Container Storage · Red Hat OpenStack Platform · Red Hat Virtualization · Red Hat Ceph Storage

Red Hat

Future uses, same architecture... bring the data science to your data!

Upstream

Midstream

Downstream

Exploration

Leak Detection

Refinery Incident Monitoring

Drilling Planning

Logistics
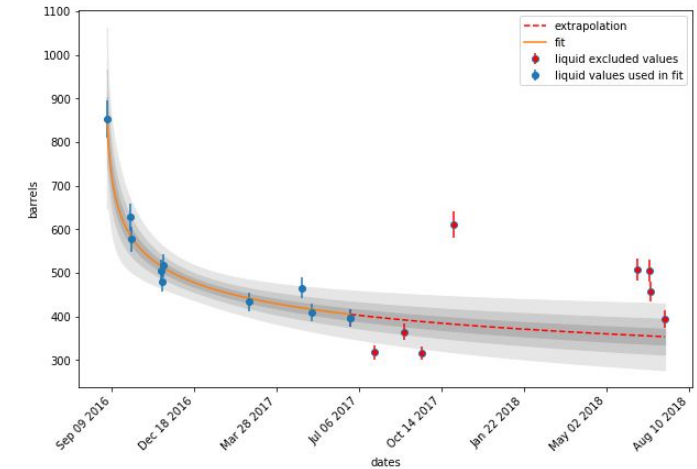
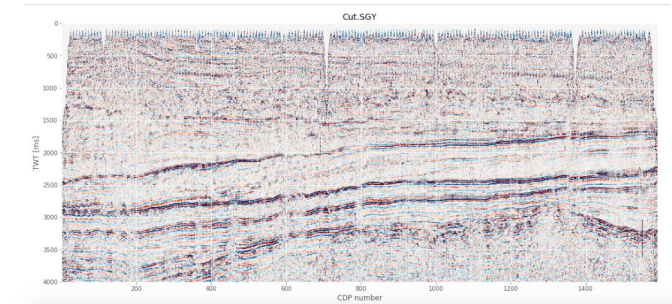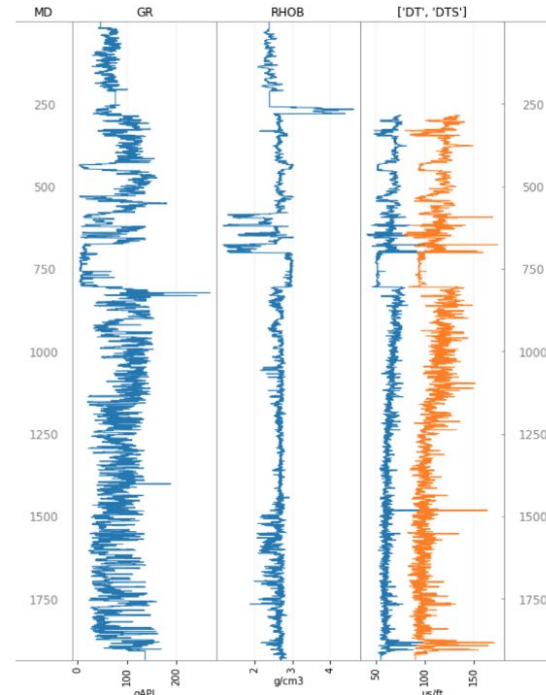Synthetic Well Log Generation

Flare Monitoring

Data Qualify and Preparation

Safety Incident Response

Computer Vision

Red Hat

Upstream Data Science on Red Hat OpenShift

Unique performance computing requirements for Artificial Intelligence, Machine Learning, Neural Networks and GPUs
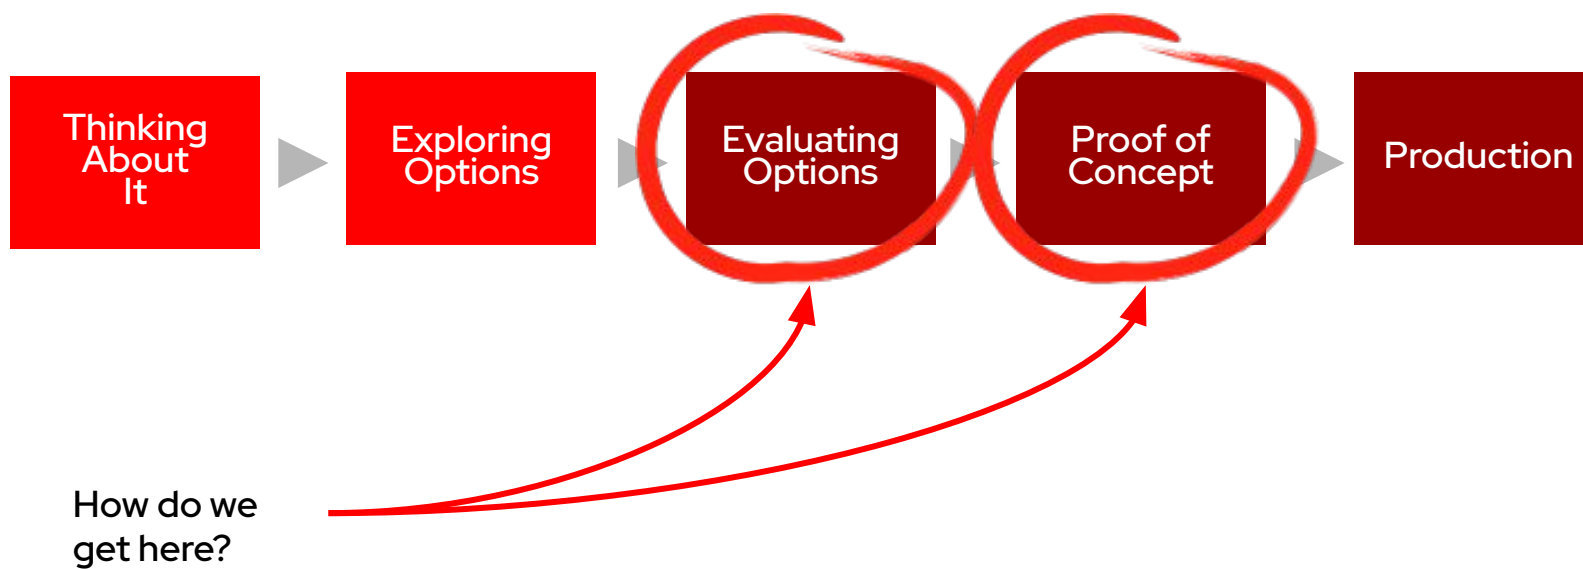
Efficiently leverage your data with the Open Subsurface Data Universe

- Pre-built data integrations
- Get to your data easily
- On Prem or in the cloud
- Filters and APIs
- Fine-grained access to data
- Terms of use

Red Hat

Take the next steps to AI/ML productivity with Red Hat!

Thinking About It ▶ Exploring Options ▶ Evaluating Options ▶ Proof of Concept ▶ Production

How do we get here?

Red Hat

Thank you

Red Hat is the world's leading provider of enterprise open source software solutions. Award-winning support, training, and consulting services make Red Hat a trusted adviser to the Fortune 500.

Steve Meldrum

Senior Solutions Architect

smeldrum@redhat.com

in   linkedin.com/company/red-hat

▶   youtube.com/user/RedHatVideos

f   facebook.com/redhatinc

🐦   twitter.com/RedHat

Red Hat