Open for Innovation

KNIME

Welcome to
Sharing and Deploying Data Science with KNIME Server

Going live at:

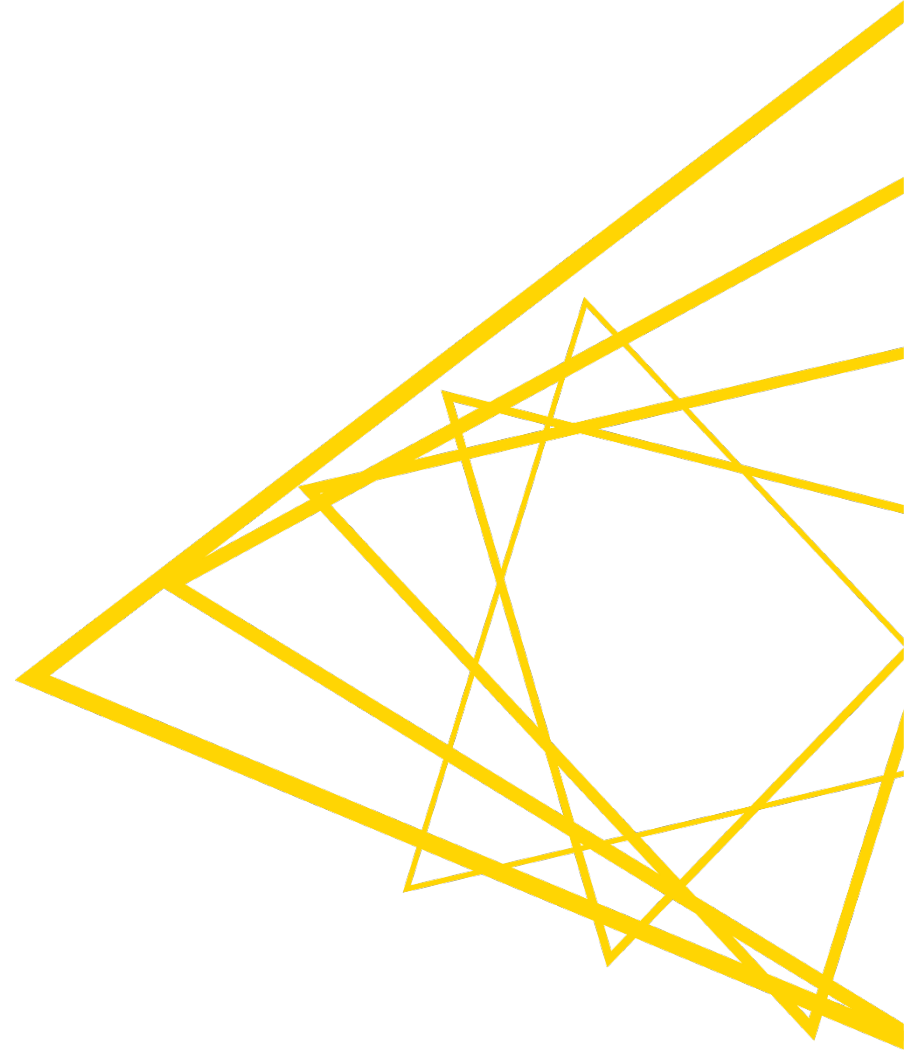**Berlin** 6:00 PM (CET)

New York City 12:00 PM (EDT)

Austin 11:00 AM (CDT)

London 5:00 PM (GMT)

Before we start…

- Please use the Q&A section to post your questions.
- Upvote for your favorite questions.

Challenges Operationalizing Data Science

Creating Great Data Science

Large, complex variety of technologies

Providing skills and resources

Continuous, high quality delivery

Mainstreaming across the organization

Data Science Driven Applications

R&D	MFG	IT	Marketing	Sales	… more

KNIME
Open for Innovation

Challenges Operationalizing Data Science

Productionizing Data Science

| Large, complex variety of technologies | Providing skills and resources | Continuous, high quality delivery | Mainstreaming across the organization |

Data Science Driven Applications

| R&D | MFG | IT | Marketing | Sales | … more |

KNIME
Open for Innovation
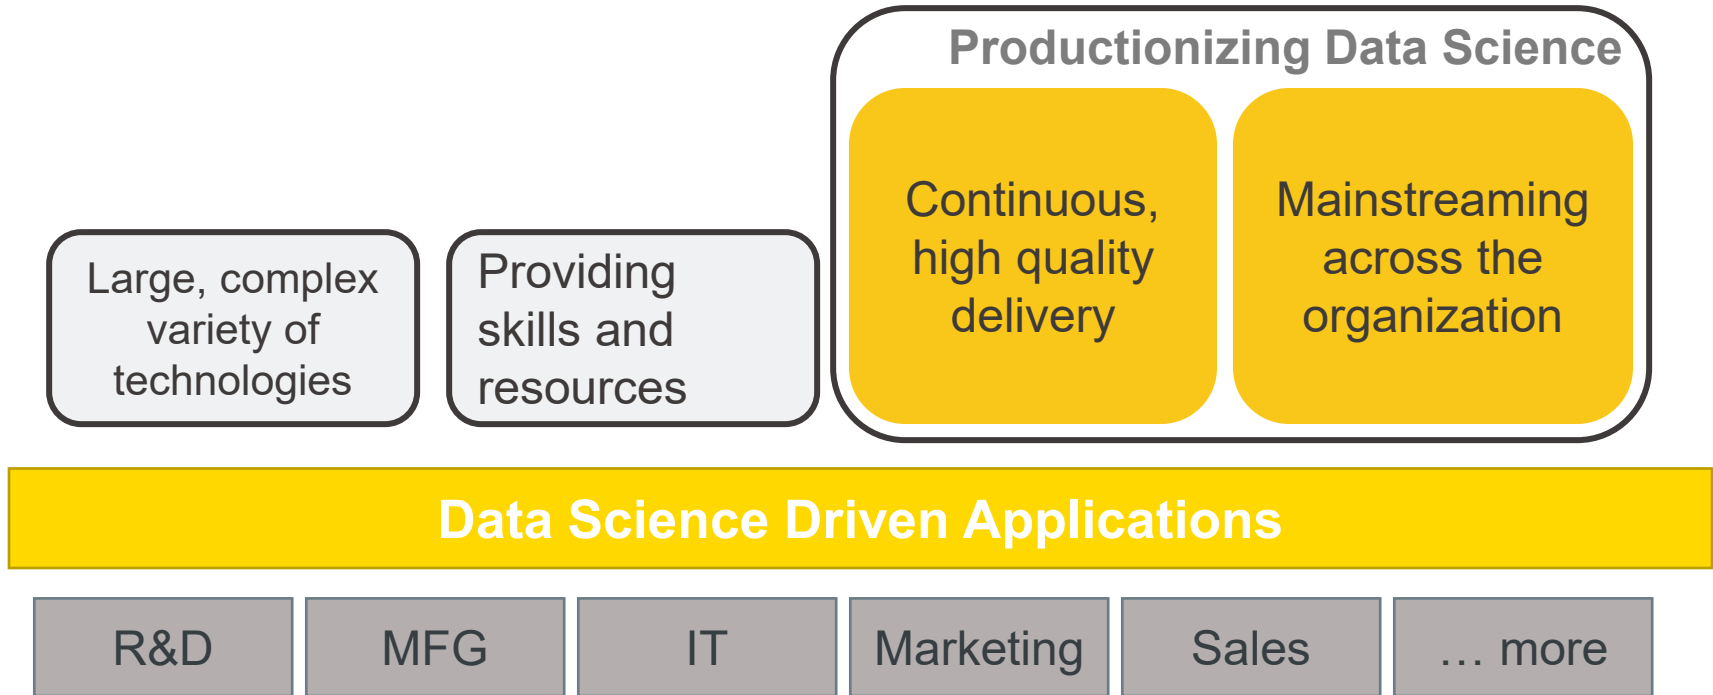
KNIME Software: for Individuals & the Business

Create

Productionize

Gather & Wrangle

Model & Visualize

Deploy & Manage

Consume & Optimize
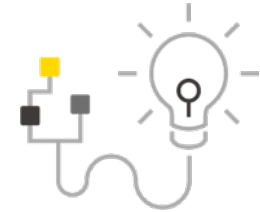
**Open Source
KNIME Analytics Platform**

| KNIME Extensions | KNIME Integrations | Community Extensions | Partner Extensions |

**Commercial
KNIME Server**

| KNIME WebPortal | Data Science as a Service |

Open for Innovation
KNIME

KNIME Server

- Collaboration
  - Share Expertise and Templates/Blueprints

- Automation
  - Schedule, Monitor, Update

- Deployment
  - To Humans: Analytical Applications ("Guided Analytics")
  - To Machines: via REST and custom APIs

- Management
  - Manage distributed setups
  - Integrate with existing enterprise setup (authentication)

KNIME Server

Collaboration

Workflow Hub | Workflows | Metanodes, Nodes | Data

Shared Repository, Versioning

Automation

Provisioning — Scheduled Execution | Remote, Client Control — Remote Execution | Distributed Executors

Deployment

Web Front-ends — Web Browser Access | Guided Analytic Apps | RESTful Services — Custom Application Access

Management

Nodes, Files, Applications — Security | Configuration & Versioning | Client & User Preferences

Open for Innovation
KNIME

Live Demo

Guided Analytics for ML/AI Automation

- Interaction & Automation
- Data Scientist's Choice: right mix for target audience
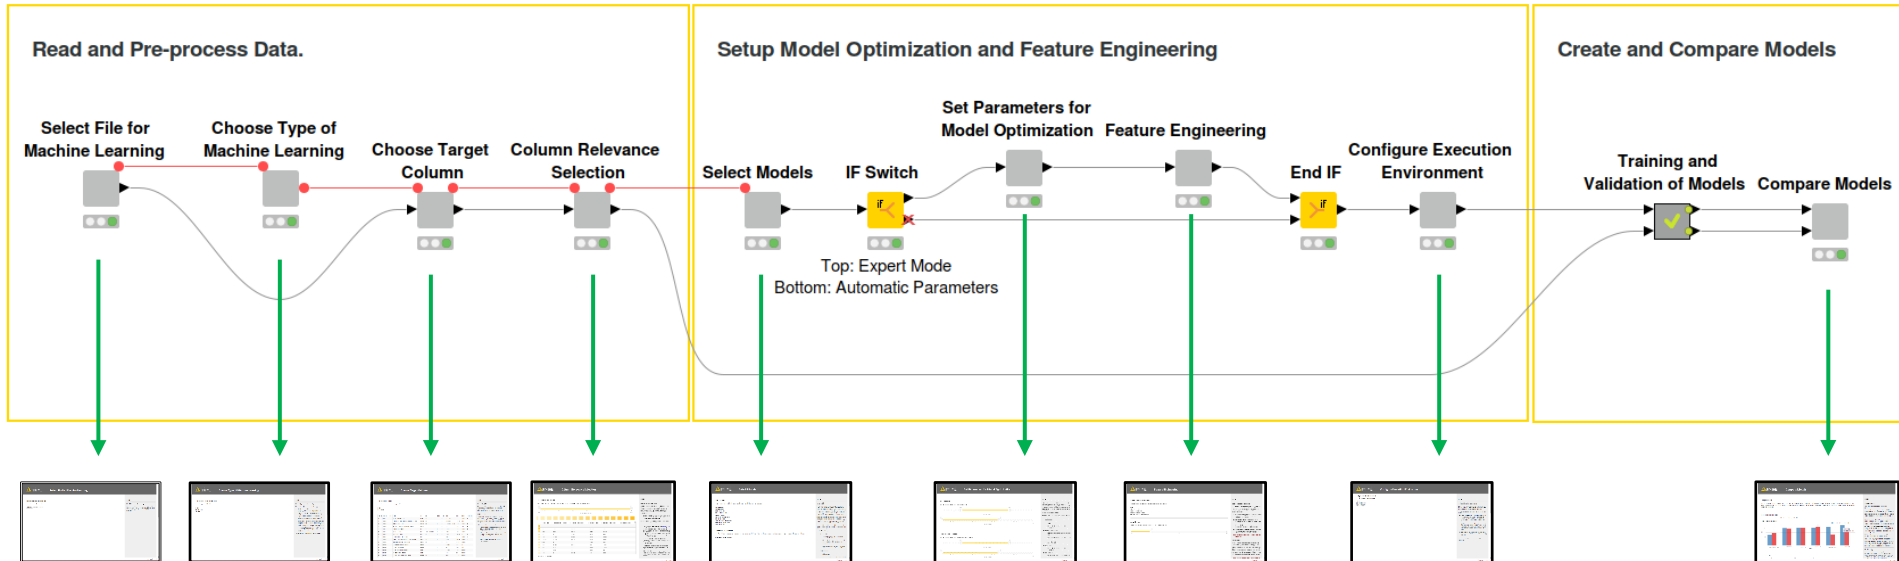
Compare Models Time

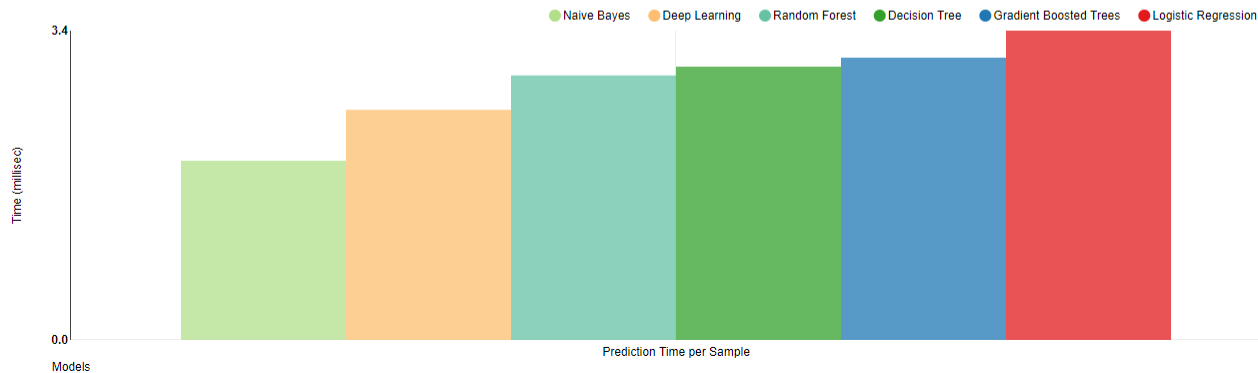Compare in those bar charts the time each model took to train and to compute a sample prediction.

Training Time

Prediction Time per Sample

Details for Experts

You can compare the models now by different metrics. Above you can see a bar plot with 3 main measures of performance: *Accuracy* is the percentage of correct predictions among all predictions. *Area Under Curve (AUC)* measures of the area under the Receiver Operating Characteristic curve.

The ROC plot describes the Receiver operating characteristic curves, one for each model. On the y-axis you have the true positive rate, on the x-axis you have the false posive rate based on the class you precedently provided. The greater the area under the model curve (AUC) the better is the performance of the model. More info about this plot can be found at this link.

Advanced Assessment of Models

The advanced assessment of models sections shows four additional charts per model.

- **1. Performance Metrics Bar Charts**

    For this visualization we measured the following metrics:

    - *Recall* (or True Positive Rate) (% of "casualty" rows correctly classified)
    - *Precision* (or Positive Predicted Value) (% of predicted "casualty" rows correctly classified)
    - *Specificity* (or True Negative Rate) (% of not "casualty" rows correctly classified)
    - *F-measure* (harmonic average between Recall and Precision)

- **2. Cumulative Gain Chart and Lift Chart**

    This chart is drawing a curve that reflects how well the model is doing compared to a random classifier. You are selecting rows from the test ranked by the probability of class "casualty". On the x-axis you have the percentage of top ranked rows by the model that define the partition of rows you are considering. On the y-axis you measure the response as the percentage of "casualty" rows over their total number in your partition of top ranked rows. If the model is bad, the curve will be close to the black line, where the percentage of original "casualty" rows is exactly equal the percentage of selected rows, (e.g. where the model ranking is random). The cumulative gain curve should be above the bisector line and the greater the area between the cumulative gain curve and the bisector line is, the better the model is.

Reset Apply ▲ Close ▲

Open for Innovation
KNIME

Advanced Assessment of Models

Each row represents a series of additional information about each created model.

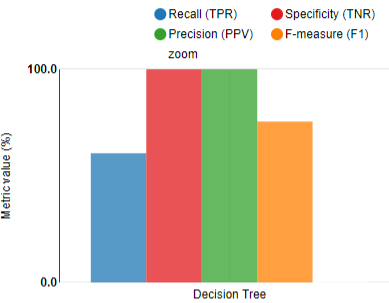## Decision Tree

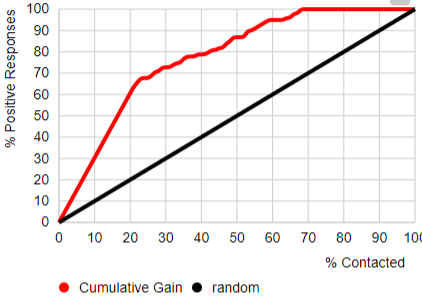Figure 1: Performance Metrics

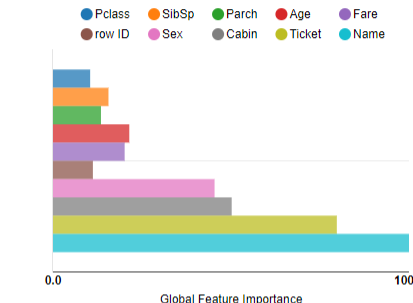Figure 2: Cumulative Gain Chart

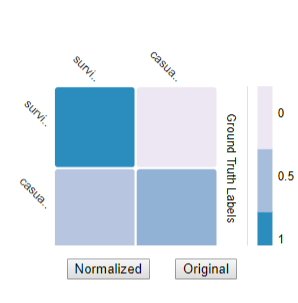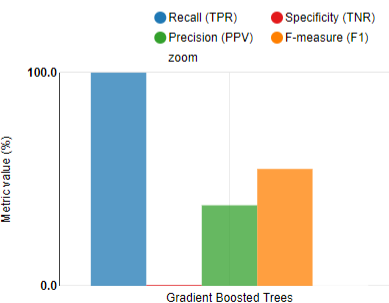Figure 3: Global Feature Importance

Figure 4: Confusion Matrix Heatmap

Gradient Boosted Trees

Figure 1: Performance Metrics

Figure 2: Cumulative Gain Chart
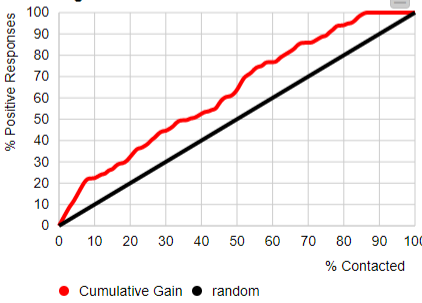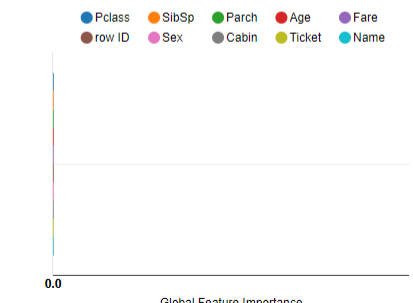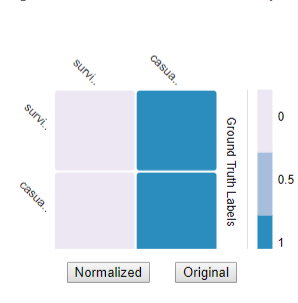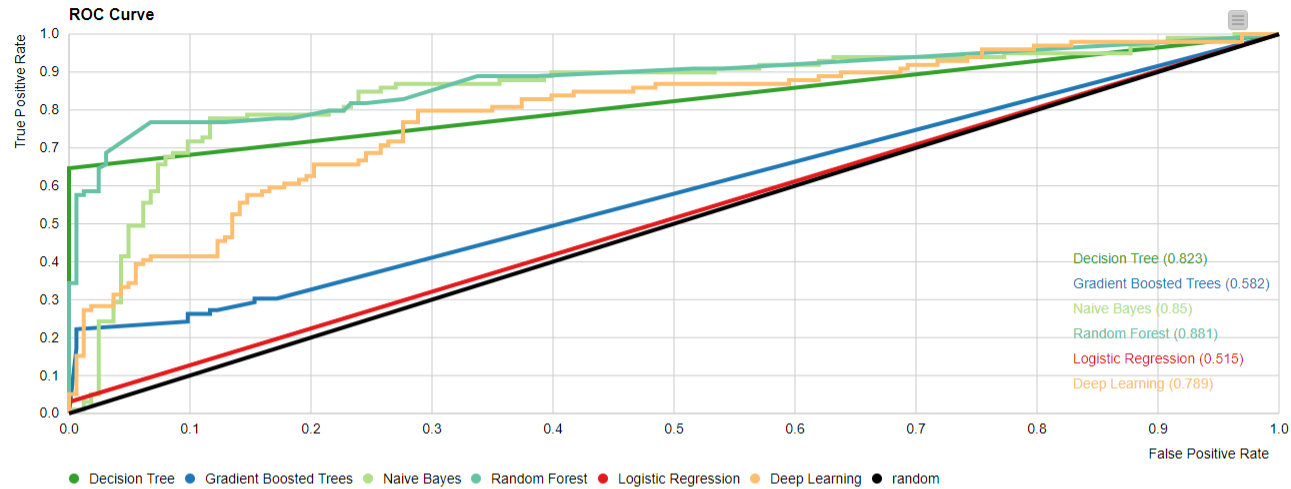

Figure 3: Global Feature Importance

Figure 4: Confusion Matrix Heatmap

Naive Bayes

Reset   Apply   Close

21

KNIME
Open for Innovation

ROC Curve Plots

Compare the curves, one for each model. The greater the area under a curve the better the model is.

ROC Curve

Decision Tree (0.823)
Gradient Boosted Trees (0.582)
Naive Bayes (0.85)
Random Forest (0.881)
Logistic Regression (0.515)
Deep Learning (0.789)

● Decision Tree ● Gradient Boosted Trees ● Naive Bayes ● Random Forest ● Logistic Regression ● Deep Learning ● random

Download Model

The following table summarizes the information in the charts. Please select from the table the model you would like to download and use.

Model	Accuracy (%)	Training Time (mins)	Prediction Time (millisec)	Area Under Curve (%)	Workflow	Parameters
Decision Tree	94.276	0.4	3	82.323	Download	Download
Logistic Regression	86.819	0.4	3.4	51.515	Download	Download
Random Forest	84.812	0.8	2.9	88.062	Download	Download
Naive Bayes	82.713	0.4	2	85.047	Download	Download
Deep Learning	81.28	4	2.6	78.881	Download	Download
Gradient Boosted Trees	50.036	0.9	3.1	58.226	Download	Download

cumulative gain chart curve and the bisector line.

- **3. Global Feature Importance Bar Chart**

  We consider for each case only the original selected features for the relative model. Then we train a surrogate random forest model overfitting the test set predicted classes. From such model it is possible to measure how often each feature is useful to outcome a prediction. More info at this link.

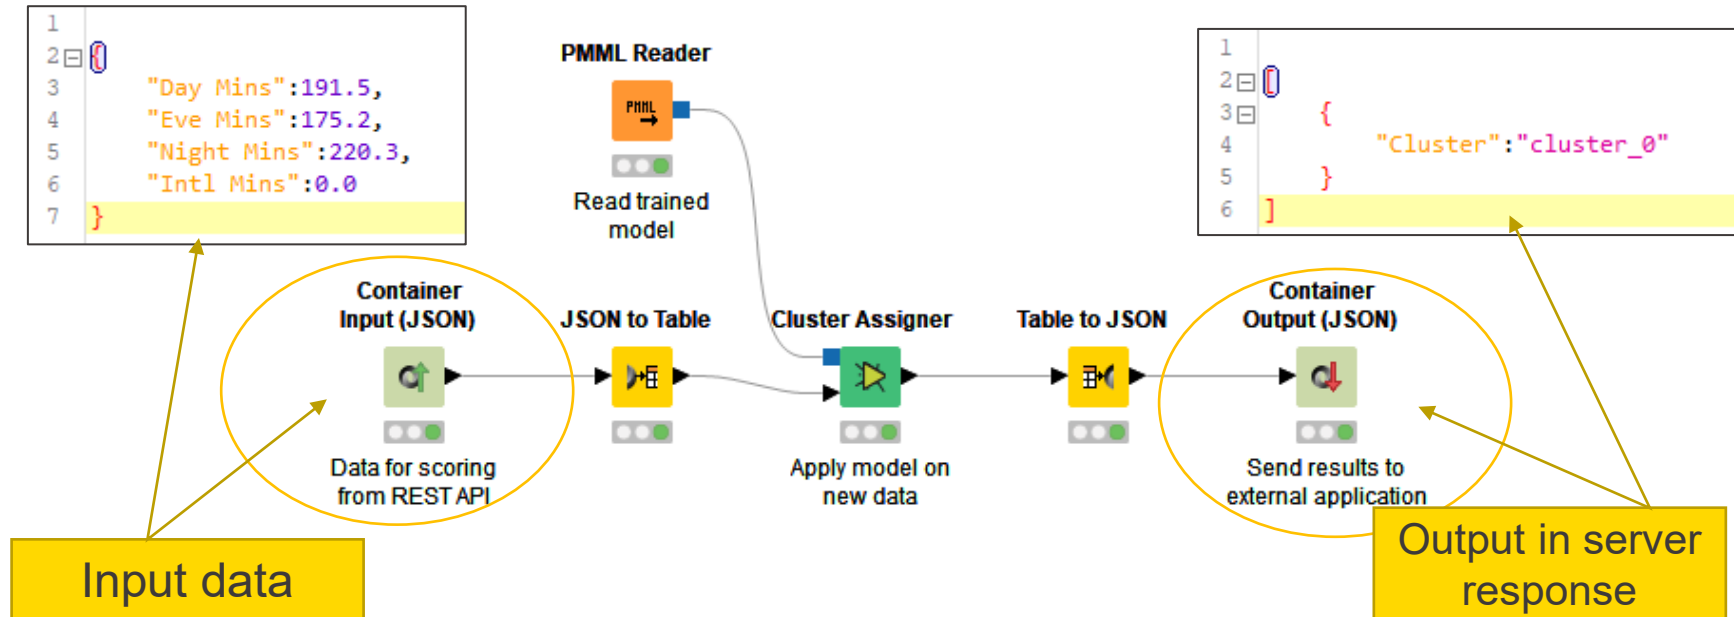- **4. Normalized Confusion Matrix Heatmap**

  A confusion matrix is summarizing all the outcomes on the test set by considering how many instance fall in each case given by prediction and original class. The heatmap is encoding with shades of blue which cell has the most instances in each row. A model with a good performance should have most of the dark blue cells on the diagonal from top-left corner to bottom right corner of the squared image. More info at this link.

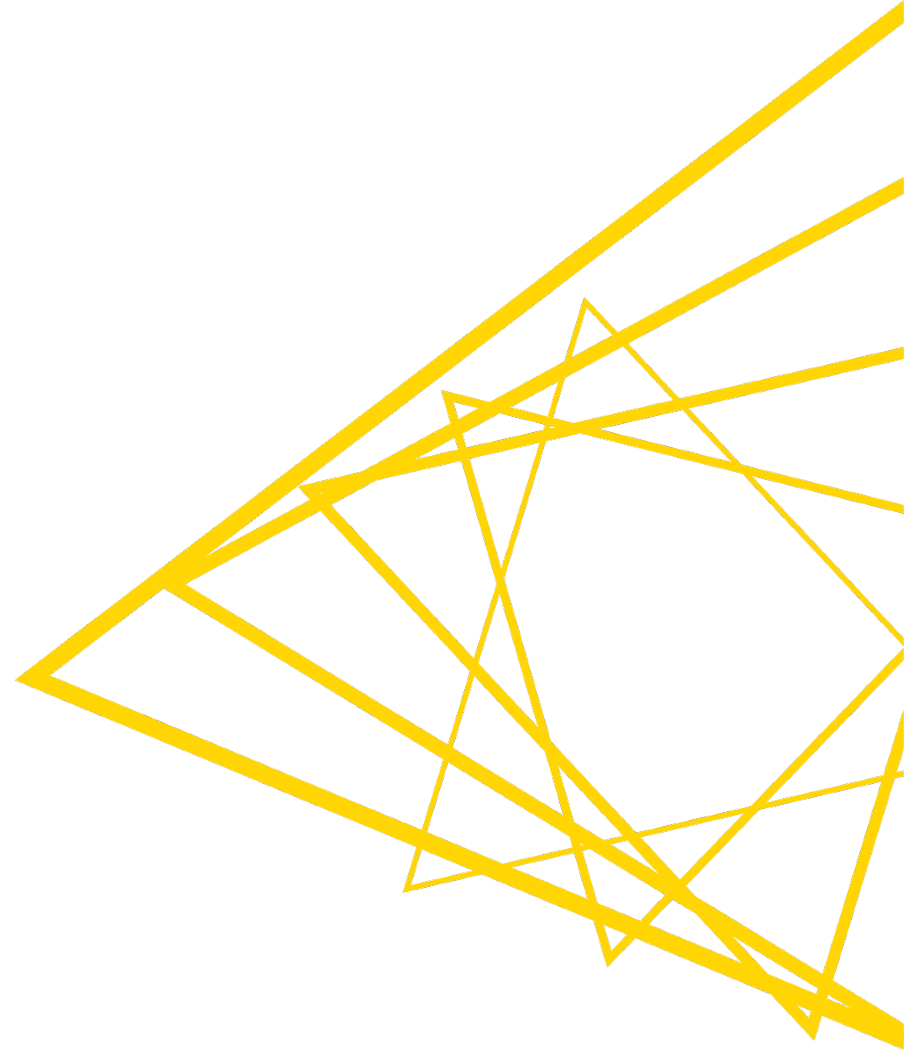Reset   Apply   Close

22

KNIME
Open for Innovation

KNIME Server REST API

- Enables external integration
- Build applications around KNIME Server
- e.g. for Microservices and real time scoring
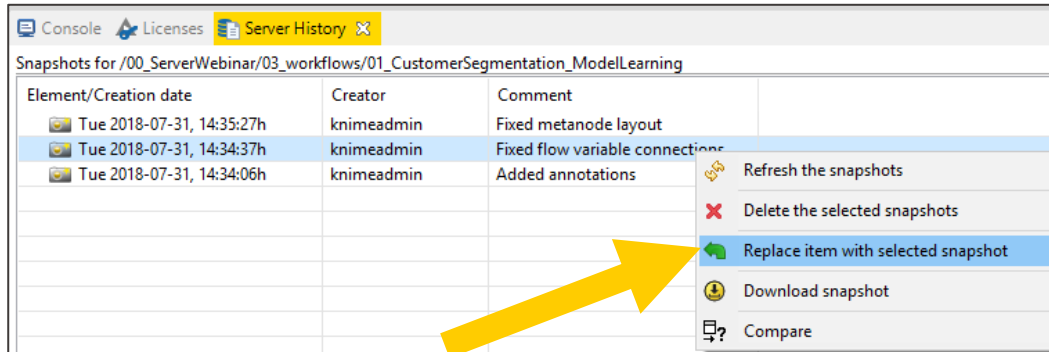
Deploy KNIME workflows as web services


```
1
2 {
3     "Day Mins":191.5,
4     "Eve Mins":175.2,
5     "Night Mins":220.3,
6     "Intl Mins":0.0
7 }
```

PMML Reader
PMML
→
Read trained model

```
1
2 [
3     {
4         "Cluster":"cluster_0"
5     }
6 ]
```

Container Input (JSON)
Data for scoring from REST API

JSON to Table

Cluster Assigner
Apply model on new data

Table to JSON

Container Output (JSON)
Send results to external application

Input data

Output in server response

23

Open for Innovation
KNIME

Live Demo

Version Control

- Create snapshots of workflow revisions

- Allow roll-back to previous versions

- Full workflow history available in client and on Workflow Hub

Workflow Difference

- Inspect differences between versions of the same workflow

Highlight differences:
- Nodes included/excluded
- Node configurations

Permission Management

- Manage permissions for workflows and data on KNIME Server

- Control which groups can access, modify, or execute a workflow

- Use **LDAP / Active Directory** groups for centralized administration
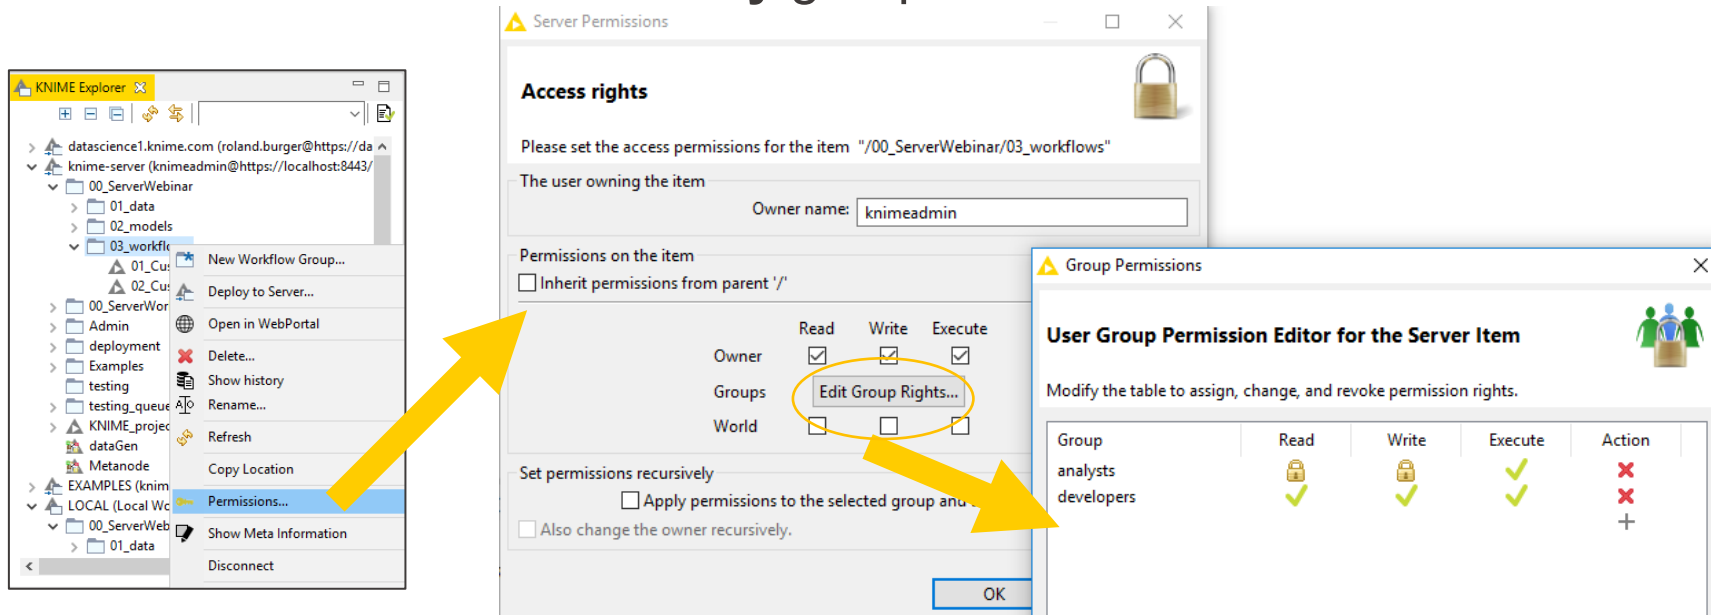
Managing Preferences – local

How do I configure Python?

How do I install database drivers?

KNIME Analytics Platform

KNIME Analytics Platform
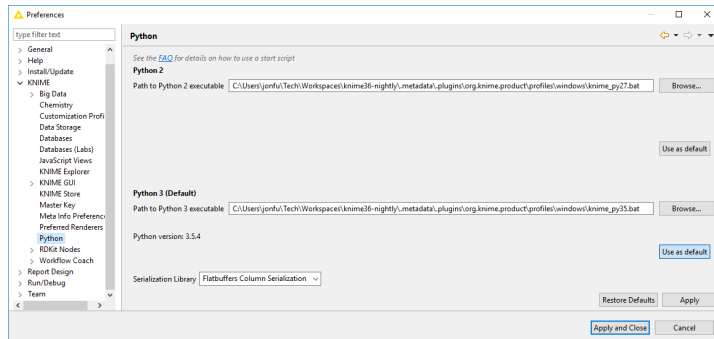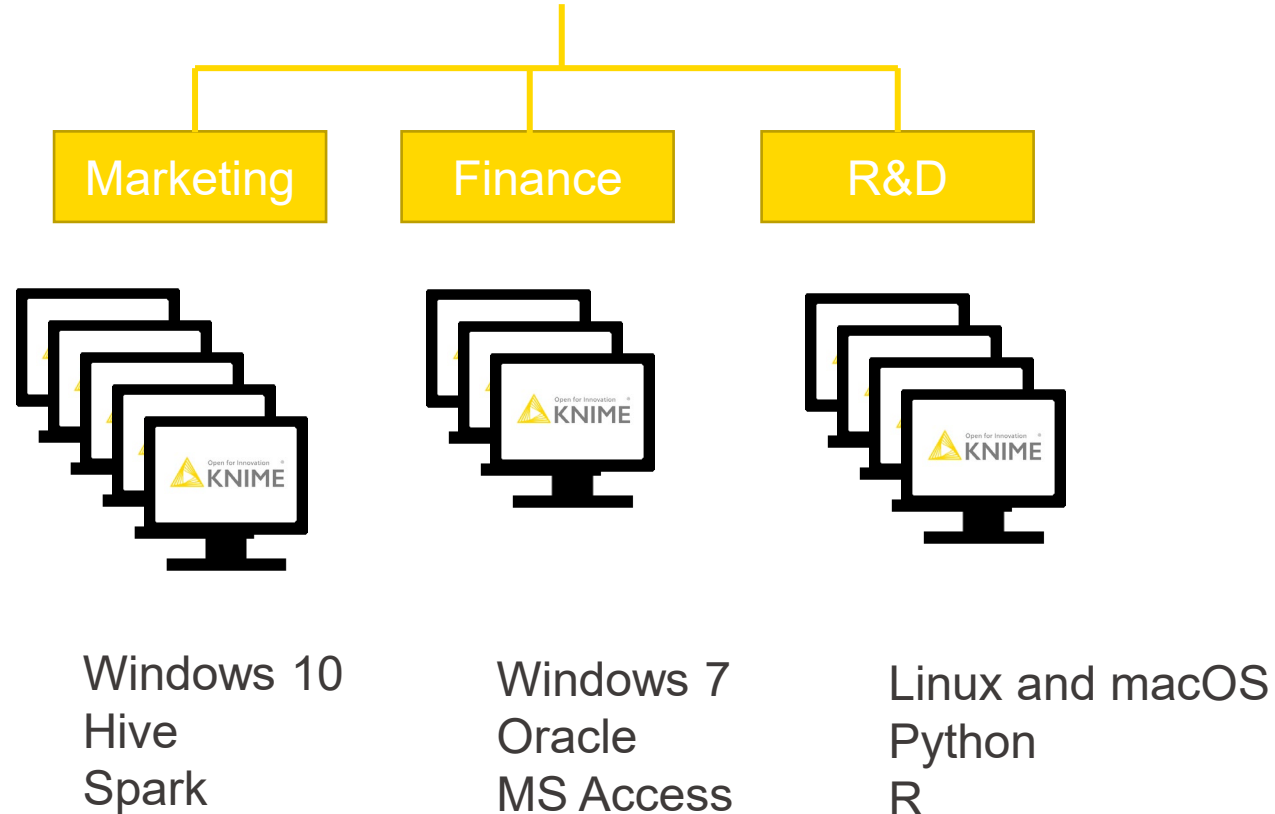
KNIME Analytics Platform
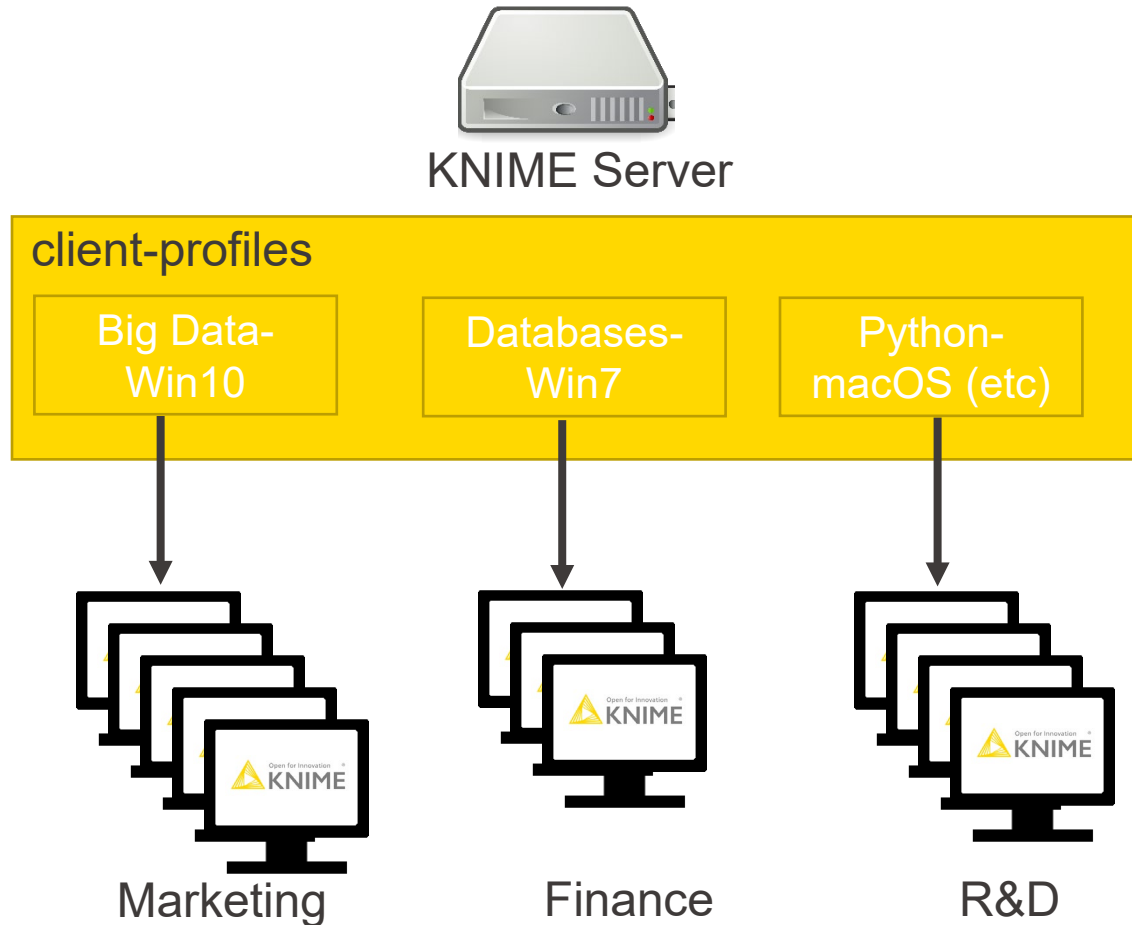
KNIME Website
Internal knowledgebase
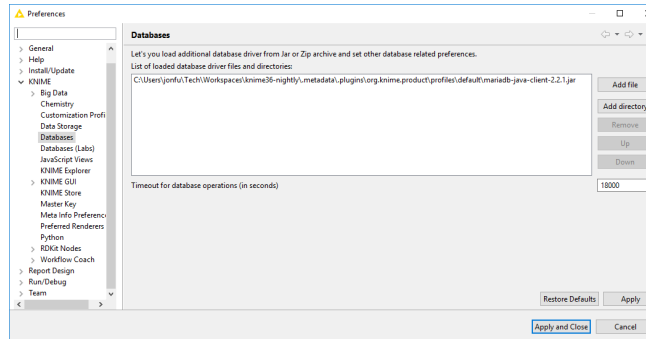
Managing Preferences – KNIME Server

- Different departments/teams have different requirements

- Multiple OS deployments
  - Windows 7
  - Windows 10
  - Linux
  - macOS

Marketing	Finance	R&D

Windows 10
Hive
Spark

Windows 7
Oracle
MS Access

Linux and macOS
Python
R

Managing Preferences – KNIME Server

KNIME Server

client-profiles

Big Data-Win10	Databases-Win7	Python-macOS (etc)

Marketing
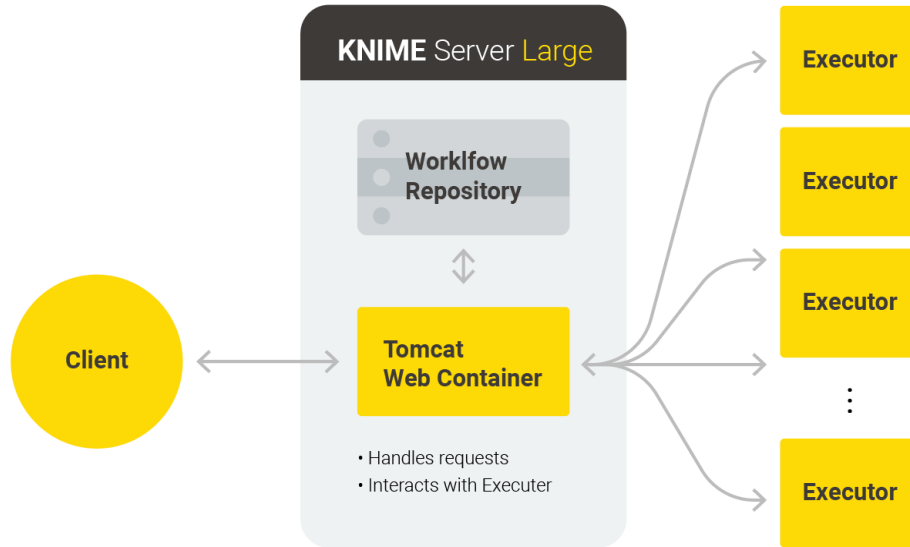
Finance

R&D

- Client-profiles
  - Python-Linux
  - Python-macOS
  - R-Linux
  - R-macOS
  - Databases-Win7
  - Big Data-Win10

Profiles can include:
Preferences, drivers, and more

On-Prem or Cloud, Scalable, Flexible
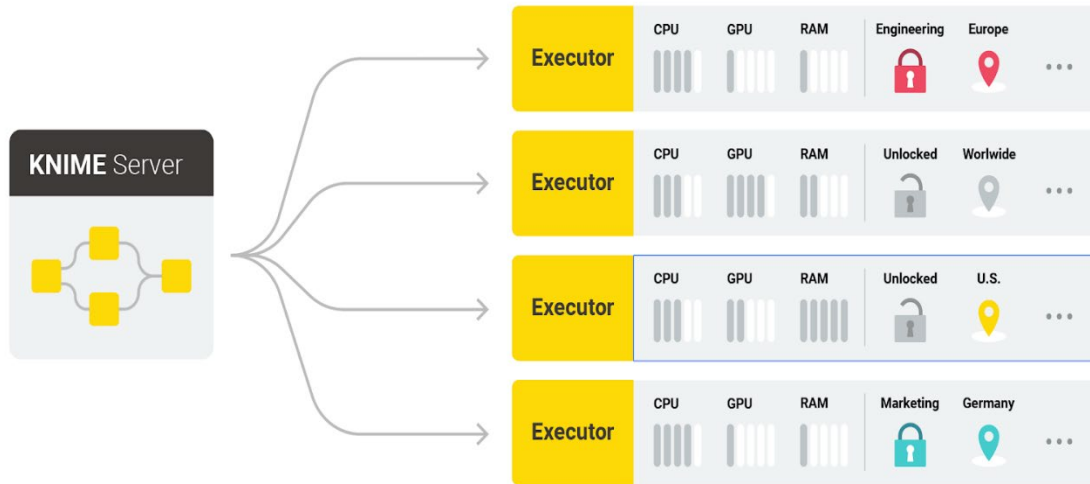
KNIME Server Large with multiple Executors

Features:
- Offload compute from KNIME Server
- Scale Executors to compute needs
- Support Executors with varying configurations
- Flexible deployment options

Capacity planning, Routing, Monitoring
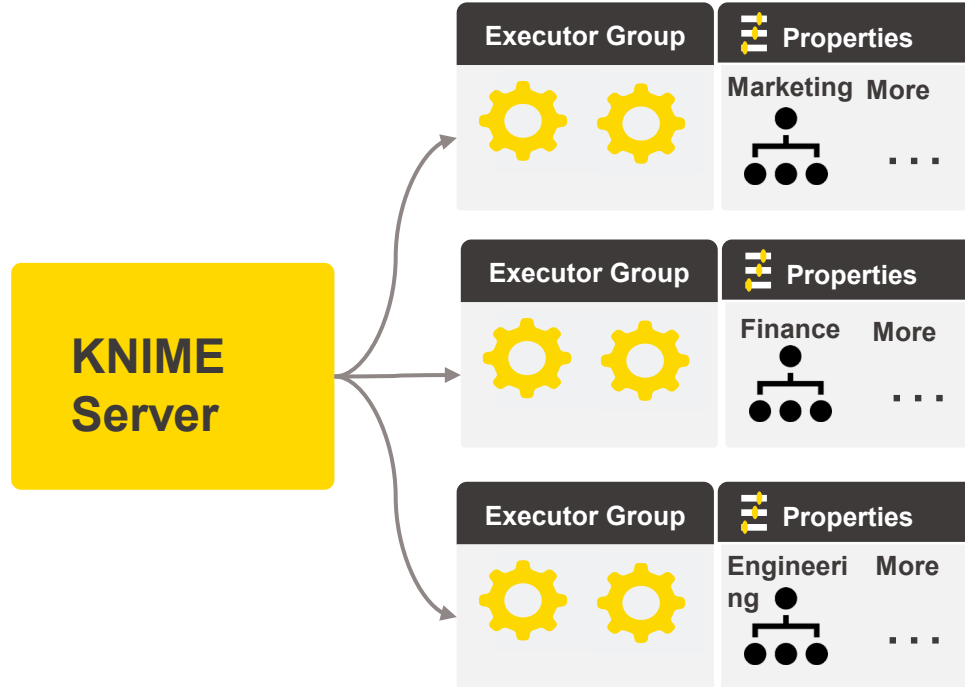
Pinning Workflows to KNIME Executors

Features:

- Match workflow needs to Executor capabilities
- Partition compute resources by capability, department, usage, …
- Workflow needs determined by workflow publisher

Expanded routing, Partitioning resources, Monitoring
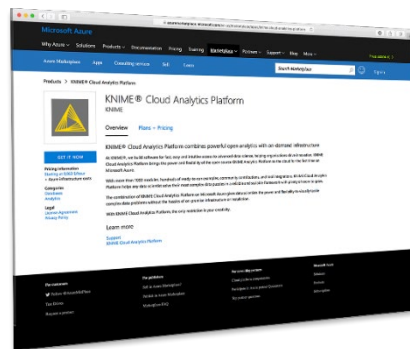
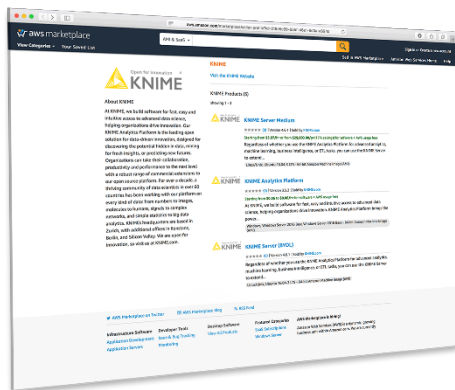Supporting Executors in the Enterprise[*]

Features:
- Logical groupings of Executors
- Match users/groups to Executor Groups
- Partition compute resources by groups, department, …
- Partitioning managed by Server administrators

*production July 2020

Open for Innovation
KNIME

Elastic Scalability, Cloud

KNIME Cloud Offerings – AWS and Azure

Features:

- KNIME Analytics Platform
- KNIME Server Large BYOL
  - Supports Server Large with multiple Executors
  - Has an embedded Executor so can be stand-alone
- KNIME Executors
  - Multiple Executors that can be used by KNIME Server Large
  - Pay as you go (PAYG) offering supports elastic scaling
  - Bring your own license (BYOL) offering uses cores from your Server license

Open for Innovation
KNIME

Elastic Scalability, Cloud
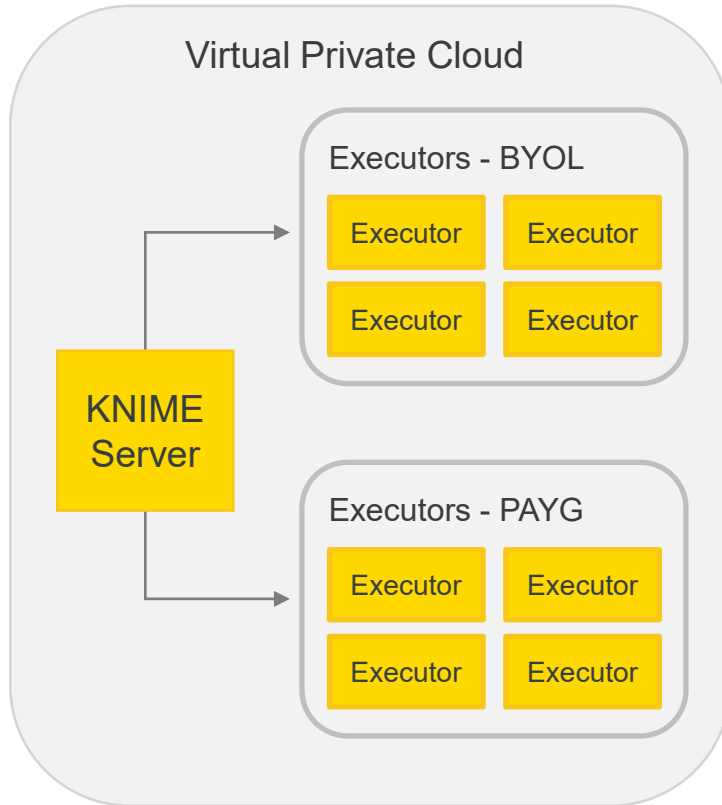
KNIME Executor Auto-scaling
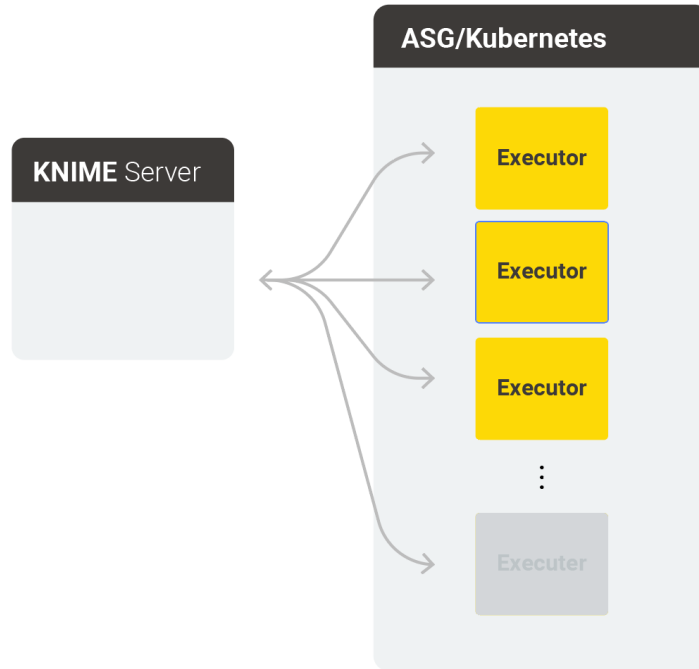
Features:
- Using cloud platforms or Kubernetes* auto-scaling capabilities
- Elastic scaling to meet demand
- Scales out to meet needed capacity
- Scales in to save costs
- Supported on AWS and Azure in PAYG licensing model

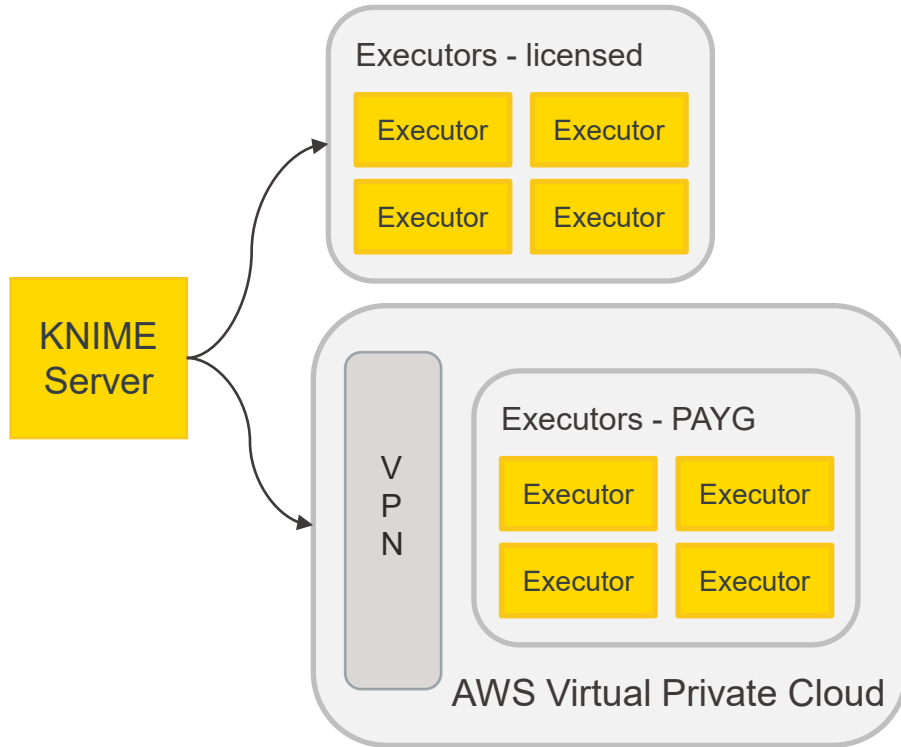# Flexible Cloud deployments to meet computing needs

Mixed Cloud Usage

Features:

- Supplement traditionally licensed Executors with Pay-as-you-Go (PAYG) model
- Meet periodic demand peaks
- Fulfill need for speciality hardware (e.g. GPU's)
- Meet budgeting needs

Flexible Hybrid deployments to meet computing needs

KNIME Hybrid Usage Model

Features:
- Mix of Enterprise data center and Cloud deployments
- Meet periodic demand peaks
- Fulfill need for speciality hardware (e.g. GPU's)
- Meet budgeting needs
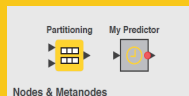
KNIME Server

Collaboration

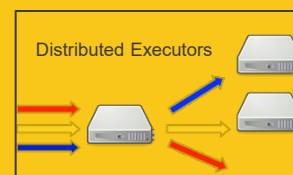Workflow Hub · Workflows · Metanodes, Nodes · Data
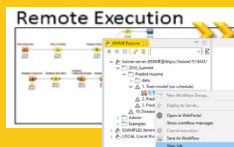
Shared Repository, Versioning

Automation

Provisioning · Remote, Client Control · Distributed Executors

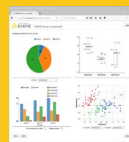Scheduled Execution · Remote Execution

Deployment

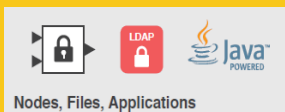Web Front-ends · Guided Analytic Apps · RESTful Services
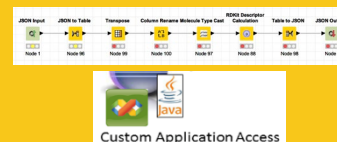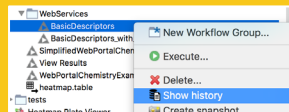
Web Browser Access · Custom Application Access

Management

Nodes, Files, Applications · Security · Configuration & Versioning · Client & User Preferences

Ready to get started with KNIME Server? Contact us for a free trial:

contact@knime.com

KNIME
Open for Innovation

Thank you!