



# What's New and Cooking in KNIME Server

Jim Falgout, KNIME

# Agenda

---


- Quick KNIME Server Intro
- What's New in KNIME Server
  - Single Sign-On using OpenID (OAuth)
  - Distributing work with Executors
  - Workflow Pinning
- What's Cooking in KNIME Server
  - Web Portal updates for Business Users
  - Supporting Executors in the Enterprise
  - Cloud mix
  - Bringing Executors to the edge

# Definitions

---

- What's New
  - Features added in the previous KNIME Server release
  - You can use them now!
- What's Cooking
  - Features being brewed for a future KNIME software release
  - May be available for limited preview


# Quick KNIME Server Intro



# KNIME Server

Create

Gather & Wrangle



Model & Visualize




Productionize

Deploy



Consume & Optimize



**KNIME  
Analytics Platform**

KNIME  
Extensions

KNIME  
Integrations

Community  
Extensions

Partner  
Extensions

**KNIME  
Server**

KNIME  
WebPortal

Data Science  
as a Service

# What's New in KNIME Server




# Single Sign-On

---

- Using OpenID (OAuth)
  - Supports using third-party identity providers such as Auth0, Google, and others
  - When enabled, logging into KNIME Server redirects to the identity provider
  - Configured in KNIME Server
  - Currently in preview
    - Additional functionality and documentation planned for the next release

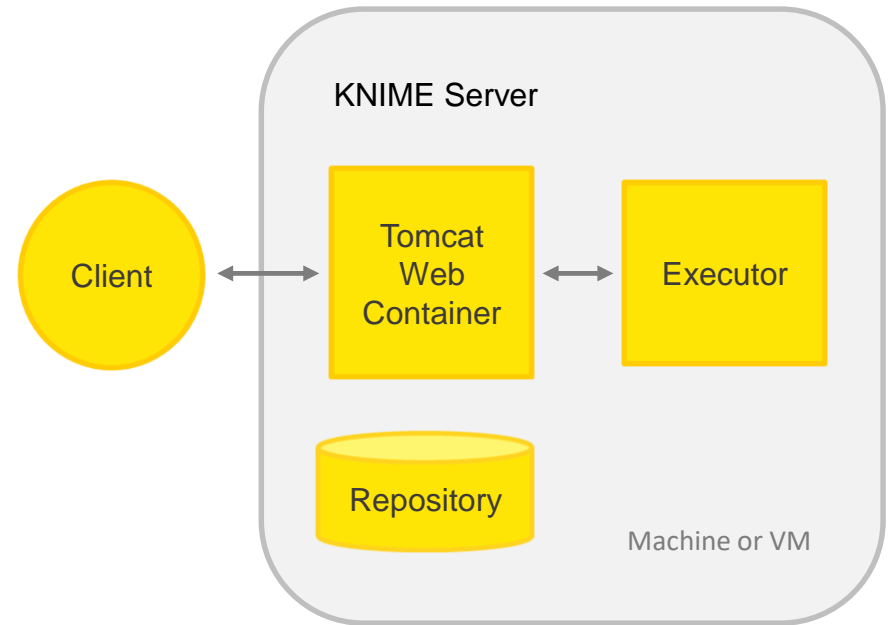
# Single Sign-On Demo





# KNIME Server Components

- KNIME Server
  - Web container with Server business logic
  - Executor for executing workflows
  - Repository for storing workflows and metadata
  - Single machine/VM



# Need for Off-boarding the Executor from Server

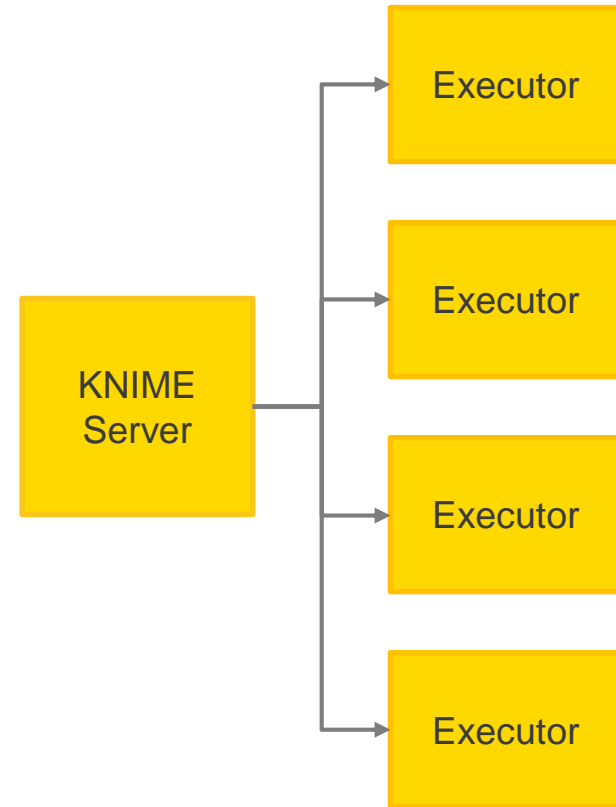
---

- Distribute load of handling Server requests from running workflows
- Support different types of Executor runtime environments
  - GPU's, high memory, ...
  - Installed extensions and runtime environments (Python)
- Add ability to scale compute resources (i.e. add more Executors)

# Distributed Executors for KNIME Server Large

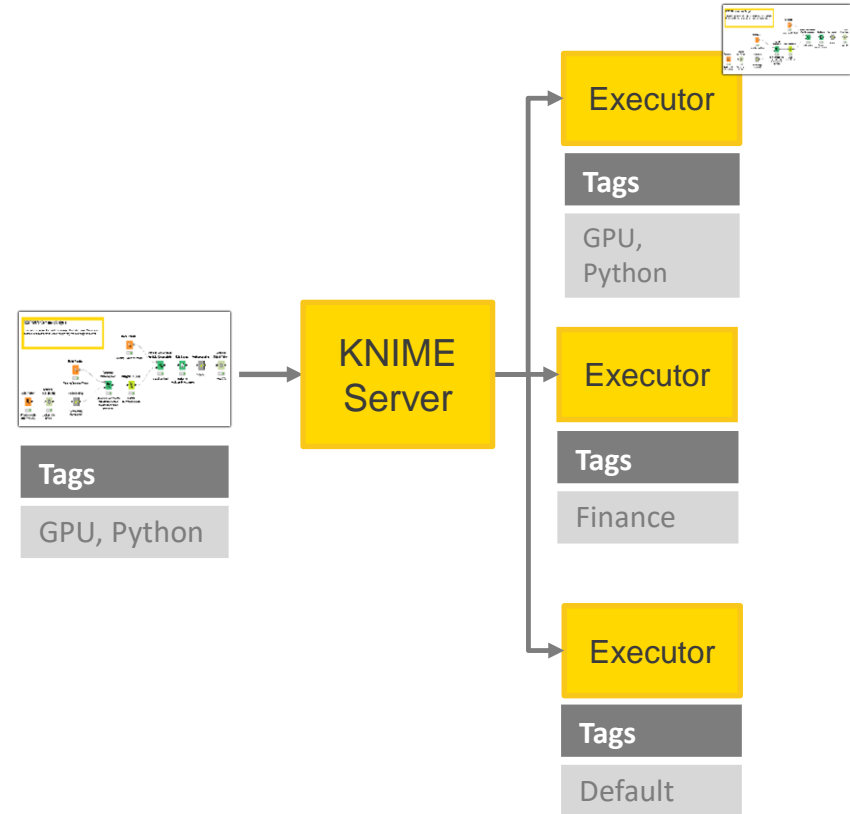
---

- Connect heterogenous executors to KNIME Server
- Executors run in their own machine or VM
- Deployment flexibility
- Licensed by the cores on executor instances



# Workflow Pinning

- Example Scenarios:
  - My workflow needs access to the Finance database
  - My workflow requires GPU's and Python
- Pinning Workflows
  - Tag a deployed workflow
  - Execute on Server
  - Workflow routed to matching Executor

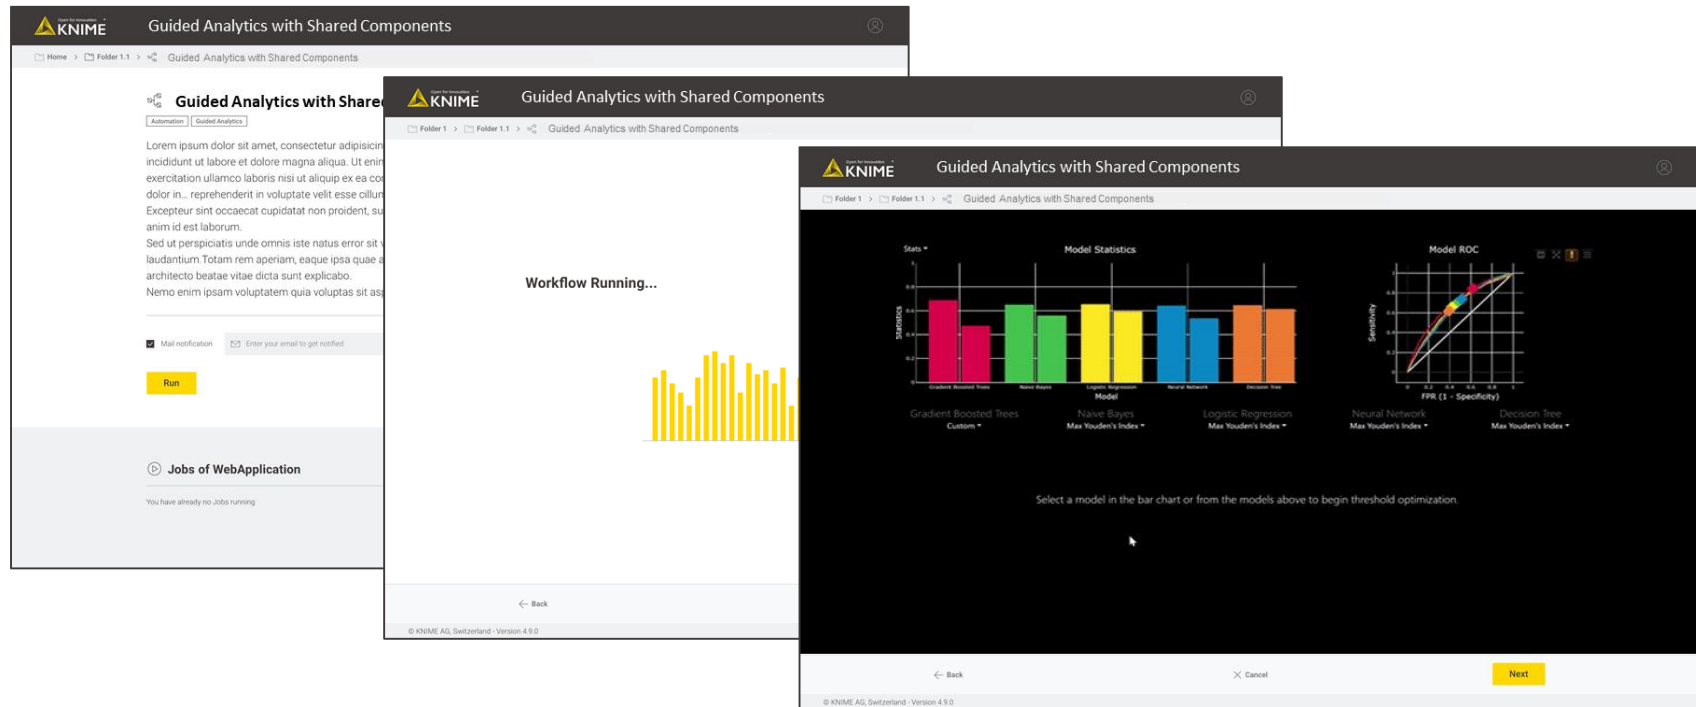


# What's Cooking in KNIME Server



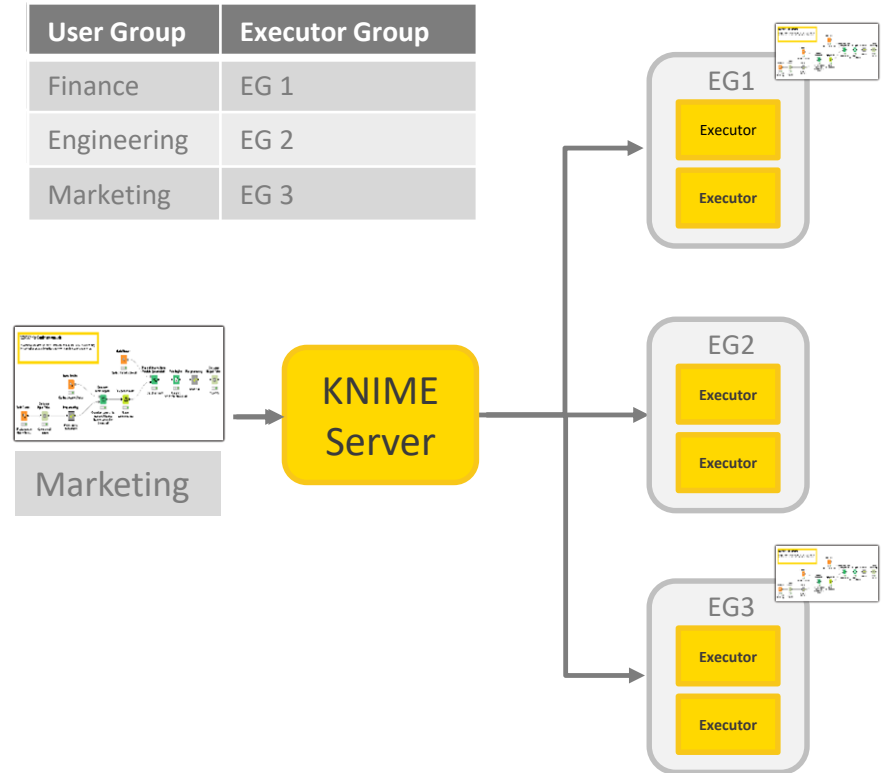
# New WebPortal

## UX Designed for Business Users



# Supporting Executors in the Enterprise

- **Scenario:** You want to segregate execution resources by department
- **Executor Groups**
  - Managed grouping of Executors
- **Managing Executor Access**
  - Assigned users/user groups to Executor Groups



# Cloud Mix – New Offerings

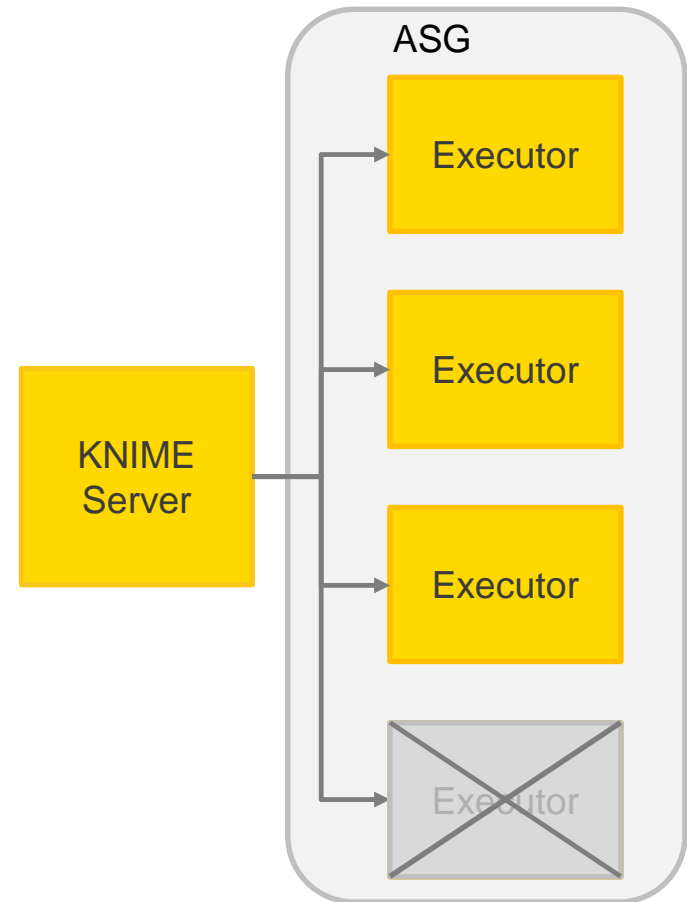
---

- New Cloud Marketplace Offerings
  - KNIME Server Large BYOL
    - Supports Server Large with Distributed Executors
    - Has an embedded Executor so can be stand-alone
  - KNIME Executors
    - Distributed 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



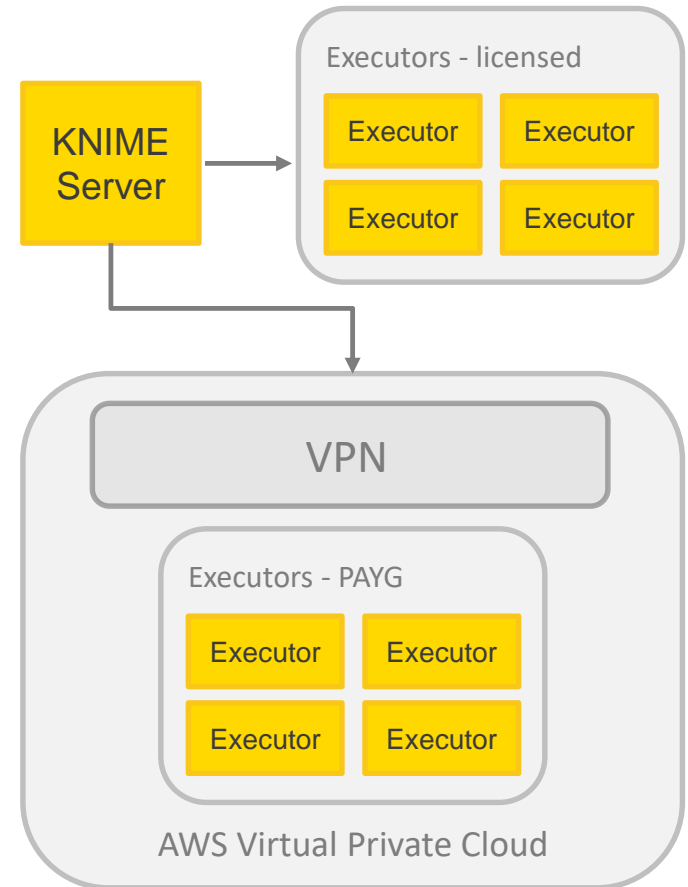
# Cloud Mix – Auto-scaling

- Scenario: Workloads vary and workload spikes are causing issues.
- Auto-Scaling PAYG Executors
  - Uses auto-scaling groups (ASG) in cloud
  - Scale out adds instance(s)
  - Scale in terminates instance(s)
  - Maintains a base capacity



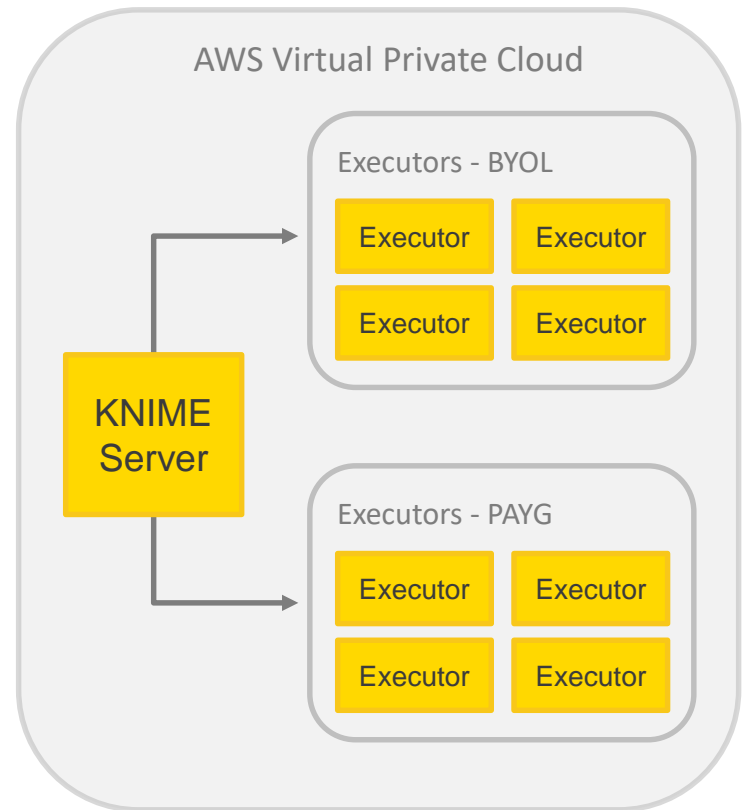
# Cloud Mix - Hybrid

- On-premise and cloud
- Scenarios:
  - Data to analyze is on cloud
  - Meet periodic demand peaks
  - Need specialty hardware
    - GPU's
    - High memory or CPU
  - Testing



# Cloud Mix – Mixed Cloud Usage

- Mix of BYOL and PAYG
- Scenarios:
  - Meet periodic demand peaks
  - Need specialty hardware
    - GPU's
    - High memory or CPU
  - Different budgeting needs
  - Executors in different cloud accounts



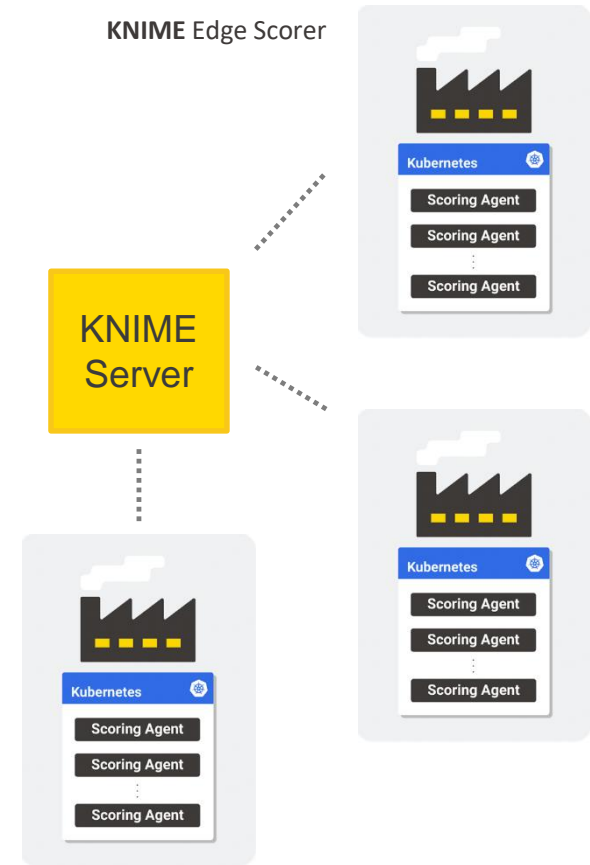
# Bringing Executors to the Edge

---

- I have a workflow that takes data, applies an algorithm/model and returns a score (answer)
- I need to deploy that to hundreds or thousands of end users, where there may be spikes in demand
- I need the latency to be low and consistent
- I need to update the model/workflow periodically
- I want to replicate deployments to multiple remote sites

# Bringing Executors to the Edge

- KNIME Scorer (TBN)
  - KNIME Server
    - Management and Monitoring
      - Workflows
      - Edge instances
    - Continuous deployment
  - KNIME Edge Scorer
    - Deploys workflows as containers
    - Provides API endpoint
    - Manages scaling



# Thank you!

