



**Trends in supercomputing architectures in the exascale era: heterogeneity, modularity, disaggregation**



**EUROPEAN TECHNOLOGY  
PLATFORM FOR HIGH  
PERFORMANCE COMPUTING**

A blurred background image showing people walking in a modern, brightly lit building with glass walls and structural elements. The image is overlaid with a semi-transparent white box containing the text "Unlock the future".

**Unlock  
the future**



# High Performance Computing **Built for the Cloud**

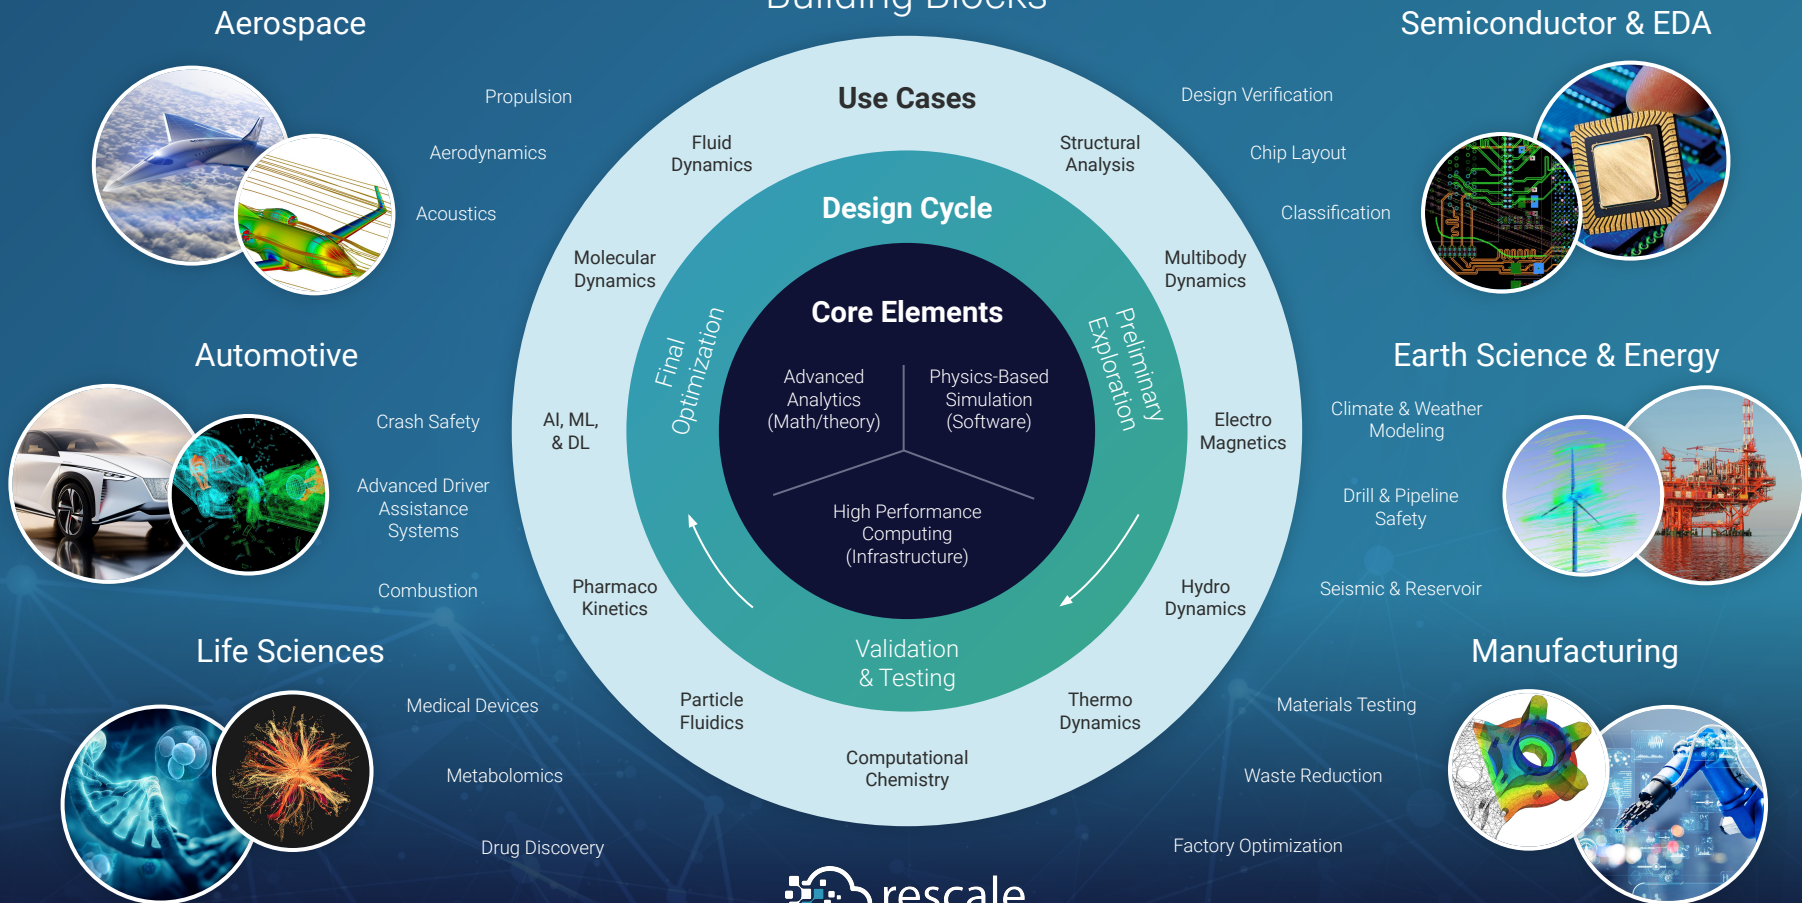
**Specialized Architectures and Computational Pipelines**

Teratec Forum - Jeudi 1<sup>er</sup> Juin 2023

[Romain Klein](#), EMEA Technical Director

# COMPUTATIONAL SCIENCE & ENGINEERING

## Building Blocks



# Engineering Innovation Has Been Underserved by Cloud

## Application Developers

Business Process Innovation  
*E.g. Ecommerce, CRM, Social, Mobile*

## Engineers & Scientists

Science & Engineering Innovation  
*E.g. Simulation, Modeling, Design Exploration*

### Cloud transformation accelerates software development

- Application developer-friendly tools
- Easy-access platform services (e.g., databases, message queues)
- Simple access to low-cost commodity hyperscale infrastructure

### Traditional HPC strategies in the cloud continue to constrain engineering innovation

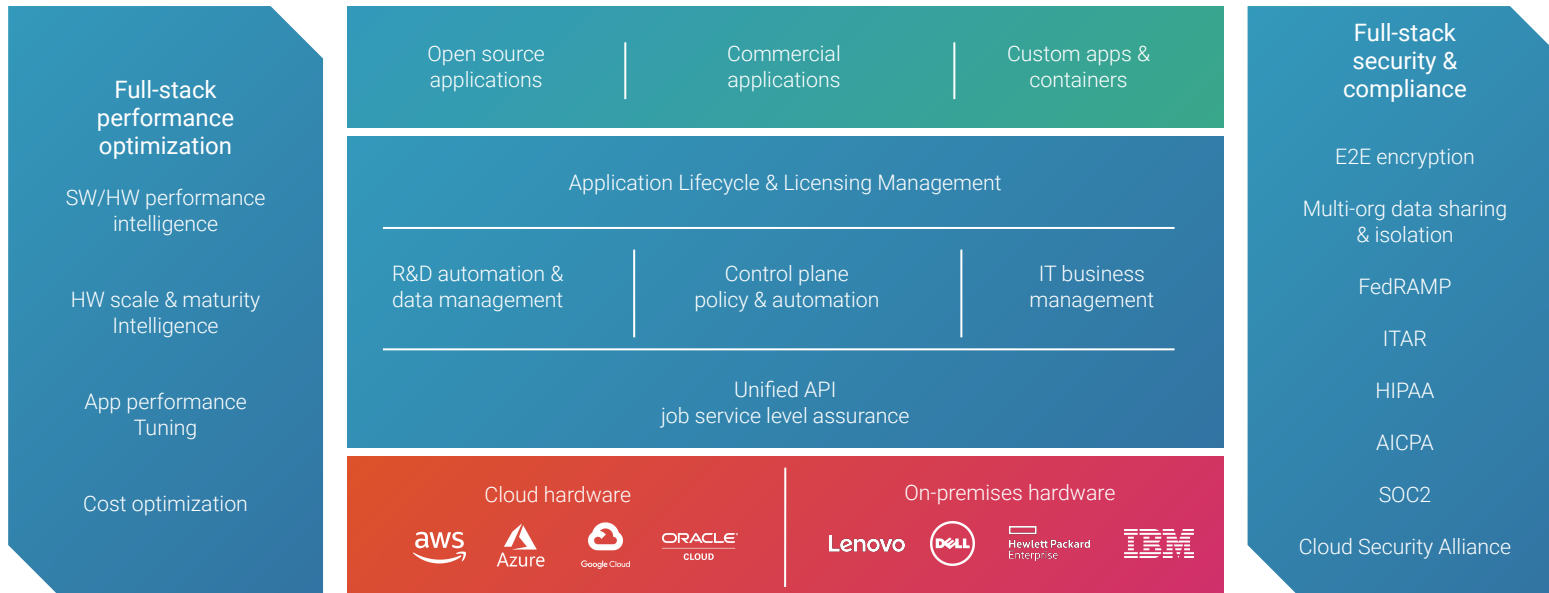
- Difficult user experience to run HPC
- Complex technology stacks and workflows
- Lack of easy access to new, specialized hardware

# Rescale Cloud HPC platform

Any software. Any hardware. Any cloud.



1000+ Applications



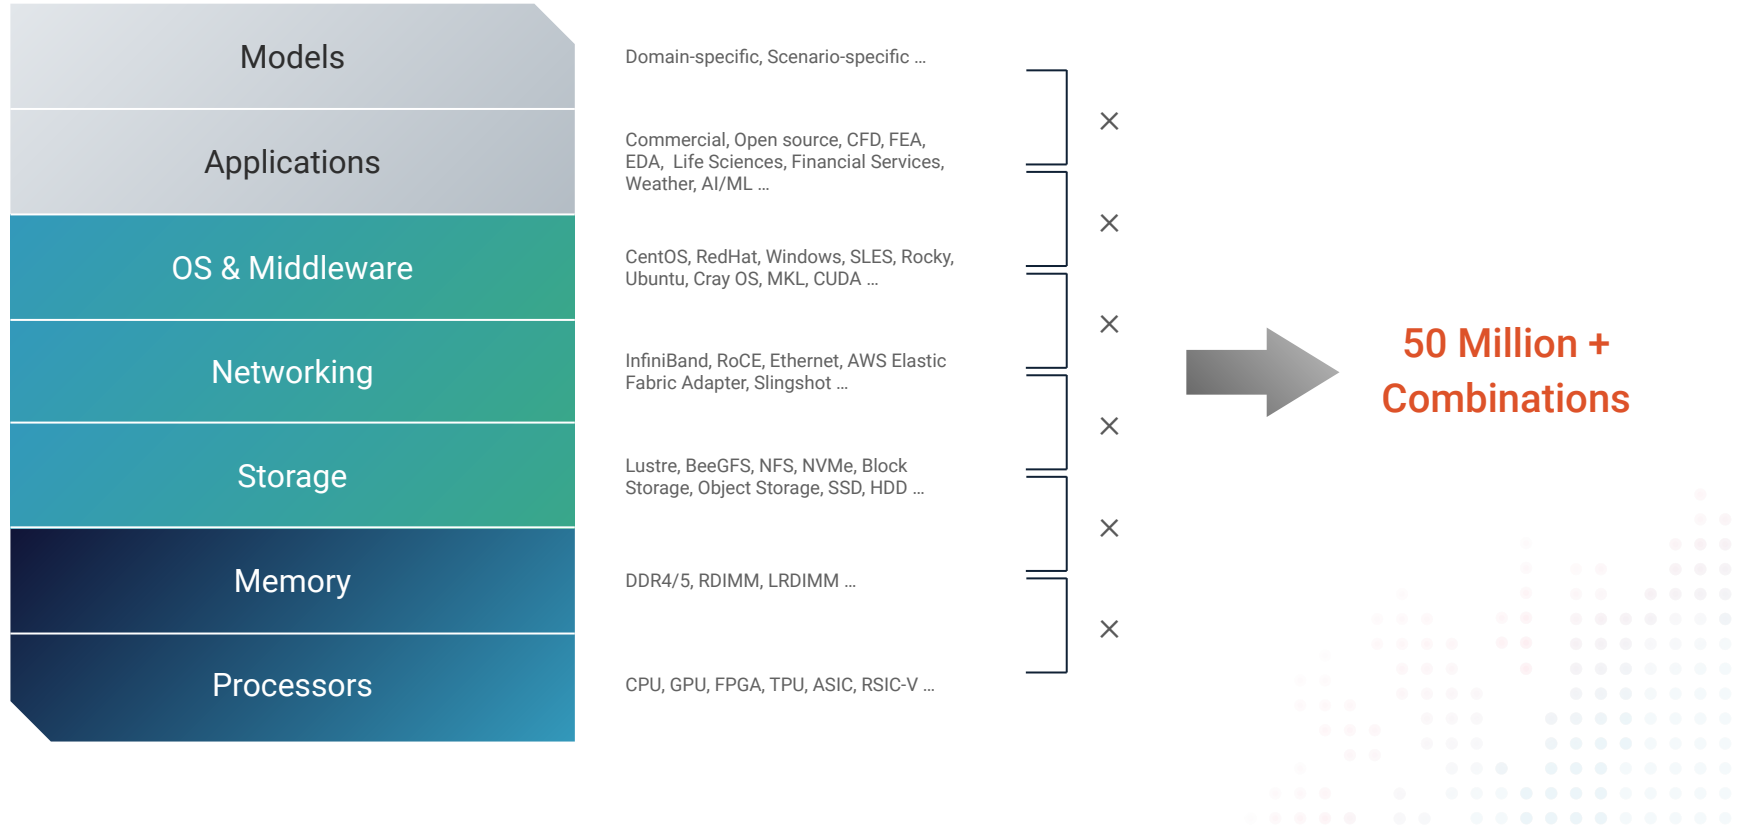
---

## The HPC Stack: Fragmentation and Specialisation

### Topics


- New Specialised Architectures Drive Performance Gains
- Take Advantage of Architectural Proliferation
- Specialised Jobs with Unique Needs
- High Throughput Computing (HTC)
- Rescale Software Publisher

# The HPC Stack: Fragmentation and Specialisation




# From Bottom Up to Top Down

Legacy hardware decisions  
constrain R&D velocity and IT flexibility



Compute Constrained

Fully optimised stack on-demand  
to meet R&D and IT objectives



Compute Empowered



# Rescale's global HPC infrastructure network

Over 100 data centers worldwide

Ontario, Canada  
California, USA  
Gov. Arizona, USA  
Illinois, USA  
Iowa, USA  
Ohio, USA  
Oregon, USA  
S. Carolina, USA  
Texas, USA  
Virginia, USA  
Washington, USA

Iceland, Europe  
Ireland, Europe  
W. C. Germany, Europe  
Frankfurt, Germany  
Paris, France  
Eemshaven, Netherlands  
Stockholm, Sweden  
London, UK

Mumbai, India  
Pune, India  
Tokyo, Japan  
Busan, Korea  
Seoul, Korea

## Leading Infrastructure Providers

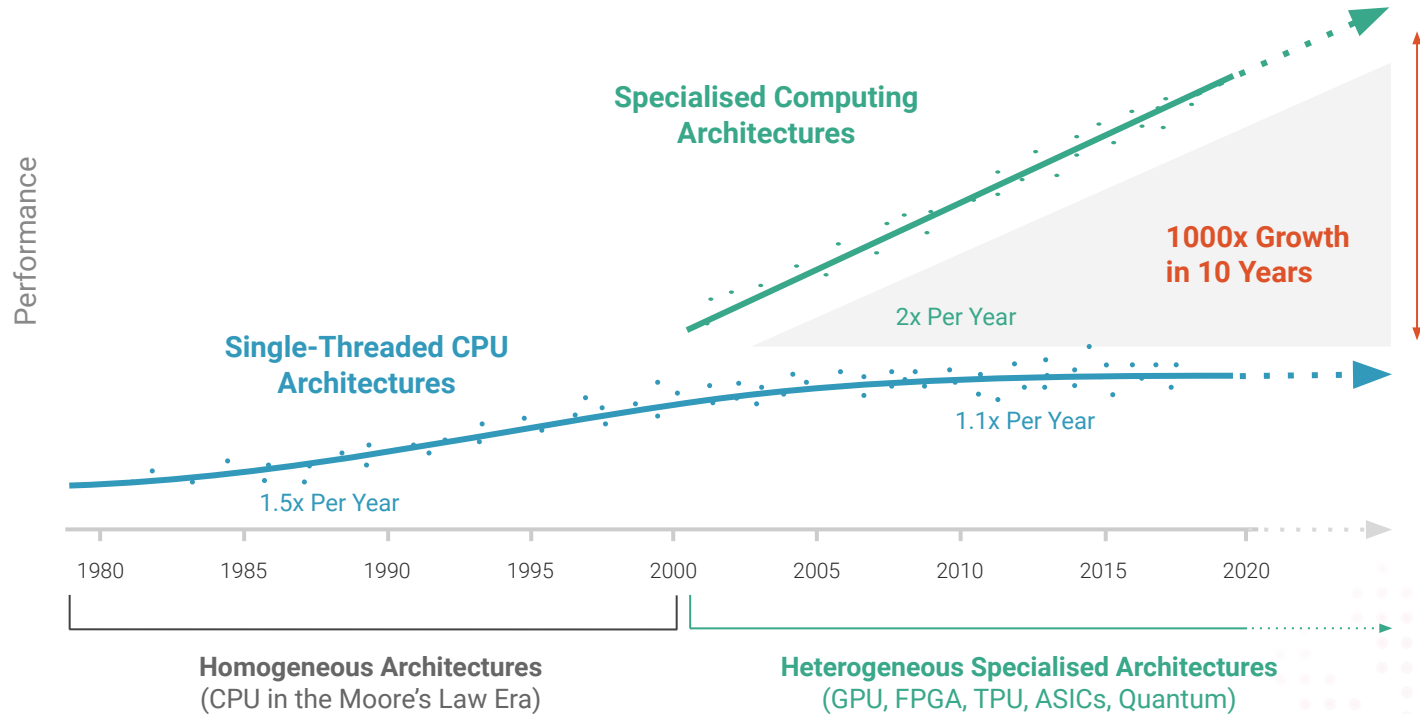


**ORACLE**  
Cloud Infrastructure

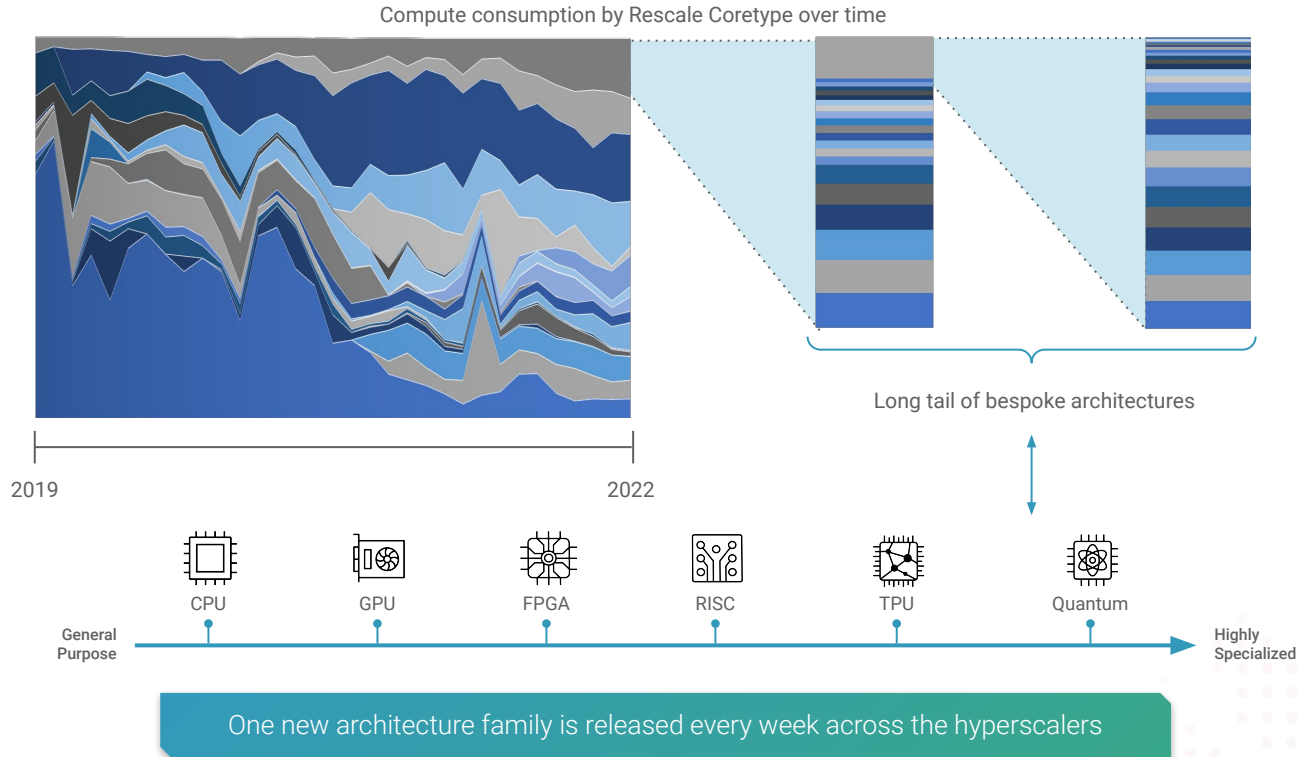
- Carbon-Neutral
- ITAR
- Non-ITAR

Turnkey deployment to the fastest local HPC infrastructure

# New Specialised Architectures Drive Performance Gains



# Rescale Users Take Advantage of Architectural Proliferation



# Specialised Computing at Cloud Scale Delivers Faster Results

## Specialised Jobs with Unique Needs



**Job 1**  
FEA - high clock speed requirement



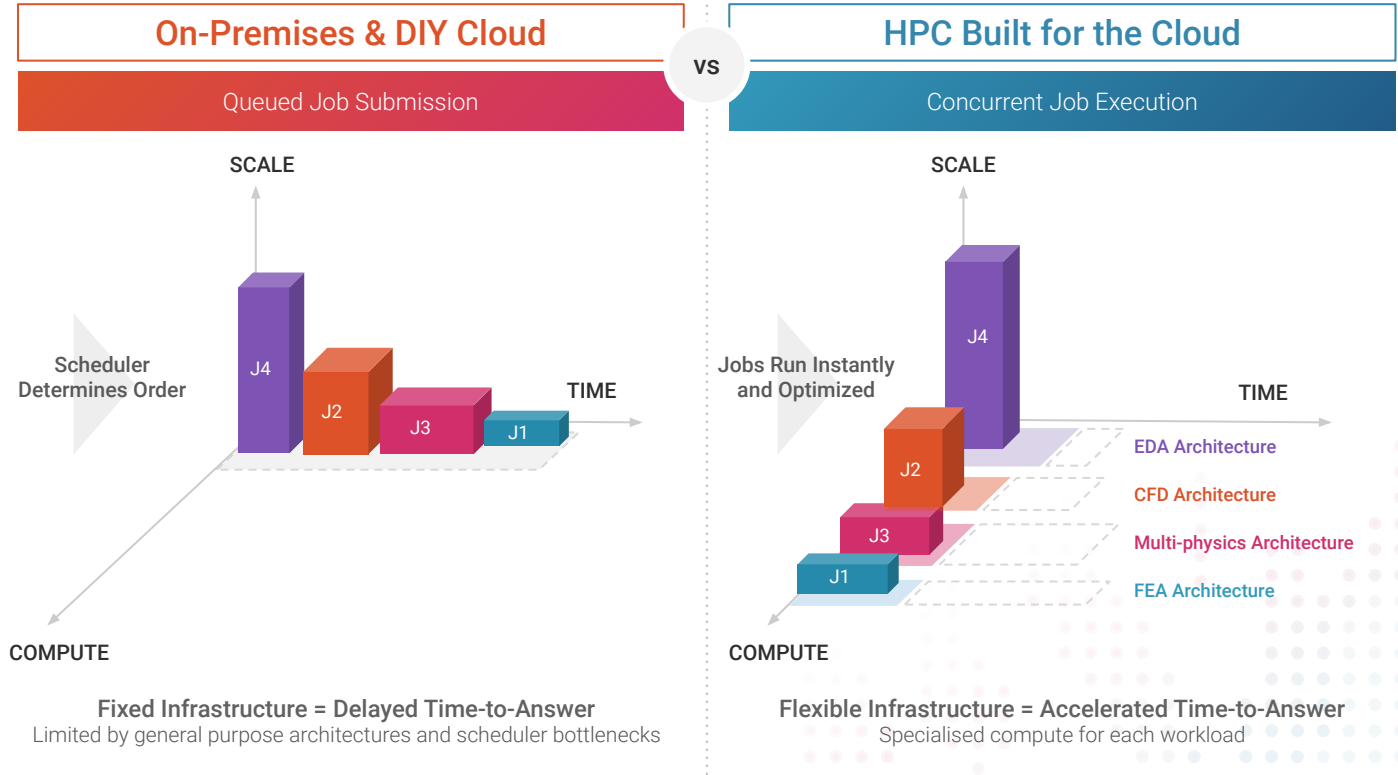
**Job 2**  
CFD - high interconnect requirement



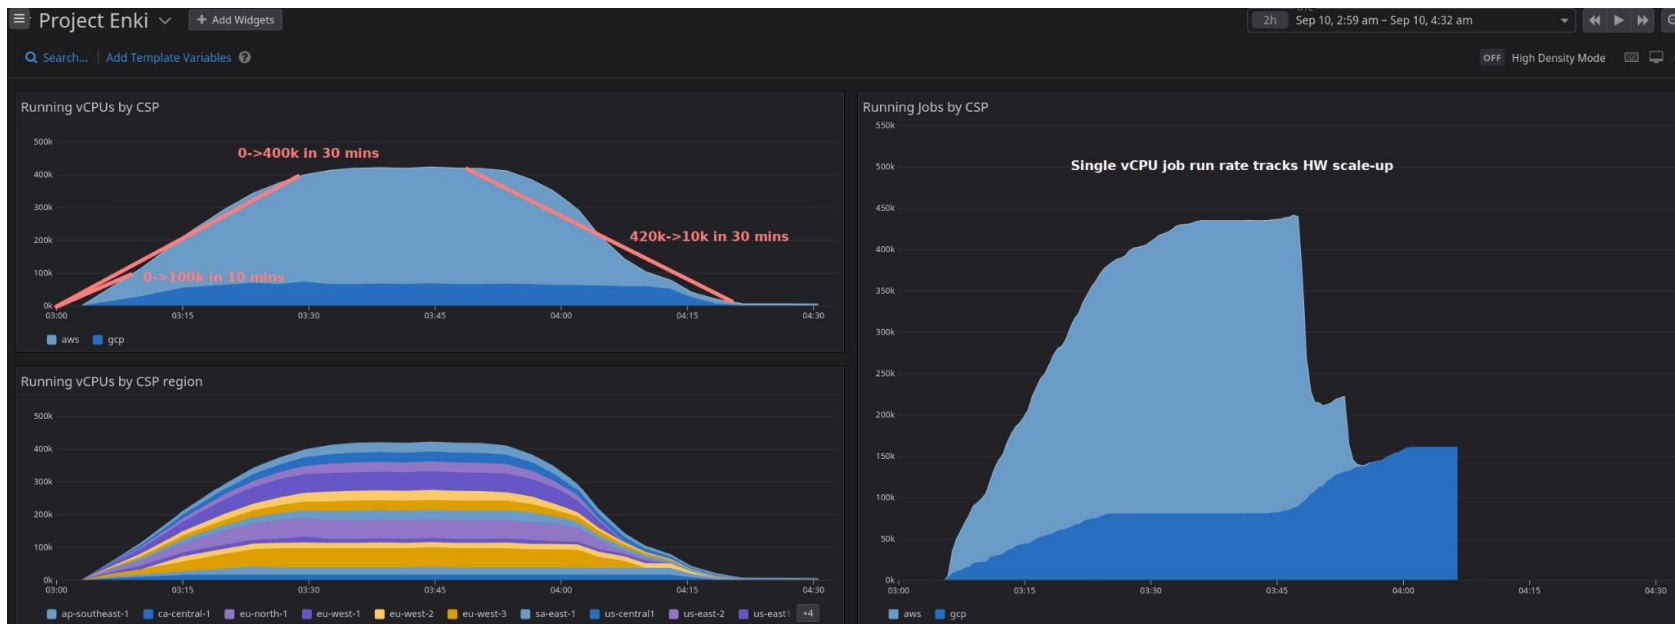
**Job 3**  
Multi-physics - large RAM requirement



**Job 4**  
EDA - Highly Parallelized Large # of cores



# Compute Capacity for Jobs: 400k+ vCPUs



- Fully managed, auto-scaling, multi-cloud HTC solution
  - Ingest and queue millions of jobs at a time
  - Currently peaks at 400,000+ concurrent jobs
  - Running on spot and preemptible VMs to minimize costs

# Rescale Software Publisher

## RESCALE SOFTWARE PORTFOLIO

*Any Commercial, Open-Source, or Custom Software*



## Overview

- Rapid, self-managed software publishing
- Authorized users can publish any custom, in-house, or commercial software
- Published software is available privately and securely for use by all or specific users/teams

## Benefits

- Increased productivity, agility, and consistency across R&D computing
- Centralized access and management of all HPC applications

---

# IT Management

Unified visibility and automated controls across performance, budgets, and security and compliance

## Topics

- Rescale Platform Intelligence
- Compute Recommendation Engine (CRE)
- From AI-Driven Insights to Action and Impact

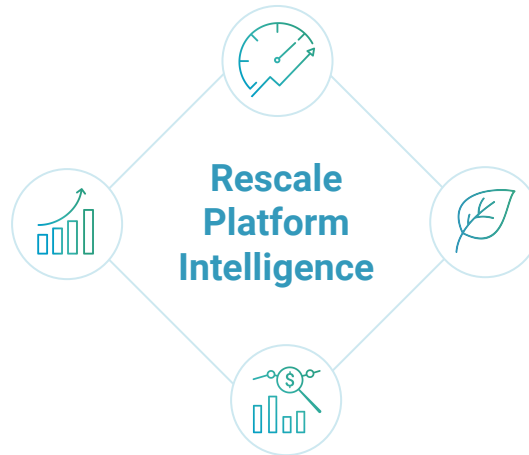
# Rescale Platform Intelligence

## Performance

Benchmark data on application and hardware performance across CSPs and available hardware types & configurations

## Maturity

Infrastructure capacity and reliability data across CSP regions and hardware types



## Sustainability

Environmental impact across available global data center infrastructure

## Cost

Pricing data across CSPs, hardware types, and service levels




# Using Rescale Platform Intelligence - Examples

A data-driven approach to optimizing workloads based on performance, maturity, and cost intelligence.

## Example A


A workload where an AWS coretype provides competitive per-core performance and value, with unmatched capacity



Source: Rescale Cloud HPC Platform Analytics, Using AWS list prices, Workload based on Caravan benchmark on LS-DYNA

## Example B


A workload where an Azure coretype delivers best per-core value and per-core performance while offering competitive capacity



Source: Rescale Cloud HPC Platform Analytics, Using Azure list prices, Workload based on Oil Rig benchmark on Ansys Fluent

## Example C

A workload where an OCI coretype delivers highest per-core and cost performance

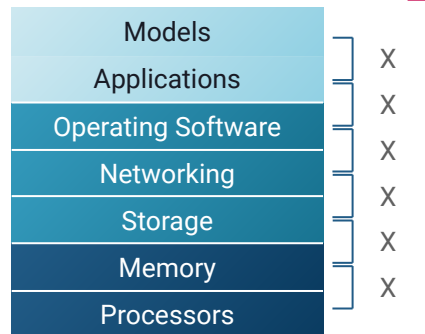


Source: Rescale Cloud HPC Platform Analytics, Using Oracle list prices, Workload based on Molecular Dynamics LJ Liquid 2M on LAMMPS

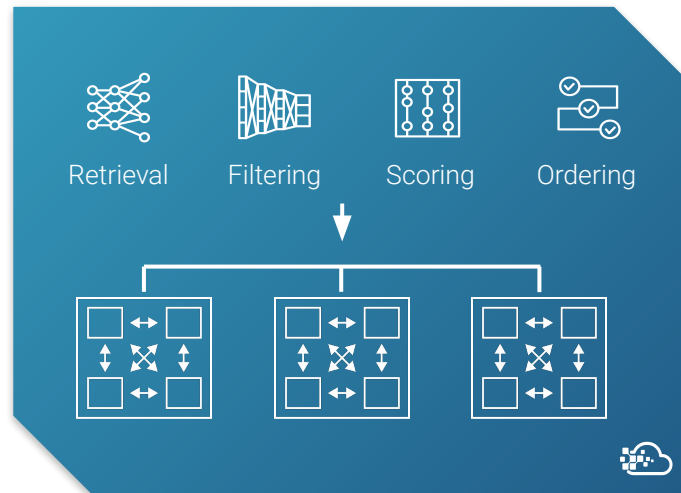
# Compute Recommendation Engine (CRE)

Performance optimization for any workload on any infrastructure

Full Stack Metadata & Telemetry  
Across Benchmarks & Jobs




Compute Recommendation Engine



User Creates New Job

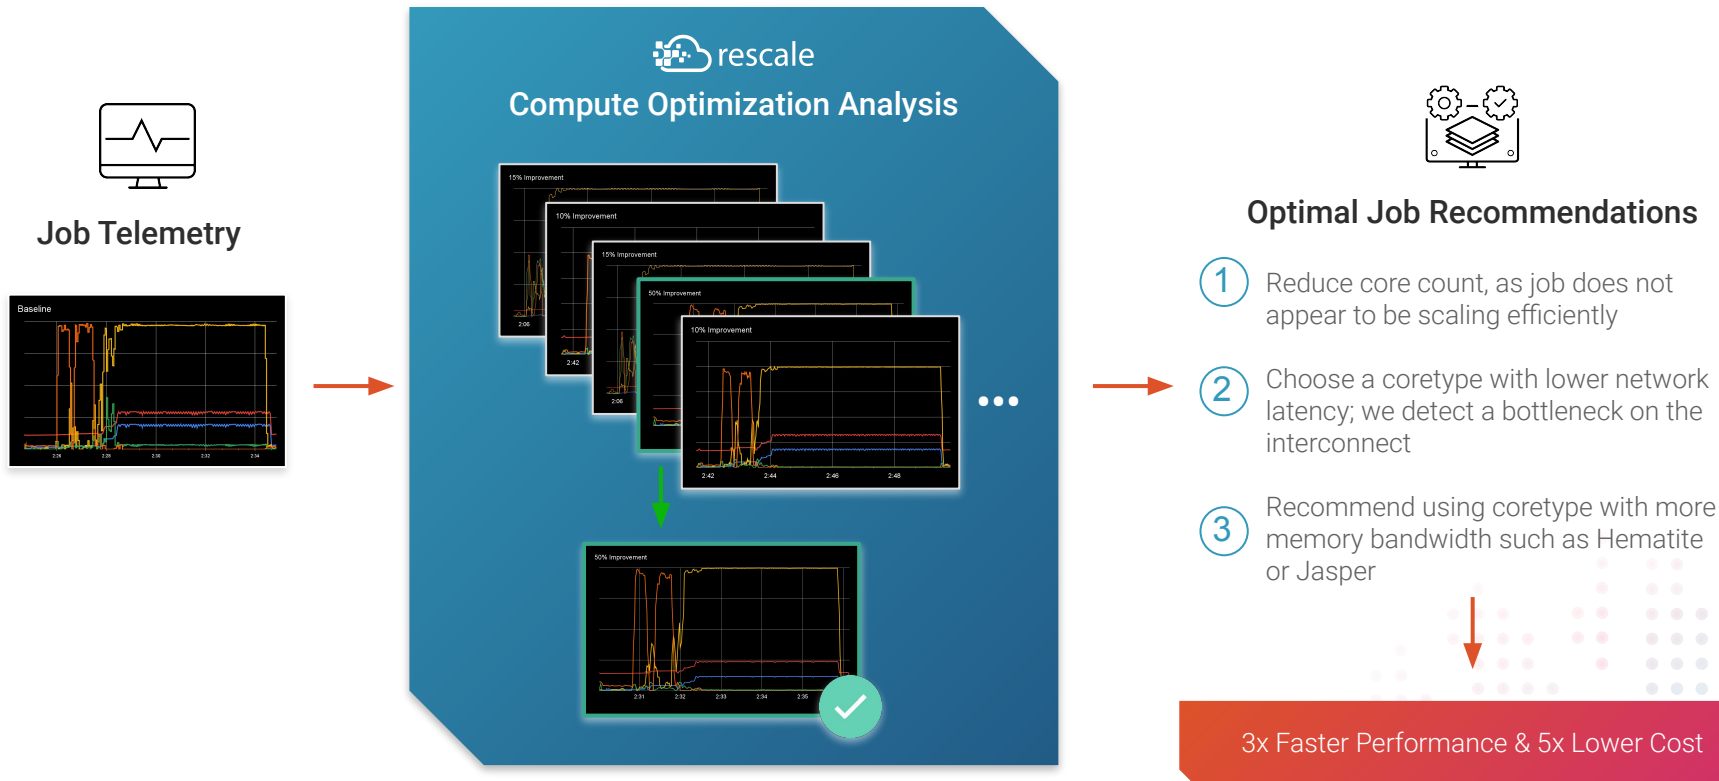
Compute Optimizer

Optimal Stack  
Recommendation



90%+ accuracy in identifying optimal architectures and optimal scalability

# From AI-Driven Insights to Action and Impact



---

# Digital Engineering

Powerful HPC+AI automations, interactive workflows, and seamless collaboration across the R&D stack

## Topics

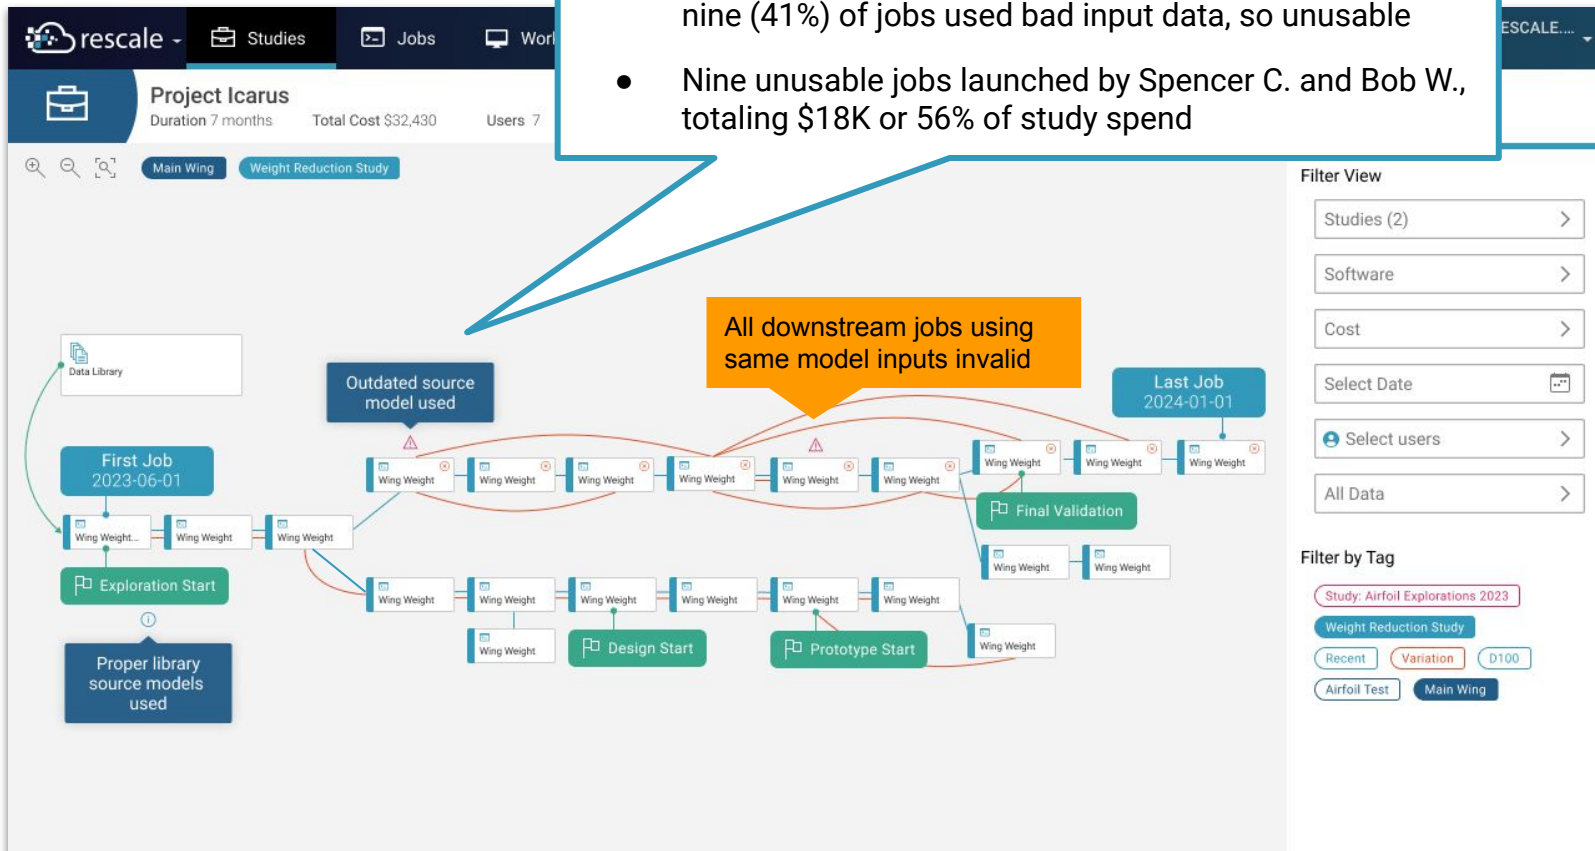
- Process Analytics
- Computational Pipelines
- Data Management for Cloud HPC

# Process Analytics

Tech Preview

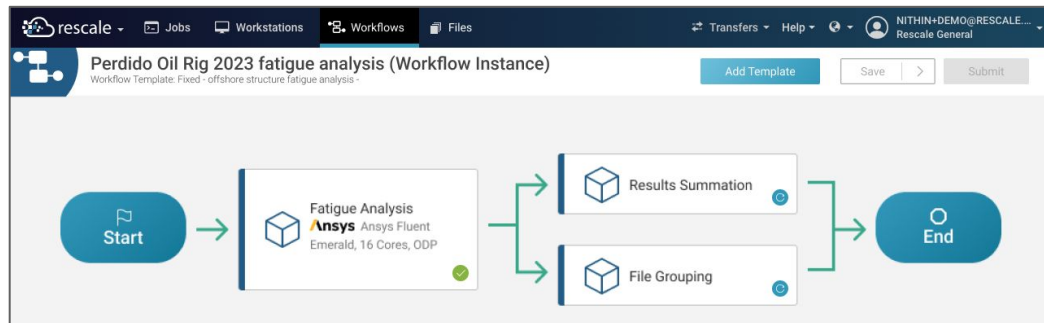


- In the “main wing - weight reduction study,” nine (41%) of jobs used bad input data, so unusable
- Nine unusable jobs launched by Spencer C. and Bob W., totaling \$18K or 56% of study spend

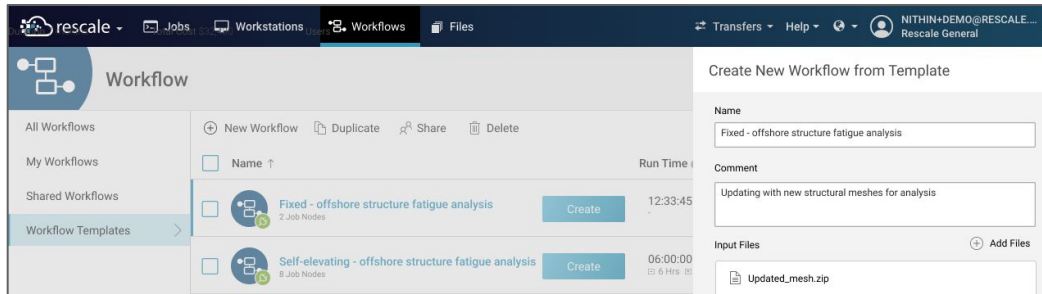


# Computational Pipelines

## Composable Workflow Template



## Configure and Orchestrate Workflow



Workflow

Create New Workflow from Template

Name: Fixed - offshore structure fatigue analysis

Comment: Updating with new structural meshes for analysis

Input Files: Updated\_mesh.zip

Workflow Name	Job Nodes	Run Time
Fixed - offshore structure fatigue analysis	2 Job Nodes	12:33:45
Self-elevating - offshore structure fatigue analysis	8 Job Nodes	06:00:00

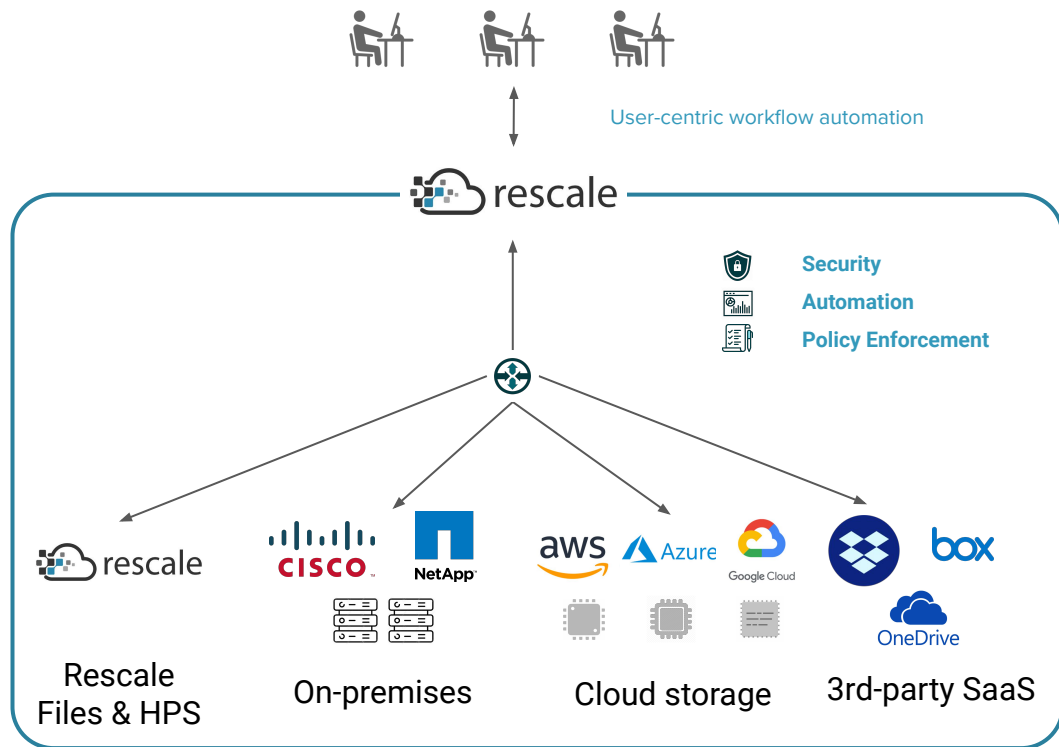
## Features

- Connect jobs into workflow to model sophisticated R&D processes

## Benefits

- Minimize the manual effort to complete a workflow and focus the talent on real problem solving
- Use right sized hardware for different jobs under a workflow for best economics
- Achieve high degree of consistency in R&D process for predictable outcome

# Rescale Optimizes Data Management for HPC



Aggregated interface for integration, management, and security

- Single pane of glass for visibility and management (permissions & policies)
- Seamless data integration & transfer by bridging data sources across on-premises, cloud, and SaaS

Data-driven performance & economic optimization

- Intelligent data & storage routing for compute performance and economics



## High Performance Computing Built for the Cloud



Digital  
Engineering



Workload  
Optimization



Intelligent  
Automation



Security &  
Compliance



# **COFFEE BREAK**

**We start again at  
11:30**

