Adaptive COMPUTING

Intelligent Workload Management Software for HPC and Cloud Environments

Best of Breed HPC Tools on Torque & Slurm

CUG 2017 Nick Ihli

5/11/13

Moab Intelligence Engine + Torque

Moab HPC Suite

- 15+ years battle tested
- Patented (70+ Patents)
- Mimics real-world decision-making

Multi-dimensional Policies Optimize Across:

- Workload/application requirements
- Priorities and SLAs
- Time (real-time and future, predictive)
- Heterogeneous resources

- **Torque – Open Source Resource Manager**
  - Executes decisions made by Moab

Product Offerings

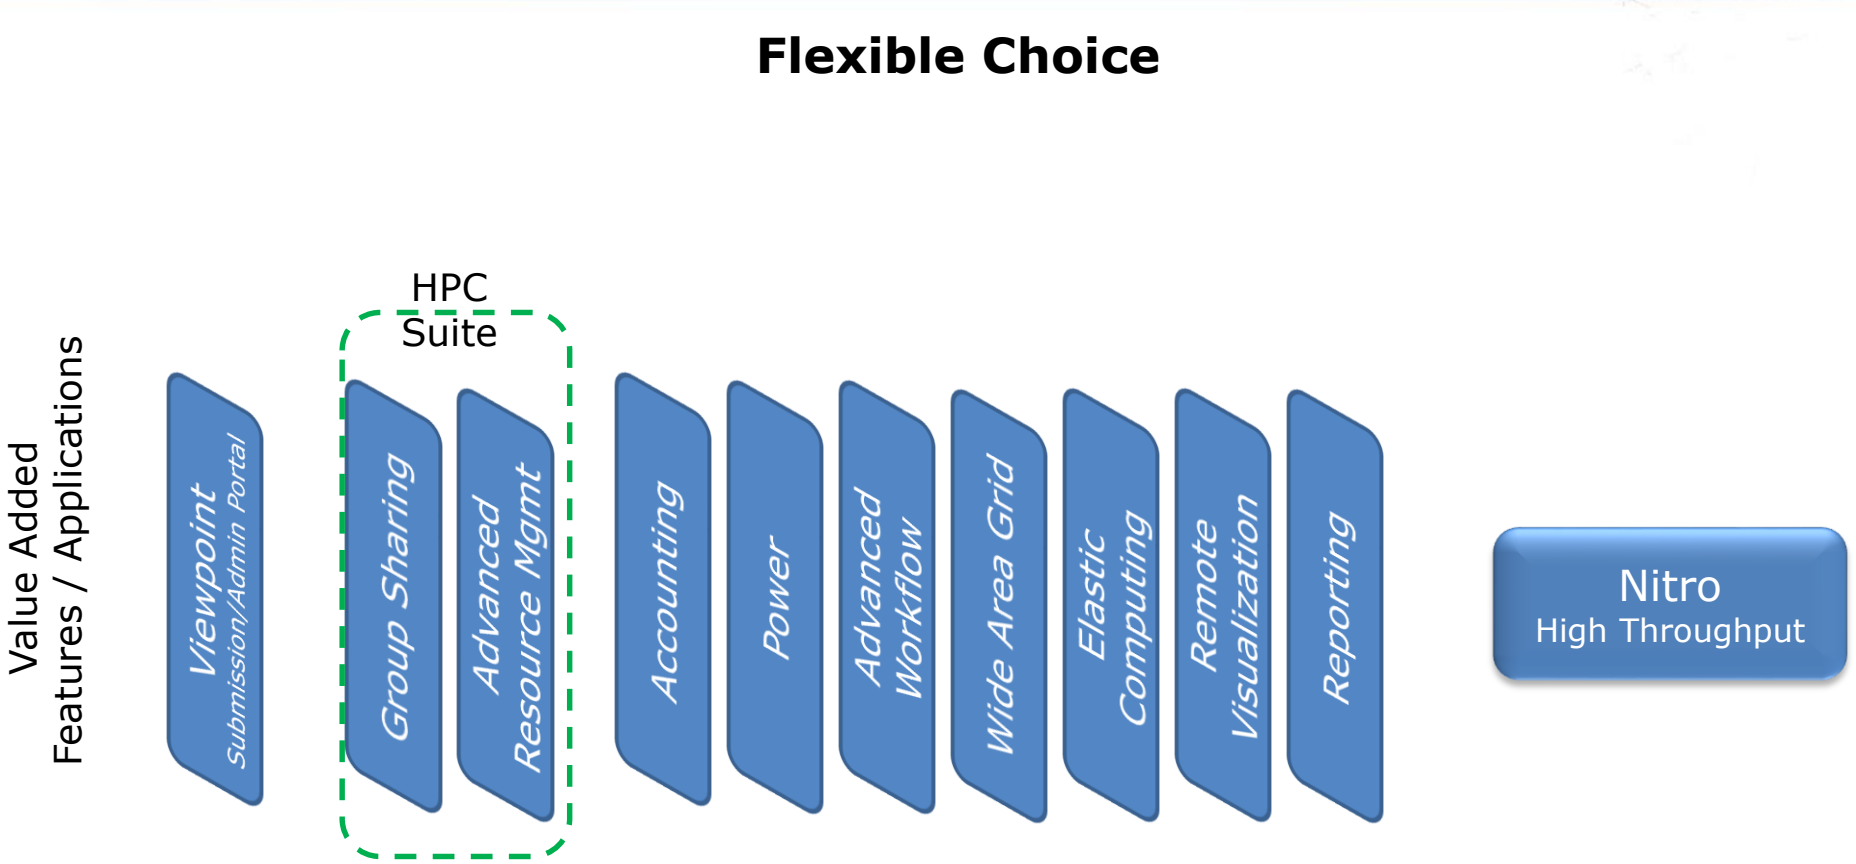

Adaptive COMPUTING

Moab/Torque Unification

- **Moab directly communicates with Torque node daemons (pbs_mom)**
- **A faster communication framework**
  - Faster:
    - Submission
    - Job start and teardown
    - MOM communication
    - Performance at scale
- **Eliminates race conditions**
- **Simplifies Everything!**
  - Synched JobIDs
  - Configuration
  - Debugging
  - Etc.

Enhanced Power Management

- **Workload-aware Power Management**
  - Reduce power state of idle nodes
  - Maintain high response times with green pool buffer policy
    - Set quantity of available nodes in power-ready state "buffer pool"
    - Dynamically increase power state of power reduced nodes to maintain "buffer pool" target
- **Per Application CPU Speed Throttling**

(New Initiative with Department of Energy)

- ***System Power Cap, Floor and Power Ramp Management***
  - Dynamic control over system energy consumption
    - Conform energy use to electrical power contracts
  - Job-level energy use data

Other Notable Features

- **Datawarp Integration**
- **KNL**
- **Docker Integration**
- **Singularity Integration**
  - Coming the summer

Open Platform

Get the **best-of-breed HPC tools** on **your choice of scheduler**. "Open Platform" enables organizations to **unify** the user, admin, and manager experience across multi-scheduler environments on a per-service basis.

- **Best-of-Breed HPC Tools:**
  - **Viewpoint Submission Portal**
  - **Remote Visualization**
  - **Reporting & Analytics**
  - **Nitro High Throughput**
  - Moab Accounting Manager
  - Grid Management
- **Cross Scheduler/RM Support:**
  - **Torque**
  - **Slurm**
  - SGE/UGE
  - PBS Pro
  - LSF

Viewpoint
Ease-of-use Driven Productivity

Add On

User

- Increase productivity of end users with **easier** and **faster portal-based submission** of tasks/jobs
- Expand HPC user base to **include non-IT skilled personnel**
- **Automate best practices** information into submission process in order to speed submission, reduce error and optimize processing speed

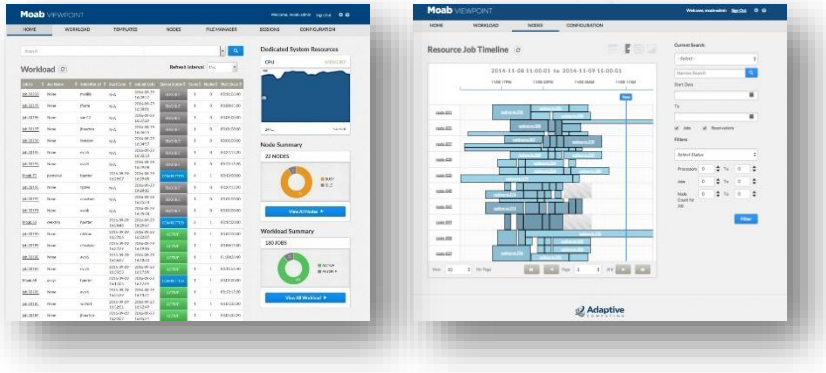

Admin

- Enable admins to **manage user requests with less time** due to best practices based templates, and user feedback which **helps users help themselves**
- Admins gain quick **visibility** into system and workload status and workload **troubleshooting**

Adaptive COMPUTING

Viewpoint
Ease-of-use Driven Productivity

End User Submission Portal

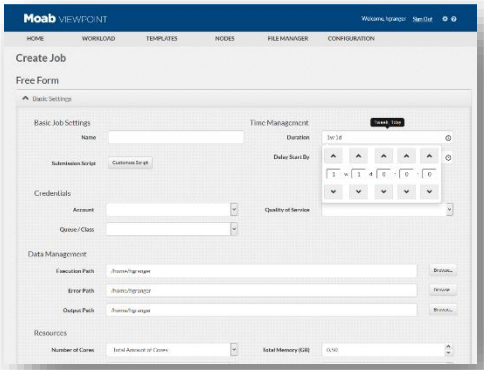

Application Template Form Builder

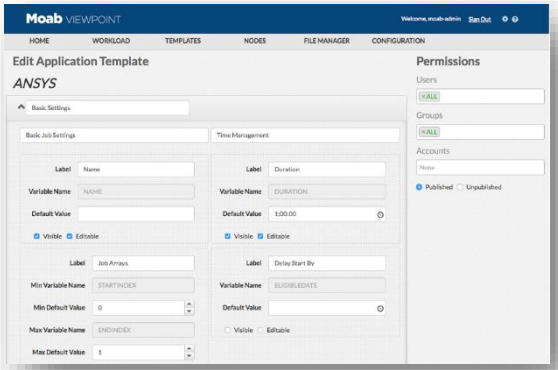

Self-help Oriented Job Details

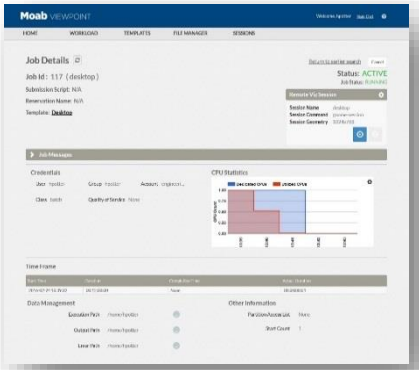

Simple Integrated File Manager

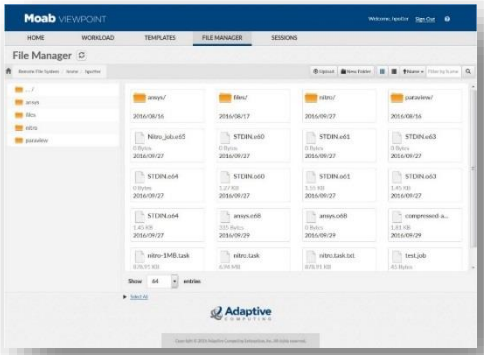

Administrative Reporting and Workload Status Tracking

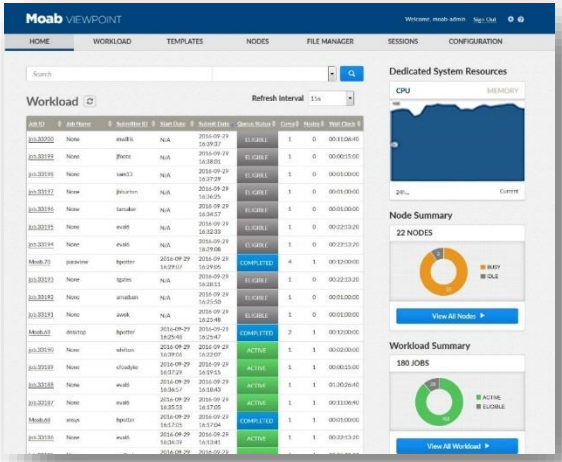

Resource Job Timeline

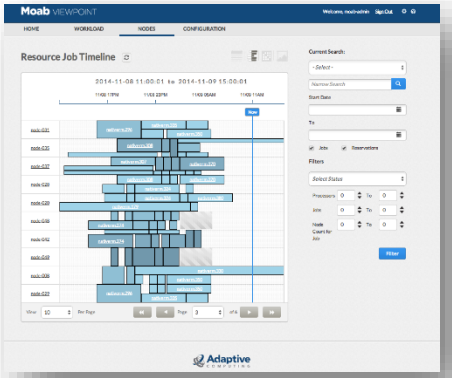

Remote Visualization

Avoid purchasing expensive licenses and GPU's for all worker, by rendering applications remotely and visualizing locally through an integrated portal.

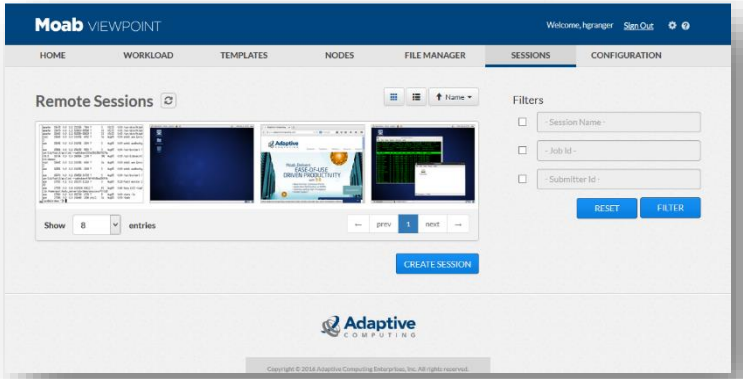

- **Improve Productivity**
  - Avoid waiting for file transfers
  - Collaborate on shared projects
  - Improve access to high value resources
- **Save Resources Integration**
  - Avoid purchasing individual high-end desktops
  - Share memory resources, accelerators, etc.
  - Share expensive application licenses
- **Improves Manageability**
  - Easier to launch, find, and manage sessions from submission portal
  - Schedule, prioritize, and guarantee SLAs between users and groups
  - Integrate into full HPC workflows, including notifications
  - Administrative controls over usage, security, sharing, etc.

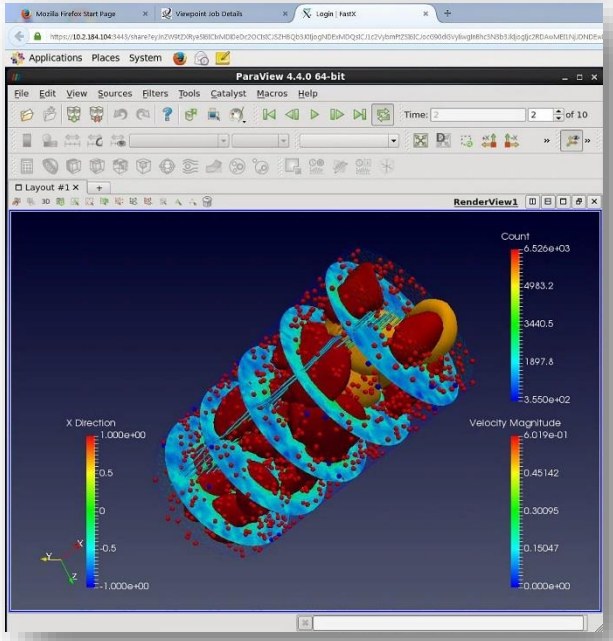

Nitro High Throughput

Accelerate launch times for short computing jobs by scheduling only once for large sets of jobs. At **500 tasks/second/core**, Nitro eliminates bottlenecks and puts high throughput performance into your computing system.

- **Supports** thousands to millions of tasks
- **Provides** simple user job submission
- Sessions can **grow** and **shrink**
- **Integrates** with all common scheduler
- **Applicable Workloads:**
  - Job arrays
  - Regression tests
  - Embarrassingly parallel workloads like Blast, Monte Carlo, and mass simulations
  - Anytime there are thousands and thousand of short run jobs

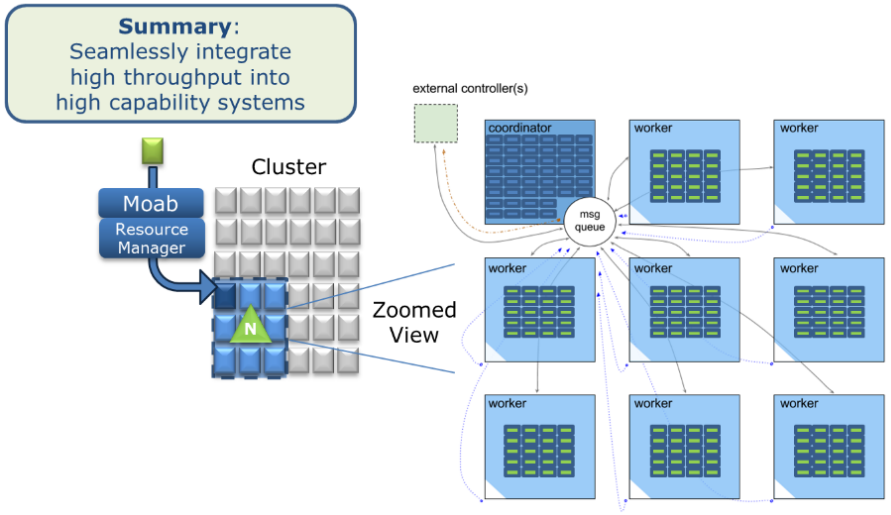

Reporting & Analytics

Gain insights by streaming resource usage and workload data into **custom reports** and **personalized dashboards**. This drives improved resource utilization and efficiency, **better capacity planning**, and greater **alignment of resources to mission objectives.**

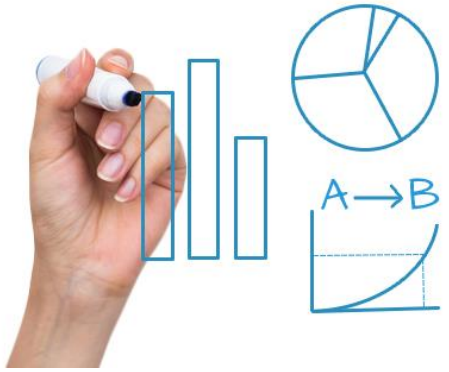

- **Stream** data with lightning-fast Apache Spark processing engine and flexible Mongo database.
- **Process** data to ensure relevance with traditional analytical functions (group, reduce, join, filter, etc.)
- **Visualize** aggregated information in chart-based reports or structured tables.
- **Monitor** key indicators in customizable dashboards
- **Drive** better decision-making & policy enforcement.

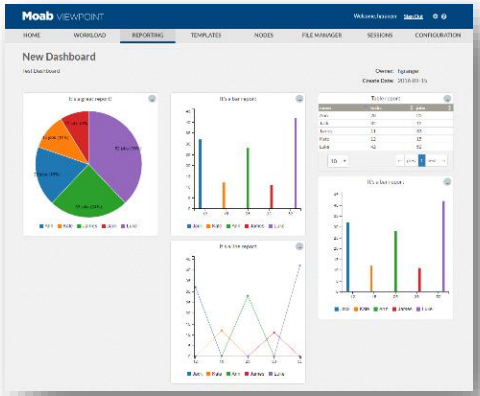

Reporting & Analytics

1. Data Stream Pipelines

2. Aggregated Views

3. Chart or Table Reports

4. Custom Dashboards

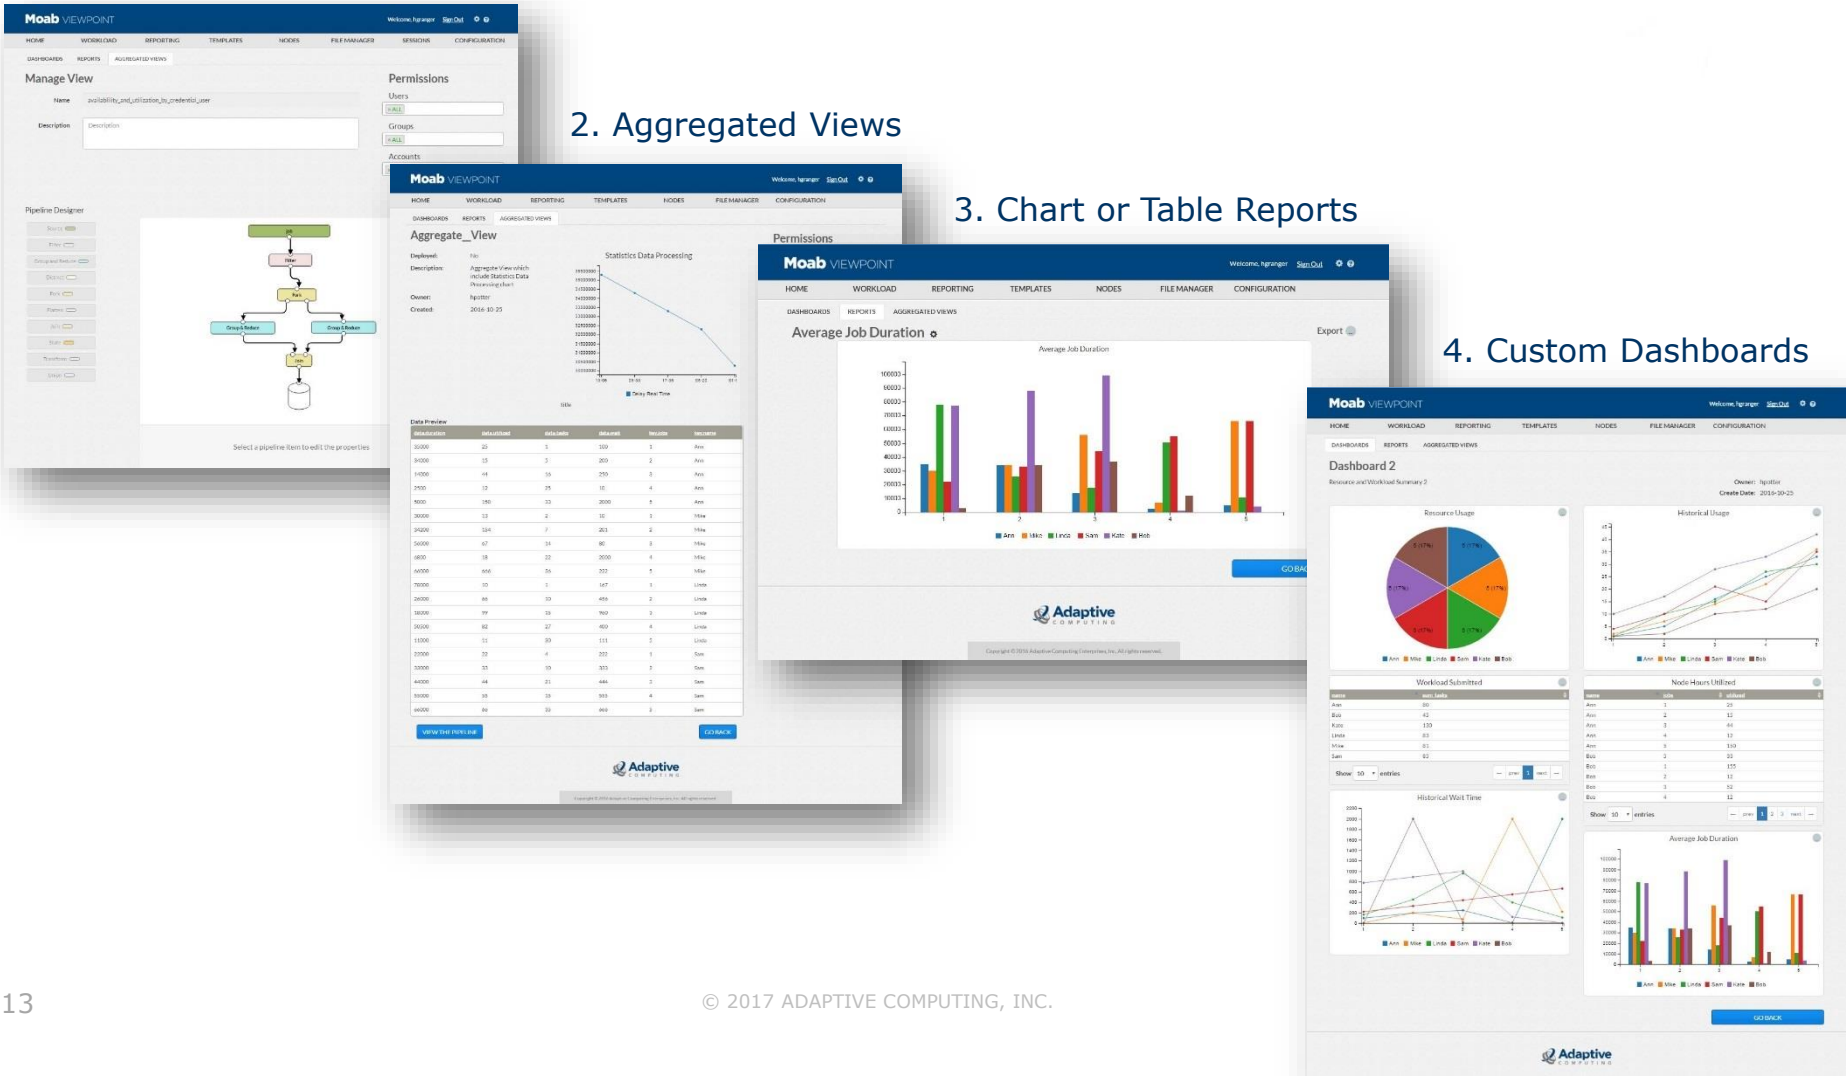

Public Grant Program
Buy Support – Get Licenses Free

Eligibility: Government and Education Organizations

- Education organizations must be diploma issuing entities
- Systems must not be used by commercial organizations (e.g. hosting)

Included Products: Moab HPC Suite, Viewpoint Submission Portal on Torque or Slurm

Engage Us

Nick Ihli

- nihli@adaptivecomputing.com
- +1 801 717 3736

Stuart Wright

- swright@adaptivecomputing.com
- +1 801 717 3371

Adaptive COMPUTING