



Ten Reasons to Switch from Maui Cluster Scheduler to Moab® HPC Suite

Comparison Brief

Many Maui users make the switch to Moab each year for key scalability, capability and support advantages that help them run a more efficient and optimized HPC system. This document provides a comparison of the additional value and capabilities Moab HPC Suite, Basic and Enterprise Editions, provides over the open source Maui Cluster Scheduler in managing HPC workloads. Moab HPC Suite is a policy-based intelligent workload management system that automates the scheduling, dynamic resource management, organizational priorities balancing, monitoring and reporting for HPC workloads on massive scale, multi-technology cluster and grid installations. Maui Cluster Scheduler is a basic open source job scheduling policy engine for clusters.

Moab HPC Suite offers ten significant benefits and capabilities over Maui Cluster Scheduler that organizations should consider as reasons to make the switch to Moab HPC Suite:

1. **Scalability for the growing needs of your cluster and organization.** Moab can scale to manage the workload of much larger environments as it handles significantly more nodes, jobs and users than Maui Cluster Scheduler. Moab easily scales to manage over 15,000 nodes, hundreds of thousands of queued job submissions, and over 500 users. The massive scalability of Moab has been proven as it currently manages workloads for about 40% of the top 10, top 25 and top 100 on the Top500 list. Organizations will encounter issues and performance problems with Maui job scheduling if they exceed 2,500 nodes or 15,000 jobs with as few as 100 users. Switching to Moab HPC Suite can eliminate these issues and enable more jobs and workload to be run on your cluster by more users without frustrating and costly errors, delays and restarts.
2. **A richer set of flexible policies that self-optimize workload scheduling and management for higher ROI.** Maui lacks many important policies to optimize workload throughput speed, resource allocation and prioritization. Moab HPC Suite offers a far richer set of policies to automate and optimize workload placement and resource management that better optimize throughput and utilization to speed the amount of results the system delivers. The Moab intelligence engine is more flexible and dynamic in its policy parameters and enforcement to deliver the behavior envisioned and desired rather than the rigid approximations that Maui provides. Some of the valuable policies customers can leverage by switching to Moab include:
 - a. Deadline based prioritization policy to meet service level guarantees and organizational priorities. Moab HPC Suite offers the ability for jobs to be prioritized according to their proximity to the deadline set for the job. This deadline prioritization can also be included in Moab Quality of Service (QoS) parameters. Deadline-based prioritization appropriately and continually adjusts priorities for jobs in the queue as their deadline gets closer. Making the switch to Moab means you can better meet and enforce service level guarantees and deliver the right results at the right time from your HPC system. Maui does not provide deadline-based prioritization.
 - b. Hierarchical fairshare usage prioritization policy enables more representative and accurate prioritization. Moab HPC Suite automatically determines the extent of a priority penalty or advantage applied for usage over or below usage target, but can also do this prioritization factor

within hierarchical fairshare groups or accounts organized by the administrator. Maui can only support basic fairshare where all users and jobs are considered together.

- c. More flexible and automated preemption policies based on a fuller context of the changing environment. Moab HPC Suite provides trigger and context based preemption that enables the preemption capability to be used in more types of situations than the limited Maui preemption which can only be done based on manual, Quality of Service (QoS,) and preemption based backfill. Moab can automatically mark a job as preemptible to get the right jobs completed faster based on additional configurable triggers such as jobs that have violated soft usage limit policies, fairshare targets or as the result of a generic event or generic metric specified by the administrator. Moab can also automatically mark jobs as a preemptor if the job's delivered or expected response time exceeds a specified threshold such as target queue time.
 - d. More flexible and granular Quality of Service policies to better deliver service level guarantees Moab HPC Suite provides a broader and more granular set of Quality of Service parameters that enables administrators to better control and deliver service level guarantees across multiple users, groups, and service levels. Moab provides for deadline proximity, ignoring of resource usage policies, how priority factor data will be tracked for jobs, and exceptions to the standard prioritization factoring. Administrators can set these once and then easily apply the full set of privileges, priorities and access to jobs based on user, group, or class. Maui does not provide these added service level and privileges, providing only a basic set of parameters.
3. **Professional services to keep your cluster up and running at peak performance.** Moab HPC Suite customers have access to 24x7 technical support, patches, and training from Adaptive Computing HPC experts to ensure questions and issues are resolved quickly so there is no disruption to or decline in performance of workloads on their cluster. As a bonus, support for the popular TORQUE job resource manager is also included with each Moab HPC Suite subscription so Moab customers can have a fully supported environment. Moab product training also ensures that customers maximize the value from their cluster investment and Moab investment by learning how to deploy, configure and use the policies and capabilities in Moab. Maui users must deploy, support and solve issues for Maui and/or TORQUE on their own, with help from the community as available, which often involves trial and error fixes which can delay jobs, results, and disrupt user and administrator productivity.
 4. **Broad and deep intelligent management and automation to reduce management complexity and burden while improving system efficiency with heterogeneous resource manager integration.** Moab HPC Suite integrates across multiple resource managers which enables vital data variables to be taken into account when making scheduling decisions. This broad and deep integration also enables Moab to orchestrate more automated actions to optimize resources to meet workload needs, reducing the management complexity and burden for administrators. Moab's integration also provides for resource manager diagnostics that can be auto-triggered. Key use cases enabled by switching to Moab include:
 - a. Integrate with provisioning managers to dynamically auto-provision OSES. Moab HPC Suite automatically determines when and orchestrates the re-provisioning of node OSES to better meet workload demand. This capability means jobs wait less, utilization goes up and admins don't have to make constant manual adjustments or maintain separate clusters for multiple OSES applications and codes need.
 - b. Integrates and resolves resource status data conflicts between multiple managers to make more accurate, efficient scheduling decisions. Moab HPC Suite automatically not only integrates the data from multiple resource managers managing and monitoring the cluster resources, but also has the intelligence to resolve conflicting resource status reports. Administrators can use the default Moab node state policy of pessimistic conflict resolution that enables it to determine a node is down if either resource manager reports a down state for a resource. An alternative optimistic policy setting is available that will allow Moab to determine and schedule a resource as

“up” if either resource manager reports it as up. This capability ensures more accurate and efficient scheduling decisions so jobs don’t get put on resources that are actually down or kept waiting when resources are actually up and available.

- c. Moab can work with license managers such as FlexLM to optimize license use and management. Moab HPC Suite can place workload based on nodes already running a license or determine floating license availability, suspending workload requests until the required license becomes available. In addition, Moab can also reserve licenses in conjunction with future jobs to ensure these jobs can run at the appropriate time.

Moab can also work with multiple network managers, power resource monitors, and many other types of resource managers to make other intelligent workload placement decisions and optimization actions. Maui can only integrate with TORQUE and SLURM as resource managers, which limits its placement policies to only that data that TORQUE can provide. Maui did have historical integration with LSF, PBSPro and SGE, but this integration has now aged and does not provide reliable support for the current versions of these workload managers.

5. **Uptime automation to maximize system reliability and productivity**– Moab HPC Suite provides several key capabilities to automate the uptime of your HPC system and the flow of jobs through it including high availability and auto-recovery trigger policies. Moab provides for high availability with a synchronized Moab head node for immediate fallback. Moab also initiates auto-recovery to resource, scheduler or resource manager incidents and events using custom trigger policies to further maximize job and system uptime. For instance, with Moab customers can create a trigger policy that automatically notifies of a node going down, executes a diagnostic command, and uses the results of the diagnostic command to trigger a reboot or restart of the failed node so it can begin to accept workload again. Moab also undergoes much more stringent quality testing as a product which means greater stability and fewer restart requirements. Maui does not provide any of these capabilities. This leaves Maui managed systems more susceptible to downtime, high administrator management and response, and ultimately being less productive and disruptive to the productivity of its users.
6. **Rich administrator dashboard tools for more efficient management and ROI reporting.** Moab HPC Suite provides an administrator dashboard that visualizes and filters cluster resources and status, enables visual policy management, and provides graphical charts and reports on utilization and workloads. These graphical management tools make management, issue diagnoses and performance optimization much more efficient. Maui Cluster Scheduler does not provide a graphical cluster workload management dashboard or reporting which makes management time consuming and complex for administrators. Maui’s limited management, policies, and formatted text outputs must be done via text-based commands, without the option for XML, using nearly 300 pages of documentation.
7. **Simplified HPC job submission and management for users to increase productivity and decrease the training and support burden on administrators.** Moab simplifies job submission and management for users with capabilities like advanced job array management, job templates, and a unified self-service submission portal to submit jobs without training, anytime, anywhere. Job templates make it quick for users to submit common jobs, avoiding the repetitive work to specify resources and requirements. Maui Cluster Scheduler provides no job array capabilities of its own, but can only leverage the basic job array submission provided by a resource manager like TORQUE. This also means that Maui has no job-array level policies or master job level management to simplify the ongoing management of the jobs for users and administrators. Maui also provides no job templates for users.
8. **Ongoing HPC workload management enhancements and development to enable organizations to adopt new HPC technologies and capabilities to maximize system productivity.** Moab HPC Suite has over 11 years of active development and a dedicated engineering and product management team to continue to deliver the new enhancements and capabilities as HPC needs and complexity continues to evolve. This commitment to development helps Moab customers get the most out of their cluster investments and

continually improve the results they deliver their organization. Maui Cluster Scheduler has not had significant enhancements in the past 10 years and has not kept up with new HPC productivity advancements. An example of this includes:

- a. GPGPU scheduling and policy-based management to maximize their utilization and ROI. Moab has added capabilities to allow for workload to be effectively managed across hybrid systems where GPGPUs have been adopted on some nodes. Customers who switch to Moab can leverage GPGPU productivity acceleration and maximize their utilization of the GPGPUs with auto detection and scheduling of GPGPUs to ensure the jobs that need and can leverage the GPGPUs get them first. Moab also provides policy-based scheduling of GPGPUs, dynamically making decisions of the optimal GPGPU resource to use for a specific job based on utilization, memory, temperature, etc. to ensure job needs are met and avoid hardware errors or failures.
9. **Grid and multi-cluster workload management capabilities to maximize utilization and throughput while consolidating and reducing the administration burden.** Moab HPC Suite supports both local area grid and wide area grid to enable unified workload and policy management across multiple clusters. This improves the utilization across the multiple clusters as well as the job throughput and results speed by giving users a broader pool of resources their jobs can be placed in. Maui cannot provide any of these grid or multi-cluster management use cases and capabilities. Customers who switch to Moab can use powerful and flexible multi-cluster and grid management use cases such as:
- a. Flexible management options for grids and multiple clusters management. Moab HPC Suite can be flexibly configured to manage globally, global and local, or local peer-to-peer for grids or multi-clusters. This flexibility enables organizations to use the level of workload management and policies that meet their needs in consolidating workload management across multiple clusters, while still protecting the sovereignty of the individual clusters if needed.
 - b. A global queue where all jobs can be queued or submitted for multi-cluster placement. Moab HPC Suite enables all jobs to be submitted to or consolidated into a global queue with job translation allowing jobs submitted in one resource manager language to be executed on a different resource manager on a different cluster.
 - c. Unified reporting of resource usage and the state and status of workloads and resources. Moab HPC Suite provides consolidated and unified reporting of workload and resource usage and status across multiple clusters to enable faster and more efficient management of workload across the multiple clusters or grid.
 - d. Data migration and staging. Moab HPC Suite provides for the migration and staging of data across multiple cluster locations that works with data management services, such as scp, NFS and gridFTP, and storage managers. Moab's optimized data staging ensures that remote data transfers between clusters are synchronized with resource availability to keep throughput and utilization rates high.
10. **Enforce departmental budgeting and provide pay-for use showback and chargeback to ensure SLAs are met, influence usage, and recapture usage costs to reinvest in capacity to keep up with growth.** Moab HPC Suite – Enterprise Edition provides usage accounting management that automatically schedules resources in line with resource sharing agreements and usage budget allotments, (%age, credits, resource amounts), for specific users, groups, accounts or projects. Customers who switch to Moab can use this capability to enable users or groups to pre-pay/budget, make installment payments/budget adjustments, or post-pay for cluster resource usage with the option to make those budgeted resources available all at once or phase availability in and out over the time period of the budget. Moab Accounting Manager also tracks and reports actual resource usage and budget amount usage by each user, group, account, etc. for showback or chargeback. Moab auto-determines appropriate charges based on flexible charge rates and

multipliers set-up by the administrator (resource types X usage amount + any multipliers and fees for things like service level). Maui does not provide this integrated use case and set of capabilities.

Capabilities Comparison

The following chart presents a detailed comparison of the major capabilities of Moab HPC Suite to Maui Cluster Scheduler for managing HPC workload in key customer value and use case areas.

Key Value Propositions and Enabling Capabilities ▼	Moab HPC Suite (all editions)	Maui Cluster Scheduler
PRODUCTIVITY ACCELERATION		
1. Massive Scalability: Solution has real world performance to efficiently manage (throughout and responsiveness) 500,000-1,000,000 jobs, petaflop systems of 100,000+ cores, 10,000+ nodes, 1,000+ users. This includes high throughput computing capabilities.	✓	No
2. Workload-optimized allocation policies: Matching and placement of job workload on optimal resources according to the job, code or applications' full granular needs and to maximize resource utilization (<i>timeline availability, network, processor, memory, storage, location, cost, software licenses, GPGPUs, load level, etc.</i>). This includes policies that can consider and even prioritize multiple allocation dimensions and attributes at one time.	✓	Limited <i>Only based on basic resource attributes and timeline availability. No GPGPUs, software license, etc.</i>
3. Workload-optimized provisioning: Policy-based dynamic OS provisioning on-demand to meet the OS needs of workloads in the queue via preferred provisioning tool (<i>dual boot, diskful imaging and diskfree imaging as well as multi-Linux and Windows OSes</i>).	✓ <i>Enterprise Edition</i>	No
4. Unified workload management across heterogeneous clusters: Workloads are automatically managed across heterogeneous cluster resources as if one optimized cluster, even if in siloed clusters (<i>mixed hardware/processors, multiple operating systems, (multiple flavors of Linux and Windows), GPGPUs and non GPGPUs, multiple resource managers and types of resource managers, etc.</i>)	✓	Very Limited <i>Only heterogeneous nodes. No GPGPUs, multiple resource managers and types of resource managers, etc.</i>
5. Simplified HPC submission and control for both users and administrators	✓	No
a. Choice of command line interface or easy-to-use web portal for user job submission and control: flexible and easy job submission and management option via web portal to reduce user learning curve and IT administration burden	✓	No <i>Jobs must be submitted via a separate resource manager such as TORQUE, which provides only basic command line submission</i>

Key Value Propositions and Enabling Capabilities ▼	Moab HPC Suite (all editions)	Maui Cluster Scheduler
b. Integrates with Windows HPC Server 2008 Job Scheduler: Allows users to submit and manage jobs through Windows scheduler in a hybrid Windows and Linux environment so users don't have to learn a new interface and administrators don't have added burden of supporting user re-learning or re- training.	✓	No
c. Job array submission, management controls: Users can submit up to 100,000 sub-jobs with one submission, track and manage at master job level with job array usage limit SLA policies and optimal grid placement selection options (single or optimal multi-cluster placement)	✓	No <i>Only basic job array submission can be provided through a separate resource manager such as TORQUE. No job array usage and auto-cancel policies or optimal grid placement.</i>
d. Job templates to rapidly submit common jobs, codes and apps: Job templates specify the variety of needed resources from servers and storage, to licenses and memory, etc., needed for each job in a standard and consistent way to reduce duplicate work for common jobs, simplifying job submissions for users.	✓	No
e. Manage application parameters & metadata: Can collect metadata and parameters specific to an application at time of request and transfer them into the set-up process and application as needed to further automate the set-up process	✓	No
f. Administrator dashboard and tools simplify cluster management: visual cluster environment views and reporting charts on status, usage and utilization of workload and the cluster as well as management of policies.	✓	No
6. Optimized intelligent scheduling including workload packing and backfill policies that maximize utilization while working in conjunction with SLA & priority policies	✓	✓
7. Advanced scheduling and management of GPGPUs for jobs: Maximizing their utilization and output with auto-detection, policy-based GPGPU scheduling, and GPGPU metrics reporting	✓	No
8. Workload-aware auto-power management policies that consolidate workload, intelligently activate power-saving modes on and off in line with workload demand, and placing workload on resources with highest performance-per-watt first	✓	No
UPTIME AUTOMATION		

Key Value Propositions and Enabling Capabilities ▼	Moab HPC Suite (all editions)	Maui Cluster Scheduler
9. Intelligent resource placement to prevent job failures including granular resource modeling to job and code requirements, avoidance of at-risk resources, and node set/location considerations	✓	Very Limited <i>Only node based attributes modeled and matched. No node health/at-risk filtering or storage and network, etc. preferences</i>
10. Auto-response to incidents and events to maximize job and system uptime including configurable diagnostic and recovery actions to pre-failure conditions, amber alerts, or other flexible metrics and monitors across resources, resource managers, schedulers, etc.)	✓	No
11. Workload-aware maintenance reservation scheduling to reserve resources for planned <u>future</u> maintenance periods and maintain a stable HPC system, intelligently allocating workloads elsewhere to avoid disruption during the maintenance period	✓	✓
12. Basic implementation, training and 24x7 support included at no extra charge with product purchase/subscription to ensure fast time to value and system uptime	✓ <i>Enterprise Edition for included implementation & training</i>	No
AUTO SLA ENFORCEMENT AND BALANCING		
13. Department budget enforcement to schedule resources in line with resource sharing agreements and usage budget allotments	✓ <i>Enterprise Edition</i>	No
a. Usage tracking and accounting: Tracks resource usage by specific users and accounts and supports pre-pay budget, install payments/budget adjustments, or post-pay with real-time and historical showback reporting	✓ <i>Enterprise Edition</i>	No
14. SLA and priority polices to ensure the highest priority workloads are processed first, even as workload input continually changes	✓	Limited
a. Simultaneously factor multiple points of workload and user priority, and at differing weights, to ensure most accurate and real-world prioritization of workload (including job credentials, requested job resources, fairshare usage, service levels, job attributes, etc.)	✓	Limited <i>Missing prioritization factors such as deadline based prioritization</i>
b. Prioritize workload based on hard (maximum) usage limits or real-time soft usage limits (allow overage with reduced priority if resource not required by others)	✓	✓
c. Hierarchical Fairshare Usage Prioritization to enable granular management of the extent of the priority penalty or advantage applied for usage over or below usage target, and within each designated hierarchical group.	✓	No <i>Only basic fairshare policies</i>

Key Value Propositions and Enabling Capabilities ▼	Moab HPC Suite (all editions)	Maui Cluster Scheduler
d. Automate SLA policy configuration by applying Quality of Service (QoS) policies (special resource priorities and privileges) for classes of users, groups and departments to ease administrative burdens while ensuring SLAs are met	✓	Limited <i>Only basic service and priorities settings</i>
15. Continuous plus future scheduling to ensure priorities and guarantees are proactively met as conditions and workload levels change including future reservations and pre-emption	✓	✓
a. Permanent resource reservations: Permanently reserve resources for a user or account to meet business needs, while still enabling their utilization by other jobs when not in active use. Also includes floating reservations that reserve resources permanently at specific time intervals.	✓	✓
16. Control access to specific resources by user and account (departments or organizations) based on policies to enforce security and usage	✓	✓
GRID- and CLOUD- READY WORKLOAD MANAGEMENT		
17. Pay-for-use showback and chargeback capabilities to track actual resource usage with flexible chargeback options (<i>usage-based, allocation-based, fixed, different rates for different resources and service levels</i>) and usage reporting by user or department	✓ <i>Enterprise Edition</i>	No
a. Native Interface for billing integration: Provides a native interface for billing allowing you to use the billing and accounting solution of your choice, commercial or custom.	✓ <i>Enterprise Edition</i>	No
18. Manage and share workload across multiple remote clusters as a unified grid to meet growing workload demand or surges and maximize resource utilization (includes consolidated reporting, unified/integrated management across diverse resource managers and resources, popular security tools integration, global grid SLA policies, unified job submission, and optimized data staging between clusters)	✓	No
a. Flexible grid management structure to allow global grid policies and management, global and local cluster policies and management, or local/peer-to-peer clusters policies and management	✓ <i>With Grid Option</i>	No

Evaluating Moab HPC Suite

Customers can make the switch to the more scalable, capable and reliable Moab HPC Suite very easily and for a very manageable cost. To evaluate and experience the advantages of Moab first hand, request a Moab HPC Suite demo or evaluation software at <http://www.adaptivecomputing.com> or via email to solutions@adaptivecomputing.com.

If you have questions or would like a quote on Moab HPC Suite for your organization, please contact us and a solution advisor will assist you:

- North America, Latin America +1 (801) 717.3700
- Europe, Middle East, Africa +44 (0) 1483 243578
- Asia, Pacific, Japan, India +65 6597-7053