

UFM Adaptive Suite

Increase fabric workload capacity with Voltaire UFM™ and Moab Adaptive HPC Suite™



As HPC clusters become more pervasive and used for a wide range of applications, managing resources efficiently becomes a challenge. Traditional solutions focus on managing and scheduling compute capacity in a homogenous and fully connected environment and only monitor compute resource utilization. This approach leads to inefficient use of resources, poor application performance, and lack of visibility into the real bottlenecks.

In HPC environments, the fabric must be dynamically tuned and provisioned to reach linear workload scalability, persistent performance and proper application isolation. Diverse applications may have different requirements from the server, OS, software, storage, and network infrastructure, and in many cases the I/O has significant impact on application performance and run time.

Voltaire and Adaptive Computing (formerly Cluster Resources) offer the first and only integrated solution for managing resources that span compute, I/O, network, and applications.

Most existing clusters cannot use more than 50% of their network capacity, and have severe blocking and congestion leading to high latency. With UFM Adaptive Suite, users can double their overall bandwidth utilization, and see up to 100 times lower latency due to reduced congestion. In addition, high-priority applications can now be granted more fabric resources and avoid wait times due to fabric conflicts with other less critical jobs.

The solution enables users to:

- Automate fabric optimization to increase application performance
- Prioritize network and I/O resources to match application requirements and priorities
- Assure network isolation between workloads
- Correlate fabric monitoring with application workload performance

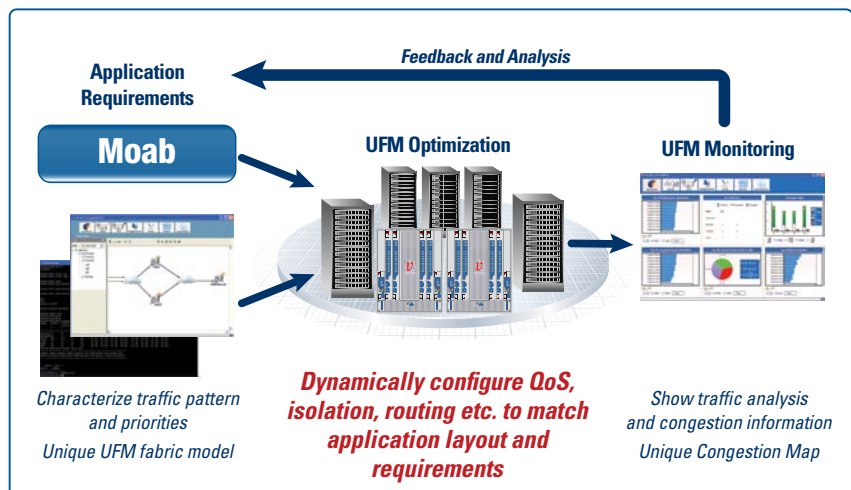
The solution consists of Adaptive Computing's Moab Adaptive HPC Suite intelligent, policy-based resource manager, Voltaire's innovative Unified Fabric Manager™ (UFM™) software, and Voltaire's smart fabric switching solutions. Based on an open architecture, it can be easily extended and customized by users.



Solution Details

The solution is composed of several integrated building blocks:

- Moab Adaptive HPC Suite intelligent, policy-based resource manager
- Voltaire switches and gateways
- Voltaire UFM fabric resource manager
- *Optional:* Voltaire OpenMPI Accelerator (OMA)
- *Optional:* Voltaire Fabric Collective Offload (FCA) add-on to UFM



UFM and Moab Adaptive HPC Suite Integrated Solution

Voltaire Unified Fabric Manager (UFM)

UFM is a central fabric resource management and monitoring solution that:

- Automatically discovers fabric resources
- Automatically provisions fabric resources to maximize performance, reduce congestion, and match application priorities
- Monitors fabric resources and traffic in real-time
- Provides a central administration and maintenance console

UFM enables users to model application and job requirements, and automatically configure the fabric to fit the application layout and unique requirements.

For example, UFM can dynamically control routing decisions, configure QoS and congestion management policy, configure Voltaire's unique switch-based collective offload, and partition fabric resources. As a result, performance increases significantly, congestion and latency is reduced by a factor of ten, and resources can be isolated and prioritized according to user needs.

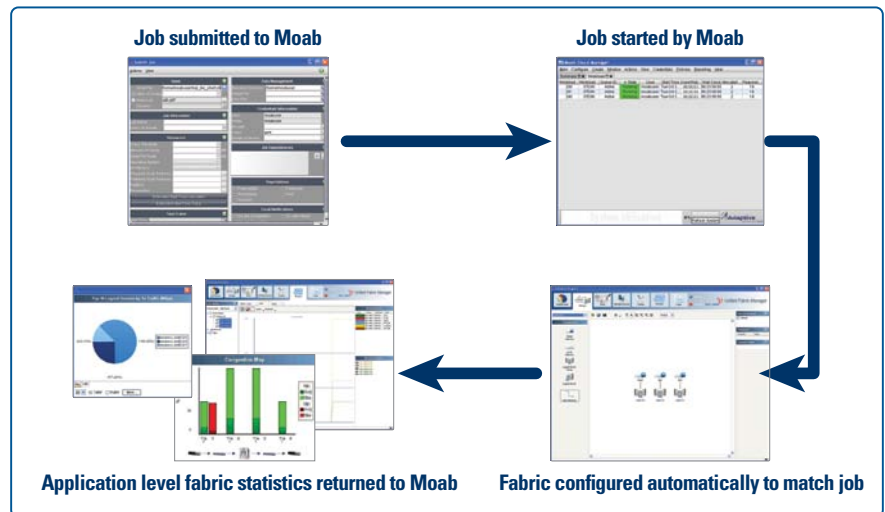
UFM incorporates a smart correlation and monitoring engine that constantly monitors statistics from hosts, adapters, switches, and third-party applications. The data is filtered, analyzed, and translated to valuable application-level statistics that can be used to identify application bottlenecks, fabric health and problem root cause. In order to scale, UFM leverages Voltaire's switch-based fabric management offload (using integrated line CPUs) to offload statistics gathering and analysis.

By integrating UFM with Moab Adaptive HPC Suite, the fabric can be managed as a dynamic resource and policy is automatically generated based on job layout (rank location) and requirements. High-level statistics can provide valuable feedback to the scheduler to further optimize resource allocation, and administrators can easily see relationships between resources and applications.

Moab Adaptive HPC Suite

Moab Adaptive HPC Suite is an intelligent, policy-based resource manager and scheduling system that allocates workloads to appropriate resources to enable end-to-end automation with minimal human intervention. Tasks and job submissions from users are prioritized and matched with the available resources according to organizational policies and high-level objectives. Moab can compare workload resource requirements with service level agreement priorities and apply corrective measures such as adding or removing a resource to an existing application environment or initiating the reprovisioning of resources with different operating systems and software stacks if needed. Moab's policy engine is flexible and allows for advanced reservations and prioritization schemes, which ensures that resources are available for key groups, departments and projects and prioritized appropriately.

When combined with UFM, Moab can manage additional resources such as software and software licenses, as well as a variety of software resource managers and server monitoring tools for a full resource management solution.



UFM and Moab Adaptive HPC Suite Integrated Solution