

Workload Management Systems: NQE and LSF Status and Plans

Jack Thompson, SGI, and Brian MacDonald, Platform Computing

ABSTRACT: *SGI is working with Platform Computing to incorporate key Network Queuing Environment (NQE) Features into Platform's Load Sharing Facility (LSF). Platform and SGI will jointly present the status of LSF for IRIX, UNICOS, and UNICOS/mk and discuss upcoming features in LSF 4.0.*

1. Network Queuing Environment (NQE) Transition

NQE 3.3 was the final feature release of NQE, SGI's proprietary workload management system. NQE will continue to be supported on SGI and Cray systems through December 31, 2004, with fixes for critical problems and call center support. It will be available for Cray SV1 systems. NQE on non-SGI systems is in retired mode.

2. Next Steps

When SGI was considering options for workload management systems after NQE 3.3, the company found that several independent software vendor (ISV) solutions were available in the marketplace. While NQE had been a stable, successful product for SGI, the company also knew that workload management software was not a core competency of the firm, particularly when considered in the context of a multi-vendor system environment. Clearly partnering with an ISV was the best choice for SGI and SGI customers.

As SGI considered the ISV solutions, Platform Computing's Load Sharing Facility (LSF) stood out as a market leader that was gaining wide acceptance. The decision was made to partner with Platform and offer LSF as SGI's workload management solution.

3. LSF Migration Program

SGI and Platform Computing recognize the level of investment that customers have made in NQE. To minimize the financial impact of the move to LSF, the two companies have developed a strong migration program.

Systems licensed for NQE before February 1, 1999 are eligible for steep discounts to upgrade to LSF. These discounts are available through January 31, 2000. In addition, SGI and Platform have jointly developed an *NQE-to-LSF Migration Guide*. The guide is available from the Platform Computing web site at:

<http://www.platform.com/Partners/SGI>

SGI and Platform also offer professional services for large and small data centers that would benefit from professional, hands-on migration assistance.

To further ease the transition to LSF, SGI has licensed a number of key NQE features to Platform Computing. Many of these features (described in Section 10 of this paper) will be available in LSF 4.0 which is scheduled for the fourth quarter of 1999.

The licensing of NQE features marked the start of a detailed technical exchange between SGI and Platform Computing. That engineering relationship has grown to include cross-company reviews of new technical plans, particularly in the resource management area. SGI and Platform understand the complex interdependencies among various system management technologies. The two companies are committed to a long-term strategic alliance that will result in tight integration of these key software products.

4. Current LSF Status on SGI and Cray Systems

LSF is now available for IRIX, UNICOS, and UNICOS/mk systems. LSF 3.2 for IRIX was released by Platform Computing in August, 1998. For UNICOS, LSF is available on Cray J90, Cray T90, and Cray C90 systems. LSF has been ported to Cray SV1 interim-rev A systems, and will be ported to Cray SV1 rev B systems as soon as hardware is available. The port for the Cray T3E was completed in May. Explicit support for MPI jobs on the T3E through LSF Parallel will be available as part of the LSF 4.0 release in Q4 1999.

LSF is also available for Linux and Windows/NT systems.

Three LSF products are available from SGI: LSF Standard Edition, LSF Parallel, and LSF Client. Customers who wish to purchase other LSF products (LSF Analyzer, MultiCluster, JobScheduler, and Make) can obtain these products from Platform Computing.

5. Creating Environments for High Performance Applications

With the LSF Suite of products, Platform Computing Corporation is focused on delivering software solutions that address the requirements for managing workload in high performance computing environments.

The LSF Suite of products consists of a layer of resource management software residing on top of UNIX and Windows/NT platforms. The LSF products extend the operating system services, creating a single system workload image across a network of computers. De-coupling the physical resources from the user, administrative and site workload requirements result in a number of clear benefits.

All administration and control can take place from any single location. Changes can be applied that affect the whole environment in a single operation. Along with the ability to control and manage the environment, users perceive a single logical entity rather than specific physical hosts.

Policies for control, management and performance goals can be applied uniformly to all computers participating in the environment. This allows for mixed workload to be scheduled within a single large machine or across clusters of heterogeneous platforms with the same level of complexity. The result is mechanisms to support capacity, capability, and response time-oriented applications mixed together in the same management framework.

6. Defining Capacity Goals

Providing an environment that can deliver on the throughput or capacity demands is a clear requirement for users of high performance applications. Achieving maximum resource utilization for CPU, memory, software licenses or other key indicators is essential to justify the investments in hardware and software.

Capacity-focused computing may come at a price. There may be more than one job running on a CPU leading to throughput being proportional to the number of tasks on that CPU. Depending on the operating system, near-capacity memory usage may result in high paging rates and drastically affect the turn around time of the applications. Good overall performance for the system may be potentially sub-optimal for applications. The result is that time-critical work may not be completed within an acceptable window of opportunity.

Examples

- A meteorological model has to be completed by the forecast deadline.
- Large multi-disciplinary codes, such as CFD, FEA, or Crash require fast turn-around time. It might take one week on a loaded system but potentially could be

accomplished overnight in an exclusive, multi-machine environment.

As work from multiple queues with differing priorities is executed in the cluster, thresholds can be used to transparently throttle the rate at which work is executed, accepted, and suspended. The advantage of this approach is that unlike using job counters, thresholds can be applied using any resource index supported by LSF. Job flow can be metered by metrics such as utilization, run queue length, free memory, paging, or sit- specific metrics, such as NFS ops.

The thresholds are applicable to queues or to hosts. This layered approach allows long running jobs to have different thresholds than critical, shorter jobs. The host thresholds could then be used to guarantee that the service levels of the host are within acceptable ranges. Thresholds are the backbone of LSF's capacity computing offering. Over-subscription mechanisms can be supported through extensible resource definitions and will allow sites to over-commit machines that work well under these scenarios (UNICOS).

7. Defining Capability Goals

Having a job, or series of jobs finish as quickly as possible using all necessary dedicated resources is clearly a capability computing goal. In order to do so, sharing and contention must be avoided at all costs. The result is that problems can be tackled that otherwise could not be considered, and mission critical applications can gain the undivided attention of the computing infrastructure.

By using value-added application partitioning mechanisms and a deep awareness of the underlying architecture, it is possible to drive complete solutions for maximum application performance, guaranteed across a distributed environment.

The marriage of capability and capacity computing, in the same consolidated environment, will provide the customer service goals in mixed mode operation on single large machines as well as clusters of large boxes and smaller departmental machines.

The goal is to support all facets of capability computing under the roof of one cluster. Multi-box MPI, threaded parallelism on Origin 2000 systems, and clean integration with UNICOS/mk will all open the door for mixing a wide variety of workloads without sacrificing any pre-stated customer goals.

A part of this strategy includes integration with Miser, Miser CPU Sets and playing a role with the Miser follow-on product, XRS, as they become available on future operating system releases.

8. Right-sizing Job Placement

Job placement or host selection decision is a match making process. Requirements or attributes of the job are used to accurately match with the available system environments.

Examples: free memory, light disk I/O, an Ansys license, NFS ops, or the availability of a Miser CPU Set

Sophisticated decision making on the order of the selection can be made automatically by the batch scheduler.

Example: `order[r1m:pg:io]`

If the job has a resource requirement of 2 GB of free memory, it can then use the selection of one-minute run queue, paging rates, and I/O rates as the ordering criteria for host selection.

Reservation can be applied to all indexes, but most importantly memory. Memory can be reserved to ensure jobs with uneven allocation profiles will have the free memory available when needed. Processor reservation can be used to generate sophisticated policies for parallel applications.

Placement policies include locality (SMP, spanning single processors), reservation, back-filling, deadline scheduling and striping (4 nodes by 16 per node).

The single system image works to separate or virtualize the execution environment from the user. By allowing the user, administrator and application to work in terms of resources, the scheduler can ensure that the correct placement policy is made while assuring the best possible availability.

The scheduler acts as the policy center for the applications. The queues act as repositories for policies in business terms. High priority, preemptive control with user limits, run limits, and time windows would be an example of a queue that captured the business policy. In this case it is a mechanism to allow one critical run per user at a time for work that needed to be accomplished in a short time period.

Hierarchical fair-share can be used to equally or unequally distribute the work that is run. Fair-share is a strong mechanism for adjusting group throughput goals to reflect business needs.

The scheduler will fail over to any of the server hosts in event of a failure of the scheduler host. This will happen transparently to the user, thereby ensuring the uptime of all machines in the cluster. Server hosts are responsible for collecting accurate, job level accounting information and returning it to the scheduler for accounting and real-time limit enforcement.

9. System Level Integration with Operating System Services

By integrating at the system level, applications that use MPT 1.3 will transparently benefit from the available resource management of LSF. LSF Parallel makes use of SGI Array Sessions to startup the MPI tasks. All subsequent operations make use of the LSF remote execution facilities to collect accounting information, apply limits, and control the tasks.

All applications that make use of MPI can benefit from the integration without changing their codes. This is very important for commercial Independent Software Vendors who may have lengthy release cycles and may not have the resources to tightly integrate workload management with their parallel products.

If checkpointing is supported by the application (*e.g.*, Fluent, custom codes) the appropriate signals can be sent to all tasks making up the parallel job. This removes a significant constraint from being able to deploy large-scale parallel applications in production environments.

10. Product Evolution for LSF 4.0

LSF 4.0 will be available in the fourth quarter of 1999.

Enhancing the scheduler has been driven by large multi-machine sites, including film production houses as well as Government sites such as LANL. By allowing all the policies to be in place with extreme scalability will allow the big sites to benefit for the full granularity of control.

All daemons will now be configurable without restart. This applies to hosts and resources being added, queues being modified, and batch policies changed. This requirement was especially critical for the large environments that were having many machines re-reading disk files and experiencing delays due to re-configuration.

The scheduler has now been multi-threaded to support a large number of client queries. The scheduler has also been enhanced to support chunked dispatch to large machines to support very fast turn around for short, transaction-oriented jobs.

Configurations for queues will now be changeable based on time. It will be possible to have a high priority queue with tight restrictions during the day that will automatically loosen the restrictions at night.

As part of the NQE feature support, the NQE File Transfer Agent (FTA) will be integrated as an optional job submission service. Input/output spooling and output re-direction will be fully supported. Administrators will now be able to specify work and spool directories. As a further enhancement, jobs may discover working directories if needed and pass this

information back to the scheduler for accounting and job tracking.

Hierarchical work-flows and extended dependencies will be supported for logically connecting jobs. An example would be a pyramid of portfolios for risk processing. Each level (consisting of many jobs) needs to be complete before the next stage of processing takes places.

All load information manager (LIM) configuration will be propagated by the master LIM. All hosts will automatically detect the type of host, operating system, model of host and if possible some form of performance indicator.

11. Summary

The close partnership between Platform Computing Corporation and SGI has resulted in some immediate short-term product benefits for user environments. This relationship will clearly align the direction of the LSF Suite of products with those of SGI products and ensure that Workload Management solutions will have an active an vital future for High Performance Computing applications.