



89 Fifth Avenue, 7th Floor

New York, NY 10003

www.TheEdison.com

@EdisonGroupInc

212.367.7400

HPC Workload Management Tools: Tech Brief Update

**IBM Platform LSF Meets Evolving
High Performance Computing
Demands**

Printed in the United States of America

Copyright 2015 Edison Group, Inc. New York.

Edison Group offers no warranty either expressed or implied on the information contained herein and shall be held harmless for errors resulting from its use.

This report was commissioned by IBM.

The information contained in this document is based on IBM provided materials and independent research and was aggregated and validated for Edison Group, Inc. by the Edison Group Analyst team.

All products are trademarks of their respective owners.

First Publication: November 2015

Produced by: Matthew Elkourie, Analyst; Manny Frishberg, Editor; Barry Cohen, Editor-in-Chief

Table of Contents

Executive Summary 1

Competitive Benchmark Study Recap 2

The Road Forward..... 3

Conclusion..... 6

Executive Summary

Evolution in technology is a constant for organizations in all business verticals. In a quest to achieve the quickest results and stay ahead of the pack, business leaders continually look for any advantage that gives them the edge.

In High Performance Computing (HPC), hardware improvements, both at the processor and peripheral levels, often take time to be fully realized. System latencies continue to decrease with improved hardware design elements as compiler updates and software enhancements are applied to ever-increasing hardware throughput and speed.

IBM® Platform™ LSF®, a powerful workload management platform for distributed HPC environments, is a prime example of this constant evolution and enhancement. With this in mind, Edison revisited IBM Platform LSF to see how, over time, change may have enhanced and augmented scheduler performance. With the expectation that elapsed time improvements would be shown over time, this Technology Brief update focuses on how the effects of time (hardware evolution) and code optimization can greatly enhance past results.

In the original white paper¹, published in July 2014, Edison took a wide view of different metrics for judging performance and agility. In order to present a flat field to gauge Platform LSF performance against its peers in the field of HPC schedulers, performance numbers had been generated without spending a large amount of time in refining either the approach or the product features.

Illustrating why IBM Platform LSF continues to be a good fit for HPC shops looking to accelerate production, IBM Platform LSF significantly improves its previously documented throughput benchmark results by over 300%. In addition, scheduler cycle timing decreased by a similar 300% as compared to previous results. Readers may recall that the initial benchmark results showed IBM Platform LSF dominating its peers by more than 400% in throughput, so the combination of improved throughput and latency performance in the updated results clearly shows a trend of constant improvement and dedication to the IBM Platform LSF solution.

¹ http://www.theedison.com/pdf/2014_IBM_Platform_LSF_WP.pdf

Competitive Benchmark Study Recap

The scope of the initial audit white paper was broad and deep, illustrating Platform LSF performance traits that allow users to leverage the agility and performance capabilities provided by the scheduler engine.

Edison discovered that Platform LSF was capable of up to 150x greater job throughput (with 1,666,666 jobs completed per hour) and up to 2800x faster in dispatching high priority workloads than its peers. In addition, Platform LSF dominated its competitors in the area of agility, with up to 4x more queries performed while responding up to 250x faster to job queries under load.

Workload Manager	Throughput ²	Agility ³
IBM Platform LSF	1,666,666	0.03 Seconds
SLURM	258,435	10.43 Seconds
TORQUE	11,000	84.93 Seconds
Open Grid Scheduler	138,089	14.63 Seconds

Figure One: Throughput and Agility Results Summary from Previous Study

Since that time, the Platform LSF development team has not been idle. With advancements in systems processors since the publication of the original white paper, and with additional tuning to the scheduler settings, a review of current performance is warranted.

Software updates and enhancements often coincide with hardware improvements. The systems used in the benchmarks for the original publication featured Intel Westmere processors. In testament to IBM Platform LSF agility and scalability, the latest tests utilize Intel Haswell processors, with more threads running at higher frequencies, alongside significant code optimizations that allow Platform LSF 9.1.3.3 to scale with technology and achieve higher throughput numbers.

² Measurement of jobs completed in one hour where a higher number of jobs completed is the desired result.

³ Measurement of scheduler responsiveness under load, demonstrating platform agility in response to potential workload changes or user input while the system is under load. A lower number is a desired result.

The Road Forward

In the updated evaluation, the Platform LSF team performed testing on a 16-node cluster, running Platform LSF version 9.1.3.3 to evaluate improvement on the prior throughput results of 1,666,666 jobs per hour.

While both data sets were obtained using the best configuration options available at the time, the newly obtained findings illustrate significant performance improvement due to continued scheduler software improvements. A continued commitment to performance and feature enhancement by IBM is evident in significant scheduler performance improvements in a relatively short time frame.

To illustrate the agility of IBM Platform LSF, Figure Two compares the previous benchmark results with the benchmark results that leveraged newer Haswell processors running on a later version of IBM Platform LSF (9.1.3.3).

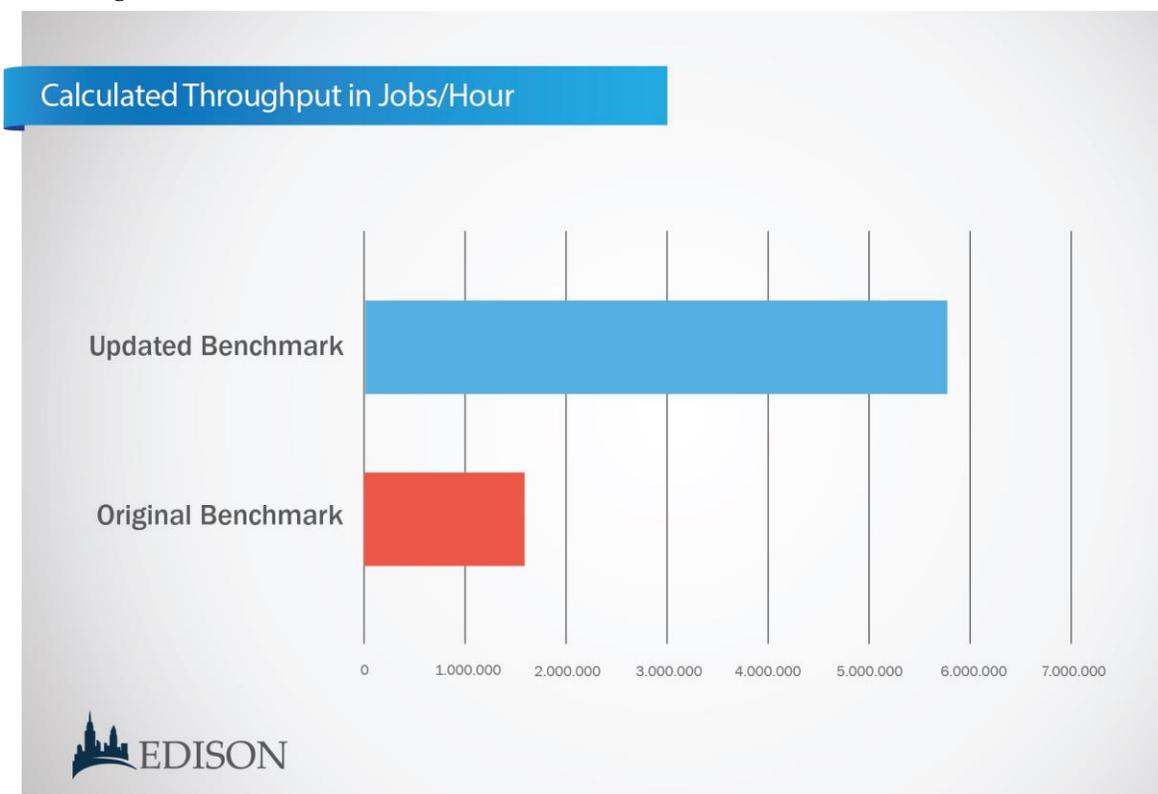


Figure Two: Previous Throughput Jobs per Hour vs. Current Jobs per Hour Capability

As noted with previous similar benchmark testing configurations, while the use of sleep 0 can be viewed as measuring scheduler performance/overheads, it is not representative of a real workload. Real workloads are not zero-length, often require software licenses to

be checked out from a license server, and do access file systems. In fact, while raw performance measurements indicate performance capability, other optimizations, related to code optimization around system and software communications tend to be more important in production systems.⁴

IBM is constantly optimizing the IBM Platform LSF software code, while applying updated and enhanced algorithms to the scheduler. Leveraging additional compiler optimizations and newer available hardware, in addition to IBM optimizations and enhancements, provides obvious and significant scheduler performance.

While having faster processors and other system elements allow IBM Platform LSF to scale, the real benefits of the IBM Platform LSF scheduler becomes apparent when viewing a real-world scenario, as illustrated in Figure Three A, which shows scheduler improvements as the product has evolved and received service packs and patches.



Figure Three A: Code and Processor Optimization – Real Client Workload Throughput in Jobs/Hour; Platform LSF 9.1.3 versus 9.1.3.3

⁴ Depending on Platform LSF configuration and workload pattern, up to 10% additional throughput was achieved with in real workload scenarios utilizing IBM® POWER8®.

Previously, Edison viewed an application of IBM Platform LSF in benchmark comparisons with other workload managers. Configuration was achieved with standard available templates to evaluate performance on an equal, non-optimized basis. As shown in Figure Three B, the scheduling cycle was reduced by 72 percent, while throughput increases 10.5 percent. Clearly, the scheduling cycle is not the bottleneck for this workload pattern.



Figure Three B: Code and Processor Optimization – Real Client Workload Scheduling Cycle in Seconds; 9.1.3 versus 9.1.3.3 with optimizations

It has already been established that IBM Platform LSF dominates its peers in previous benchmarks and the newly available testing round shows that efforts in continual optimization and enhancements applied to IBM Platform LSF have resulted in dramatic improvements in scheduler throughput, while reducing latency (Scheduling Cycle).

While the new data presented in this document does not include another round of exhaustive competitor reviews, several conclusions can be logically inferred. Performance increases, amounting to nearly 3x improvements in job throughput, are observed using the previously evaluated benchmark. In itself, this is an outstanding display of scheduler scalability keeping pace with newer hardware capabilities and enhanced software developments. What really stands out, however, is the agility to scale up performance when optimizations are applied.

Conclusion

In the world of HPC, elapsed time tends to bring the best and worst of software components to the forefront. The case is often made that software can be made to perform faster as additional and/or upgraded hardware becomes available. Too often, while increasing temporarily, performance can hit a brick wall due to framework constraints or lack of software evolution. As is common in many HPC shops (and IT shops generally), a point is reached where more and better hardware does not equate to a healthy or rational return on investment. IBM Platform LSF rises to this challenge.

IBM Platform LSF shows that, on the one hand, hardware and software revisions have and will continue to push the envelope in HPC throughput capability. On the other hand, it is equally important to note that while performance numbers today already showed improvement over previous results, due to constant enhancement and optimization of code to specific architectures, the future performance improvement on both Intel and POWER8 systems will continue to increase.

IBM Platform LSF has room to extend performance capability in the real world through available optimizations when running real customer workloads, and not just synthetic benchmarks.

It is clear that IBM Platform LSF has much to offer for HPC departments evaluating new and existing projects. With consistent growth in capability and performance as the product and related industries continue to evolve, IBM Platform LSF is well positioned in assisting HPTC departments to stay current with ever-changing demands on research and development timelines and requirements.