

Maximize Productivity and Value with IBM Spectrum LSF

Sponsored by IBM

Srini Chari, Ph.D., MBA

June 2016

<mailto:chari@cabotpartners.com>

Executive Summary

The volume, velocity and variety of data are compelling traditional High Performance Computing (HPC) environments to support a new generation of complex data-centric application workloads on highly heterogeneous infrastructures. This is blurring the boundaries between the growing HPC and Big Data Analytics markets; driving up the performance, scale, reliability and accessibility needs of distributed computing environments and clusters.

Today, even as the acquisition cost/performance ratios of large computing clusters continue to improve, the operating costs associated with application licenses, under-utilization, labor, facilities and electricity are escalating. Workload/resource management software solutions such as IBM® Spectrum LSF (formerly known as IBM® Platform™ LSF™) are designed to reduce these operating costs while boosting client productivity and delivering the following value components:

- *Business Value: e.g. better products/services, increased business insight, faster time to market, and new breakthrough capability*
- *Operational Value: e.g. faster time to results, more accurate analyses, more users supported, improved user productivity, better capacity planning*
- *IT Value: e.g. improved system utilization, manageability, administration, and provisioning, scalability, reduced downtime, better access to robust proven technology and expertise.*

IBM Spectrum LSF is a flagship member of the IBM Spectrum Computing Family of software defined computing solutions that share a common foundation. IBM Spectrum LSF delivers cutting-edge, productivity-centric, enterprise-class workload management for HPC environments with intelligent policy-based data management. It is optimized for both high throughput and traditional HPC applications. For over two decades, clients have leveraged IBM Spectrum LSF to reduce costs and maximize value and productivity from distributed computing environments.

IBM Spectrum LSF version 10.1 is the latest release with several innovative capabilities that deliver unparalleled Performance at Scale, Ease of Use and HPC Accessibility. These new innovations will enable clients to further accelerate time to results, maximize productivity and value, and get unprecedented flexibility and insights to better manage a range of complex workloads on heterogeneous infrastructures.

Copyright© 2016. Cabot Partners Group, Inc. All rights reserved. Other companies' product names, trademarks, or service marks are used herein for identification only and belong to their respective owner. All images and supporting data were obtained from IBM or from public sources. The information and product recommendations made by the Cabot Partners Group are based upon public information and sources and may also include personal opinions both of the Cabot Partners Group and others, all of which we believe to be accurate and reliable. However, as market conditions change and not within our control, the information and recommendations are made without warranty of any kind. The Cabot Partners Group, Inc. assumes no responsibility or liability for any damages whatsoever (including incidental, consequential or otherwise), caused by your or your client's use of, or reliance upon, the information and recommendations presented herein, nor for any inadvertent errors which may appear in this document. This paper was developed with IBM funding. Although the paper may utilize publicly available material from various vendors, including IBM, it does not necessarily reflect the positions of such vendors on the issues addressed in this document.

Smart Workload/Resource Management to Maximize HPC Value

The relentless rate and pace of technology-enabled business transformation and innovation are astounding. Several intertwined technology trends in Social, Mobile, Instrumentation and the Internet of Things (IoT) are making data volumes grow exponentially. In 2018, about 4.3 exabytes (10^{18} bytes) of data is expected to be created daily – over 90% will be unstructured.¹

But to consistently extract timely insights from this growing data requires faster, large-scale, cost-effective, highly-productive and reliable High Performance Computing (HPC) systems. Rapid advances in distributed clustering technologies with high-performance processors and accelerators are making large-scale HPC systems affordable to clients; spurring growth in HPC.

The lines between HPC and Big Data Analytics are blurring as High Performance Data Analytics (HPDA) continues to grow 3 to 4 times the rate of traditional HPC.² In many industries, HPDA use cases are rapidly emerging – galvanizing growth by 5.45% per year in the HPC market which is expected to reach \$31 billion by 2019³.



Figure 1: High Value Insights from Integration and Analysis of Structured and Unstructured Data

These HPDA use cases combine Systems of Records (structured data) with Systems of Engagements (unstructured data – images, videos, text, emails, social, sensors, etc.) to produce new *High Value Systems of Insights* (Figure 1).

But maximizing the productivity and value of large-scale, clustered HPC systems, across the entire range of workloads, pose several challenges in how to effectively:

- Address new business needs while adequately maximizing productivity of current application deployments with the available processing capacity
- Manage rising cluster complexity and scale of HPC clusters with hundreds or thousands of compute servers, storage and network interconnect components, and
- Control rising people and operating costs especially as clusters scale, new workloads emerge and additional capital investment and technology upgrades are needed.

¹<https://storageservers.wordpress.com/2016/02/06/how-much-data-is-created-daily/>

² Earl Joseph, et. al., "IDC's Top Ten HPC Market Predictions for 2015, January, 2015

³<http://www.idc.com/getdoc.jsp?containerId=259211>

In 2018, 4.3 exabytes created daily

Lines between HPC and Big Data Analytics blurring

Challenging to maximize productivity and lower costs for HPC

To maximize value, cluster resources must be shared and managed across several users and departments in an organization. Some workload/resource management solutions can:

- Boost resource sharing and improve resource utilization and Quality of Service (QoS) necessary for meeting business priorities and Service-level agreements (SLAs).
- Provide job scheduling and fine grained control over shared resources so that jobs don't compete with each other for limited shared resources.

However, very few workload/resource management solutions can also simultaneously:

- Deliver performance at scale while also maximizing throughput (completed jobs/time unit)
- Enhance ease-of-use with deeper insights on why a workload is stalled or when it will run
- Improve HPC accessibility with web enabled workflows and simpler cloud migration.

IBM Spectrum LSF is a portfolio (part of the [IBM Spectrum Computing](#) family – Figure 2) of smart workload/resource management software that provides these and other unique capabilities to maximize productivity and value from distributed computing HPC environments.

Introducing IBM Spectrum Computing and Spectrum LSF Families

The IBM Spectrum Computing Family is a full range of software defined computing offerings to help organizations eliminate cluster sprawl and reduce complexity. IBM Spectrum LSF is the flagship member of this family and offers a powerful and comprehensive technical computing cluster workload/resource management platform. At its core is a very powerful, time-tested, robust and resilient job scheduler with several value-add components. For over two decades, this industry leading solution has been deployed to support diverse workloads on computationally distributed systems across many industry verticals.

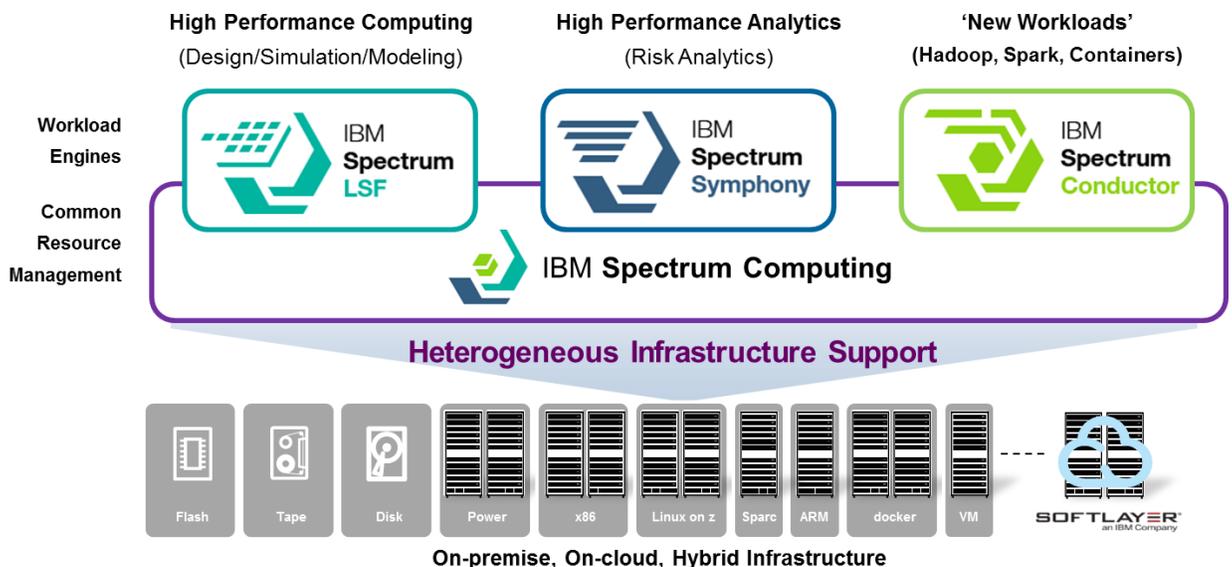


Figure 2: The IBM Spectrum Computing Family Eliminates Cluster Sprawl and Reduces Complexity

*Workload/
resource
management
solutions
improve value
and
productivity*

*But few
perform at
scale with
better ease-
of-use and
access*

*For over two
decades, IBM
Spectrum LSF
is the leading
workload/
resource
management
solution*

IBM Spectrum LSF has add-on components (Figure 3) with additional capabilities: Application-centric user portal, Computation workflow automation, Optimized software licenses, High-throughput resource scheduling, Intelligent data staging, Operational dashboard and Business intelligence & capacity planning. Together, these further enhance the productivity-centric, intelligent policy-based enterprise-class workload management features of IBM Spectrum LSF.

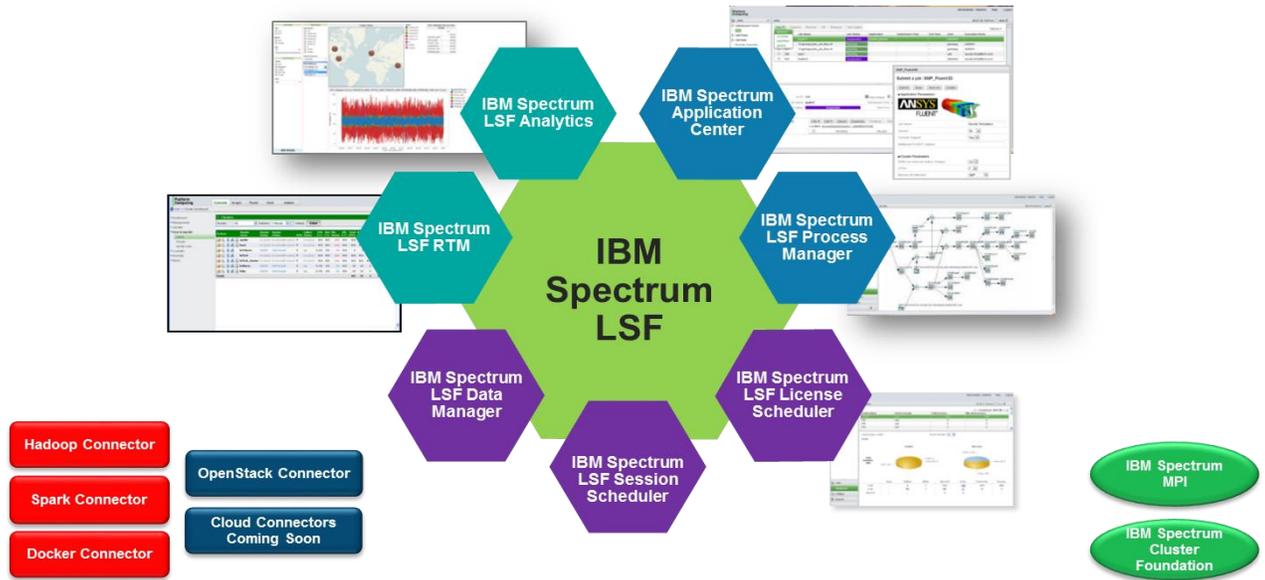


Figure 3: IBM Spectrum LSF Family with Add-on Components

Compared to its competitors, IBM Spectrum LSF has several noteworthy differentiators:

- Resource allocation is always aligned with business priorities; improving cycle times and user productivity
- Better sharing of heterogeneous resources including accelerators and containers; reducing complexity and cost
- Does not “lose” jobs, is very stable, even in large clusters with thousands of nodes; enhancing resilience
- Superior performance and job throughput along with repeatability at scale; vastly improving utilization of cluster resources and time to solution.
- Sophisticated yet provides simple-to-use resource management, operations and reporting, detailed troubleshooting information and other advanced scheduling features; enhancing shared cluster user/administrator productivity and HPC learning curve
- Simple easy-to-use, anywhere/anytime, mobile yet secure access and insights to HPC workloads status/operations; promoting remarkable flexibility and business agility.

In addition, all members of the IBM Spectrum Computing family share a common foundation that allow physical resources to be shared / reallocated between the workload heads. This means that a Spectrum LSF cluster can request additional resources from other members of the Spectrum Computing family or from OpenStack as required.

IBM Spectrum LSF version 10.1 is the latest release with several innovative capabilities that further enhance value and productivity for clients.

IBM Spectrum LSF Family with value-adds for enterprise workload management

Better value than competitors

A common foundation to promote greater sharing

New Innovations in the Latest Version of the IBM Spectrum LSF

IBM Spectrum LSF version 10.1 continues to excel in improving productivity and value (Figure 4) by better delivering: *Performance at Scale*, *Ease of Use* and *HPC Accessibility*.

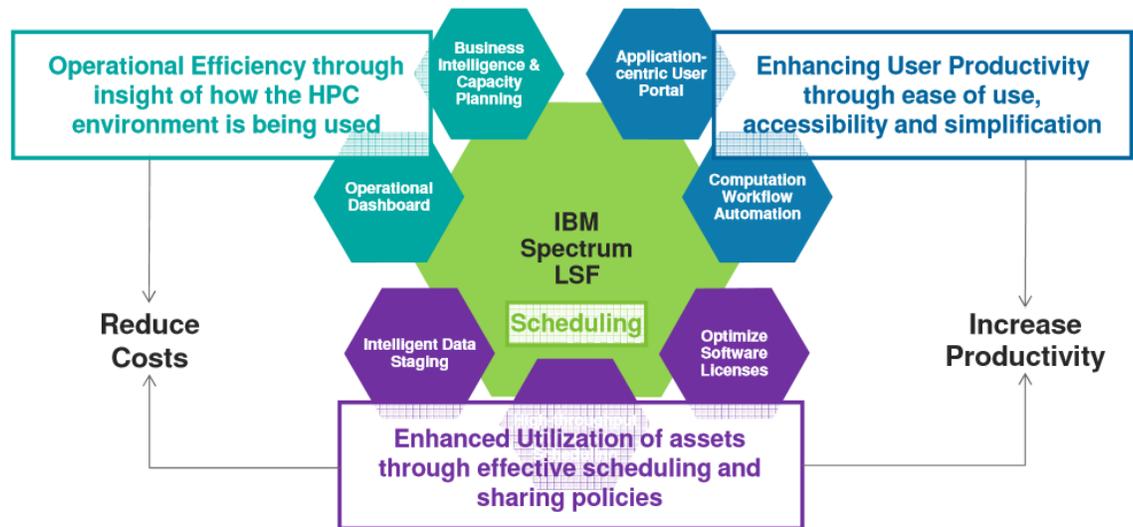


Figure 4: The Core Value Provided by the IBM Spectrum LSF Family

Performance at Scale: A very hard challenge in any compute cluster environment is to simultaneously maximize throughput while maintaining optimal application performance. High throughput requires the elimination of load imbalance among the compute nodes. On the other hand, optimal application performance demands reduction in communication overhead and careful mapping of the application workload to the best available compute resources. Achieving this in tandem is hard. IBM Spectrum LSF addresses this challenge effectively.

Specifically, the current Spectrum LSF version reduces “lost” scheduling time by decoupling job dispatch from the scheduling cycle. The scheduler can re-use an allocation and dispatch the next job based on the “current” priority list, rather than waiting for the next cycle. This scheduler-based chunking is very effective for short high throughput workloads.

Further, disk I/O costs and performance bottlenecks are eliminated by compressing and keeping the large number of, almost identical, small active job files in memory; accelerating job dispatch times. Inter-host and inter-process communication costs are reduced with a new multi-threaded communication library, lower TCP (re)connection overheads, no blocking communications, shared memory optimizations on the master host and tuned bulk operations.

According to recent IBM internal benchmarks and other studies⁴, IBM Spectrum LSF delivers unparalleled performance and scalability: 5.8 times increase in throughput, more than 95% reduction in the scheduling interval, drastic reduction in disk I/O costs and better inter-process communication for large clusters compared to prior versions. These performance advantages result in faster scheduling, higher resource utilization and improved productivity across many industries.

⁴http://www.ibm.com/common/ssi/cgi-bin/ssialias?subtype=WH&infotype=SA&apname=STGE_DC_ZO_USEN&htmlfid=DCW03058USEN&attachment=DCW03058USEN.PDF

Performance at Scale, Ease of Use and HPC Accessibility.

Optimized for highly parallel and high throughput jobs simultaneously

5.8 times increase in throughput, more than 95% reduction in the scheduling interval

For instance, in high throughput environments (the workloads are primarily sequential) such as Electronic Design Automation (EDA) and Life Sciences, IBM Spectrum LSF has driven slot utilization levels greater than 97.5% with over 100,000 concurrent jobs on dual-socket IBM POWER 8 or x86 clusters of over 100,000 cores with pending workloads greater than 20X the running workload. On the other hand, in highly parallel HPC workloads (Computer Aided Engineering (CAE) or Research), IBM Spectrum LSF has driven slot utilization levels greater than 99% with concurrent jobs on 10% of the 270,000 cores on POWER8 and x86 systems with pending workload of over 20 times the running workload. These compelling performance results are a clear testament to the ability of IBM Spectrum LSF to perform at scale.

Ease of Use: Another key challenge in workload management is for users to get granular and more precise information on their jobs' status in the queue: Why are my jobs not running? When will they execute? Getting such detail on pending jobs is the top requirement from the user community. IBM Spectrum LSF provides these details to users and administrators.

Users can easily identify which jobs are actually eligible for scheduling: Is it blocked by a future start time, dependency, error condition or something else? They can also identify the most important reason why a job is pending. In addition, users can also drill deeper and get other more precise granular reasons for why *their* specific job is pending. Also, since some resources (software licenses, hardware emulators, etc.) are more important (or scarce) than others, organizations are now able to prioritize and customize pending messages for specific resources. Lastly a pending messages summary gives users a holistic personalized view.

IBM Spectrum LSF Application Center (Figure 5) now allows a user to graphically view all the dependencies for a job, and all the jobs that depend on the current job.

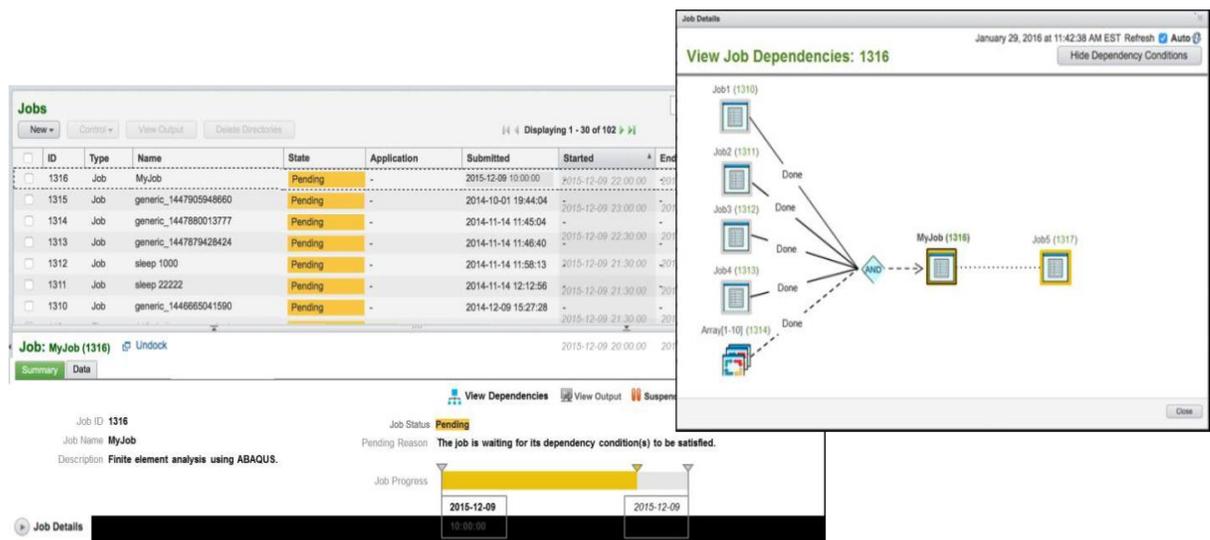


Figure 5: Viewing Job Dependencies with the IBM Spectrum LSF Application Center

Utilization levels of up to 99% or greater

Improved viewing of job dependencies and workflows

Better insights on why jobs don't run, or when they will run

IBM Spectrum LSF also raises an alarm if a job has been pending for too long. This works in conjunction with the existing alarming functionality in IBM Spectrum LSF RTM. The user can define two pending time limits: Absolute limit – total pending time from submission, and Eligible limit – eligible pending time limit. In addition, job start time prediction is enhanced using the LSF Simulator technology which periodically takes a cluster snapshot and better estimates the state of the cluster into the future.

Personalized alerts and documentation to improve user experience

Finally, even product documentation is designed with the goal to assist customers help themselves find relevant information quickly, when, where, and how they want it. Product information is mapped to the users' anticipated tasks at hand; making it easier for them to search, find and use the most relevant information to achieve their immediate business goals. The IBM Knowledge Center provides a central repository for product documentation with corrections/additions available immediately. This repository is fully searchable and is customizable for on-premise use based on personal preferences. This local copy can be automatically synced with the most current IBM Knowledge Center version.

HPC Accessibility: With an increasingly global and mobile workforce, users need to be able to integrate the tasks they do using their desktops at work with their mobile phones/tablets anytime and from anywhere. This increases workplace flexibility and enhances personal and organizational productivity, collaboration and innovation.

IBM Spectrum LSF provides web enabled workflows, simplified migration of HPC workloads to cloud environments, and new user mobility options enabling access from anywhere.

- In addition to traditional Command Line Interface (CLI) desktop access to LSF, it is now possible to use CLI with secure remote data upload/download.
- A new integrated desktop client with drag and drop submission and automatic data download features improves productivity through tight integration with the user's desktop. It also increases mobility by leveraging http communication versus proprietary protocols and provides status information with Taskbar access to the job dashboard and one click access to the remote job console. This client automates upload of input files, and the download of results. The client also includes a RESTful (web services) version of the traditional IBM Spectrum LSF CLI, simplifying access to cloud environments.
- Lightweight access via mobile devices for job monitoring and notifications.
- The Flow Editor component of IBM Spectrum LSF Process Manager is now browser based, allowing users to create, modify and run flows from with their browser.

Integrated desktop, browser and mobility/cloud options to improve accessibility

The IBM Spectrum LSF 10.1 Family delivers cutting-edge productivity-centric, enterprise-class workload management for HPC environments with intelligent policy-based data management. It is optimized for both high throughput and traditional HPC applications. For over two decades, clients have used IBM Spectrum LSF to reduce costs and maximize value and productivity from distributed computing environments.

How Clients Benefit from the IBM Spectrum LSF Family

Here are some recent real life examples of how IBM Spectrum LSF handles workload management intelligently, boosts overall performance and utilization of HPC environments:

IBM Electronic Design Automation

(Accelerating time-to-market and trimming millions of dollars from chip design)

<p>Description / Challenges</p>	<ul style="list-style-type: none"> • Tremendous business pressure to deliver advanced chip capabilities to market. • Curb developer practices such as overbooking intended to beat scheduled job start time delays that caused productivity loss and resource under-utilization. • Effectively scale and speed up chip validation mechanism with no additional investment. • Increase shared HPC infrastructure utilization for design, testing and chip refinement beyond 150 million grid jobs annually. • IT resource constraints on the volume and intensity of pre-fabrication testing.
<p>Solution / Results</p>	<ul style="list-style-type: none"> • IBM Spectrum LSF workload management and IBM Spectrum LSF RTM operational dashboard together helped in optimizing job scheduling and real time resource usage monitoring. • Triage of pain points became simpler and faster with macro-level view into lifetime of tools, flows, memory, runtime process and thread counts. • Boost in overall performance and utilization of shared HPC infrastructure by cutting time to complete compute jobs. • Enabled running of additional tests within the same deadlines with the same amount of resources. • Seamless migration to the new solution while in full production, on a site-by-site basis, with no business disruption. • Enabled the launch of a new virtual desktop tool to effectively share a larger set of resources via a user-friendly environment that could be accessed by all IBM EDA teams across geographies, boosting collaboration and productivity
<p>Benefits</p>	<ul style="list-style-type: none"> • IBM EDA saved millions of dollars by finding and fixing chip fabrication bugs prior to the fabrication stage. • 10% higher utilization enabled more thorough testing without extra IT investment. • 10X faster scheduling helped them save staff time and cut time-to-market.

Accelerating time-to-market and trimming millions of dollars from chip design

Cypress Semiconductor

(Lowering costs in bringing advanced technology to market faster)

Lowering costs in bringing advanced chip technology to market faster

Description / Challenges	<ul style="list-style-type: none"> • Complicated storage infrastructure at Cypress created challenges in simulating and verifying schematics for new chip designs that were I/O intensive. • Risks of data loss associated with the manner in which simulation data was stored on local disks during verification. • Isolated file servers caused day-to-day management issues. If a user ran out of space in the middle of a 48-hour simulation, they needed to start over, raising costs and extending total cycle time (TCT). • Storage performance, absence of centralized and scalable solutions and capacity problems impacted Cypress' capabilities to deliver to market faster. • Had to reduce chip design costs without impacting capability to bring new products to market ahead of competitors' offerings. • Need to constantly cut total cycle time (TCT) and increase both potential revenues and profit margins.
Solution / Results	<ul style="list-style-type: none"> • Cypress was able to remove storage bottlenecks and set up a robust job-scheduling environment using IBM Spectrum Scale to support global chip design and simulation HPC clusters managed by IBM Spectrum LSF and its add-on module IBM Spectrum LSF RTM. • Significant performance improvements and cost savings, along with an unexpected extra benefit of added resiliency with IBM Spectrum LSF and IBM Spectrum Scale • More predictable performance helped size capacity requirements more accurately. • An order of magnitude performance improvement on the existing hardware versus the previous distributed file system. • Complete elimination of expensive unplanned downtime during EDA simulation jobs boosted productivity and eliminated downtime costs. • With faster and more reliable storage along with intelligent and robust job scheduling offered by the IBM solution, Cypress could achieve better ROI on their expensive EDA software license investments. • Additional benefit of solution: Capability to perform planned maintenance, taking nodes off line without impacting IBM Spectrum LSF jobs.
Benefits	<ul style="list-style-type: none"> • 10X better performance on their current hardware helped in cutting time to market. • Reduced TCT provides major annual savings. • Better resiliency keeps EDA simulation jobs running and cuts administration time and cost.

Conclusions and Recommendations

For over two decades, clients in Financial Services, Life Sciences, Computer Aided Engineering, Electronic Design Automation, Oil and Gas, and other industries have deployed the IBM Spectrum LSF Family of workload/resource management solutions to reduce costs and improve productivity and value from their HPC applications and infrastructure.

Today, with the growth of the volume, velocity and variety of data and the blurring of the lines between HPC and Big Data Analytics, clients are increasingly deploying High Performance Data Analytics (HPDA) solutions. These solutions combine structured and unstructured data; creating an unrelenting appetite and need to solve larger and more time-critical complex problems on heterogeneous cloud-like infrastructures.

The IBM Spectrum Computing Family is a full range of software defined computing offerings with a common foundation to help clients eliminate cluster sprawl and reduce complexity. A flagship member of this family, IBM Spectrum LSF version 10.1 delivers better Performance at Scale, Ease of Use and HPC Accessibility compared to prior versions or competitors.

Organizations should actively consider deploying IBM Spectrum LSF version 10.1 to:

- **Accelerate time to results by:**
 - Managing different workload types including from traditional HPC to high throughput computing and big data while leveraging accelerators and containers
 - Allocating the right resource to the right users at the right time
 - Ensure priority/urgent work goes first
- **Enhance Ease of Use by:**
 - Better understanding why a workload will not run, or when it will run
 - Personalizing the environment (including product documentation) to business tasks/priorities
- **Improve HPC accessibility and learning curve by:**
 - Better integrating the desktop environment
 - Supporting web-enabled workflows
 - Simplifying migration of HPC workloads to cloud environments
 - Providing new user mobility options enabling access from anywhere.

With the latest IBM Spectrum LSF Family, clients will be able to maximize productivity and value from their HPC environments.

Cabot Partners is a collaborative consultancy and an independent IT analyst firm. We specialize in advising technology companies and their clients on how to build and grow a customer base, how to achieve desired revenue and profitability results, and how to make effective use of emerging technologies including HPC, Cloud Computing, and Analytics. To find out more, please go to www.cabotpartners.com.