

# The hidden costs of open source

*Rethinking the economics of HPC infrastructure software*



## Introduction

Clusters based on open-source software and the Linux® operating system have come to dominate high performance computing (HPC). This is due in part to their superior performance, cost-effectiveness and flexibility. The same factors that make open-source software the choice of HPC professionals have also made it less accessible to smaller centers. The complexity and associated cost of deploying and managing open-source clusters threatens to erode the very cost benefits that have made them compelling in the first place.

As customers choose between open-source and commercial alternatives, there are many different costs related to administration and productivity that should be considered. These are explored in this paper in order to give a true cost perspective. We also examine how a commercial management product, such as IBM® Platform™ HPC, enables HPC customers to side-step many overhead cost and support issues that often plague open-source environments and enable them to deploy powerful, easy to use clusters.

## HPC: Why this debate is different?

The pros and cons of open-source versus commercial software have been debated in corporate IT for years. While sometimes viewed as a decision between Windows® versus Linux®, the issues are more complex than just the choice of an operating system.

In HPC circles the Linux versus Windows debate has been largely settled, with Linux enjoying 75 percent market share compared to less than 6 percent share for Microsoft Windows<sup>1</sup>.

Linux has an even more commanding position among TOP500® supercomputing sites<sup>2</sup>. For this reason, in HPC the debate more often centers on whether to deploy a pure open-source management environment or to use a supported commercial product built for open-source Linux environments.

## Accounting for real costs

Despite the many benefits of open-source software, it is not without its pitfalls. This is particularly true for smaller organizations that may lack staff with requisite expertise to operate and maintain open-source software and middleware environments.

## When free is not really free

The costs associated with deploying and maintaining open source software can manifest themselves in several ways. These include:

- Increased time spent on system administration.
- Excessive time spent troubleshooting problems due to the lack of a formal support channel.
- Costs related to handling functions such as quality assurance and regression testing in-house.
- Reduced productivity owing to cluster downtime, sub-optimal cluster performance and low utilization.
- Additional expenses in education if new software development skills are required to maintain an open source environment.
- Creeping operational costs as organizations find themselves not just in the research business, but in the software maintenance business.
- Unexpected costs that can arise when needed functionality cannot be fulfilled by available open-source software.

### Other considerations

Aside from the factors above, other items contribute to cost and risk when using open source software, including:

- Lack of technical roadmaps for open-source software, making planning a challenge.
- Important software maintenance such as performing updates and applying security patches may be avoided for fear of breaking fragile software environments.
- Open-source software choices may preclude the selection of new hardware or layered compatible software, constraining options and leading to costly re-engineering efforts downstream.
- The quality of open source community support varies and, if available at all, is often narrowly focused on specific software components. This leaves administrators to shoulder the burden of software integration issues themselves.

### The problem of “sparseness”

Comparisons between open-source and commercial software often assume that the alternatives have equal technical merit. This may hold true for components in widespread use like operating environments, databases and scripting languages where the quality of open source software is well recognized. It is often not true, however, for more specialized software that is less widely deployed and less thoroughly exercised.

In more specialized areas, open source products may not exist at all. If they do exist, they may provide less functionality and more complexity, and require more effort to adapt and install than their commercial counterparts.

Some specific areas in HPC where open-source solutions can be challenging to source and successfully integrate include:

- Web-management consoles and portals
- Drivers optimized for high-speed interconnects
- Reporting and analysis tools
- Tools to manage engineering workflows
- Tools to facilitate application integrations
- Workload driven provisioning solutions

These and other capabilities often emerge as critical requirements. However, because open-source solutions are unavailable or lack needed features, organizations can face significant costs to develop, purchase or integrate the needed functionality.

### Real TCO: More like calculus

The decision whether to deploy open-source or commercial software for HPC is sometimes painted as a binary choice. In practice, however, organizations have a range of different options.

Figure 1 illustrates a range of alternatives between purely open and commercial solutions. The shape of the total cost of ownership (TCO) curve will vary depending on the environment, so this is not meant to be definitive, but rather to illustrate that different organizations will weigh costs and benefits differently.

For most organizations, being at one extreme or the other is likely to be expensive and limits their options and downstream flexibility.

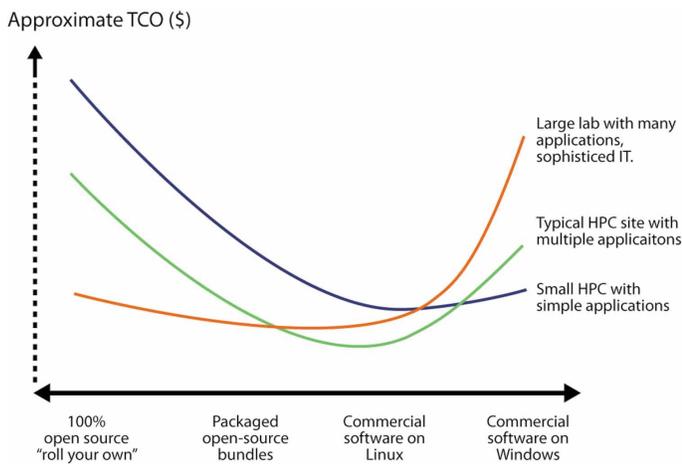


Figure 1: A range of deployment choices for HPC

Organizations that deploy and support their own pure open-source environments (operating at the left end of the diagram in Figure 1) take on a significant amount of in-house integration, development and support cost. For organizations that have deep technical expertise and are already developing applications this may be reasonable, but for smaller organizations this can be prohibitively expensive.

At the other extreme, with a commercial solution based on a proprietary operating environment, like Windows HPC Server, costs arise in different places. Organizations will pay higher costs for software licensing, maintenance and professional services, but will enjoy a better supported and integrated system.

A risk that can impact TCO in Windows environments is lack of access to the vast number of open-source HPC tools that already exist on Linux or UNIX®. Organizations operating a fully proprietary environment can incur additional costs as they are forced to purchase or develop needed solutions themselves.

There can also be surprises related to personnel costs. For example, it is simple to assume that the same skills suitable for administering a Microsoft Office or Microsoft Exchange environment are readily portable to the more complex task of administering Windows HPC clusters. Regardless of the environment deployed, deeper expertise impacting personnel costs will be required.

### A pragmatic approach

Many customers opt for a more pragmatic approach—blending open-source and commercial software in a fashion that minimizes total costs while supporting the full range of applications that they expect to run. A commercially supported HPC product deployed on their choice of Linux operating system provides users with the “best of both worlds.”

Using this combination of products provides freedom of choice, while also reducing operations and support costs. This is because cluster users are able to take full advantage of the rich HPC tools available on Linux. However, they do not have to deal with the challenge of integrating open source components. In addition they have access to a technical support organization to help them resolve problems quickly.

IBM® Platform Computing™ provides cluster management products that specifically occupy this TCO “sweet spot,” as illustrated in Figure 2.

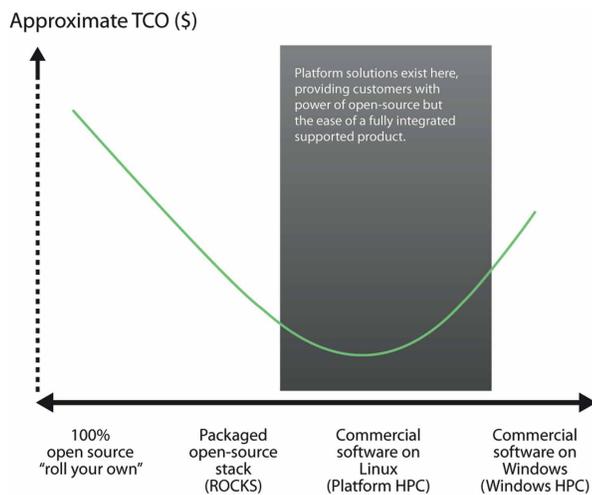


Figure 2: Platform HPC is a fully supported cluster management product for Linux environments.

## Introducing Platform HPC

Platform HPC makes it easy to harness the power, flexibility, and scalability of an HPC cluster. Whether users have a small or large cluster, Platform HPC quickly gets their cluster up and running. It simplifies the application integration process so that they can focus on their work, instead of managing their cluster.

Other HPC cluster solutions combine multiple tools and interfaces, which are not integrated, certified, or tested together. In comparison, Platform HPC is the industry's only complete management product. It includes a mature, robust set of management capabilities that are accessible through a unified portal interface. Backed by the industry's best customer support, Platform HPC also incorporates nearly two decades of product and technology leadership.

Platform HPC includes the following key capabilities:

- Easy-to-use cluster management
- Integrated application support
- Robust workload and system monitoring and reporting
- Dynamic operating system multi-boot
- Commercial MPI library
- Web-based interface for access anywhere
- GPU support including NVIDIA® CUDA kit

Because Platform HPC can be managed entirely through a browser, Linux clusters can be deployed and managed by non-specialists. This removes a longstanding barrier to Linux adoption, particularly in smaller sites.

The web interface not only simplifies management, but makes the cluster easier to use. Easy-to-use application interfaces reduce training and support requirements, reduce job submission errors, and boost productivity.

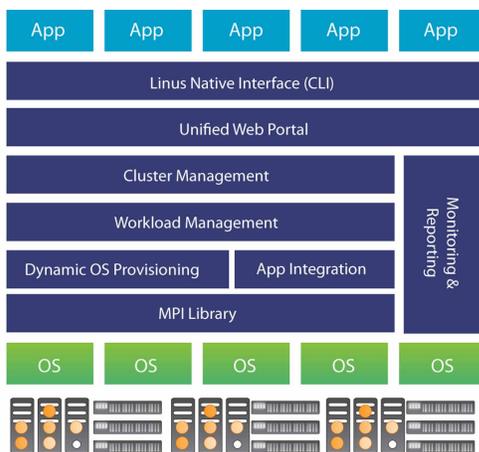


Figure 3: Platform HPC key features

Furthermore, Platform HPC enables users to easily deploy mixed Linux and Windows clusters managed through a single web interface. It can flex the cluster resources automatically, booting different operating systems on demand in response to changing workloads. This avoids costly duplication of resources, and provides an apparently larger resource pool to cluster users resulting in improved productivity.

## Sources of savings

Deploying Platform HPC can help organizations save or avoid costs in several ways. This includes:

- Reduced administration
- Reduced training and support costs
- Improved user productivity
- Avoidance of cluster down-time
- Improved cluster utilization
- Avoidance of unanticipated costs

The following sections examine some of these sources of savings in more detail.

**Personnel savings:** Even for experienced Linux administrators, building a fully functional cluster from open-source components can take weeks. Platform HPC can dramatically reduce the time to install and configure a cluster because it is a pre-certified and pre-tested product and already includes all of the software features that HPC sites need. Because of its comprehensive, easy to use web interface, Platform HPC can easily be installed and managed by non-specialists.

**Improved productivity:** Platform HPC helps make cluster users more productive, and ensures that cluster downtime is kept to a minimum. Features such as the cluster file manager<sup>3</sup> and repository snapshots take the risk out of tasks like software upgrades and patching. They enable administrators to easily roll-back to a known good configuration if anything goes wrong with a software update or patch installation. This helps reduce downtime, makes software changes simple and provides peace of mind to users concerned about breaking functionality by performing a software update.

**Hidden cost avoidance:** Some organizations do not include system and network management tools as part of the cost of their HPC deployment. Acquiring and integrating these tools can result in real costs. Many other cluster management solutions require additional expenditures for many of the capabilities that are included as standard features in Platform HPC. Examples of these extra-cost components include:

- System monitoring and alerting tools
- Workload management systems
- User-centric application-aware web portals
- Commercial grade MPIs
- Web-based interfaces for data handling
- Dynamic node provisioning and management tools

**Reduced error rates:** A source of cost that is sometimes overlooked is related to user errors on job submission. A 16-way MPI job that runs for two hours that needs to be run again because the user submitted the job incorrectly “costs” 32 hours of CPU time. During this period, other jobs may need to remain in the workload management system, waiting for available resources. By providing interfaces tailored to applications, and by allowing users to monitor and manage their jobs in real-time, the opportunity for job submission errors is reduced. If there is an error, users can identify the error quickly and take corrective action without bothering an administrator.

**Improved cluster utilization:** There is a significant difference in the cost and productivity of a cluster running at 80 percent utilization versus a cluster running at 95 percent utilization. Platform HPC is based on IBM® Platform™ LSF®, which is

widely regarded as the best in class commercial workload management system. By taking advantage of superior scheduling capabilities, customers can better align cluster resources to business needs and can achieve better levels of cluster utilization, allowing them to do more with less.

**Superior price and performance:** Because the supplied tools and libraries in Platform HPC are tuned and optimized for selected vendor hardware configurations, customers can be assured that they are getting the highest level of performance possible from their clusters. With a do-it-yourself approach to building clusters from open-source components customers risk deploying un-optimized libraries and toolsets. The components may work, but at reduced performance levels, undermining the whole purpose of deploying a high-performance cluster.

**Reduced risk:** By relying on pre-tested, certified configurations fully backed by IBM Platform Computing and hardware vendor partners, customers are assured that any issues can be addressed quickly and efficiently without the need for on-site consultants or additional support expertise that result in unbudgeted costs.

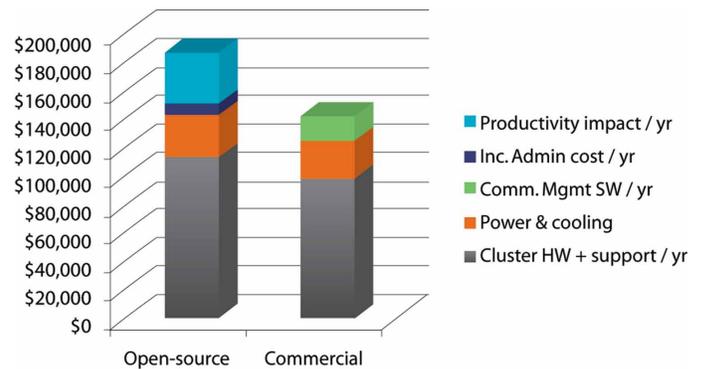
**Infrastructure cost avoidance:** Many sites run both Linux and Windows applications. Rather than provisioning separate clusters, total costs can be reduced by using Platform HPC to manage hybrid clusters comprised of both Windows and Linux nodes. Platform HPC provides workload-driven OS provisioning as a standard feature, allowing customers to use a single management tool and avoid unnecessary infrastructure spending.

**“Future-proofing”:** Over time, sustaining costs tend to dominate IT related spending. A cluster might run for a few years, but when it comes time to add nodes to that cluster, upgrades can be a challenge since new hardware platforms may require new operating systems and drivers that did not exist on the head-node when the cluster was originally installed.

Platform HPC enables customers to boot new hardware models that require updated device drivers from an existing cluster head-node without re-installing the management software from scratch. This means that clusters can be grown incrementally, avoiding costly downstream “rip and replace” scenarios and extending the useful life of the cluster infrastructure.

## Comparing the costs

TCO estimates vary based on many factors including the nature of the installation, in-house capability, types of applications and cost of down-time. For example, consider a typical 32 node cluster serving ten knowledge workers who each spend roughly 30 percent of their time working with HPC applications. For the purpose of this analysis, it is assumed that the environment is supported by a single, fully competent part-time cluster and application administrator. It is also assumed that the customer will deploy the same hardware and interconnect environment and the same open-source or commercially supported operating system<sup>4</sup>, regardless of whether the customer chooses open source software or Platform HPC.



*Figure 4:* Annual cost of ownership of a 32 node cluster with open-source versus Platform HPC software

Figure 4 compares the annual TCO for this hypothetical environment when running an open-source HPC software stack with the same cluster running a commercial supported, fully integrated HPC management product. As the figure illustrates, organizations not only need to make their own determinations of TCO based on how they account for details such as the cost of labor and depreciation, they also need to include productivity costs. Because a non-productive or idle engineering team is vastly more expensive than a part time administrator, these costs can rapidly dwarf administration

costs. Therefore, the real opportunities for savings lie in areas such as application integrations, which drive productivity, and timely risk-free upgrades that reduce down-time, and improved resource utilization.

Appendix A details the assumptions and calculations behind the TCO estimates in Figure 4. If the example depicted an environment with more active users or less in-house cluster administration expertise, the argument for commercial software would become even more compelling.

The model assumes that there was an opportunity to reduce the cluster node count by 10 percent through more effective infrastructure sharing using adaptive clustering or more sophisticated resource sharing. Even if there were no infrastructure cost savings, reductions in administrative costs and positive impacts on user productivity make the case for Platform HPC compelling. In addition, the useful life of resources was not accounted for. This might be extended in a Platform HPC environment by simplifying the management of heterogeneous clusters, further improving cost savings.

## Summary

While Linux clusters dominate HPC, there are many issues related to cost and complexity that can make open-source solutions challenging. In addition, determining real costs can be complex because every environment is different, and organizations will assess costs using their own methodologies and based on their own requirements and capabilities.

For many organizations, the most practical and cost-effective approach is to deploy a commercially supported cluster management product on a Linux platform. Platform HPC represents an attractive solution because it is a complete product that preserves the many benefits of open-source software while avoiding the headaches and management costs of open-source software.

Because it is a complete, integrated product, Platform HPC allows analysts, engineers and scientists to focus on their work rather than on less productive pursuits such as deploying, managing and supporting open-source clusters.

## Appendix A: Comparing Open Source versus HPC Cluster Manager TCO

The table below details the assumptions behind the cost comparison in Figure 4. These assumptions can be adjusted to reflect the unique characteristics of any real HPC environment.

Cost Model Assumptions		
Details about cluster infrastructure		
a)	Number of cluster nodes	32
b)	Per node hardware capital costs	\$8,000
c)	Annual cost of hardware maintenance / extended warranty per node	\$500
d)	Electrical power rate (\$/kWh)	\$0.13
e)	Per-node power consumption (watts)	550
f)	Hardware ammortization / depreciation period (years)	3
g)	Annual cost of cluster (depreciation + maintenance) -- $(a*b) / f + (a*c)$	\$101,333
h)	Annual power and cooling related costs for cluster -- $d*(e/1000)*a*24*365$	\$29,229
i)	Annual license and support cost of Platform HPC management software per node	\$550
Personnel cost assumptions		
j)	Annual burdened cost per system administrator	\$90,000
k)	Number of knowledge workers in environment	10
l)	Annual burdened cost per knowledge worker	\$130,000
m)	Percentage of time knowledge workers use the cluster	20%
Assumptions related to commercial management software benefits		
n)	Estimated reduction in cluster deployment time (weeks - over total lifespan)	3
o)	Estimated reduction in unplanned downtime per year (hours)	48
p)	Saved time per knowledge worker through self-serve access (hrs/yr)	14
q)	Saved time per knowledge worker through intuitive interfaces (hrs/yr)	16
r)	Administrator time saved / year deploying new applications (days/yr)	3
s)	Avoidable job submission errors owing to simpler interfaces (%)	2
t)	Estimated annual reduction in administrator support burden (days)	20
u)	Reduce cluster node count owing to better resource sharing (Adaptive Clustering + LSF)	10%
v)	Unplanned incremental costs, software, consulting (annual)	\$5,000

Given the above assumptions, the cost of maintaining an open-source cluster deployment and management environment can be compared with that of Platform HPC, a commercially supported cluster management product.

Comparative Cost Calculations	Open-source	Commercial
Infrastructure hardware, system & mgmt SW, support costs		
Annual HW depreciation cost	\$101,333	\$91,200
Annual HW support cost	\$8,000	\$7,200
Power / cooling	\$29,229	\$26,306
Cost of management software	\$0	\$17,600.00
Unforeseen incremental costs (per year)	\$5,000	\$0
	\$143,563	\$142,306
Incremental administration cost of open source environment		
Incremental costs related to cluster administrator -- $(t/220)*j$	\$6,923	(baseline)
Incremental costs related to new application deployments -- $(r/200)*j$	\$1,350	(baseline)
	\$6,923	\$0
Incremental cost of reduced knowledge worker productivity		
Productivity cost related to delayed or problematic deployment $(k*i*m)*n/52*1/f$	\$25,000	(baseline)
Productivity cost related to avoidable downtime $(k*i*m)*(o/(220*24))$	\$2,364	(baseline)
Productivity cost of reduced self-serve access and delays $(k*i*m)*(p/8)/220$	\$2,068	(baseline)
Productivity cost of non-intuitive interfaces and associating learning curve $k*i*m*((q/8)/220)$	\$2,364	(baseline)
Lost productivity owing erroneous submissions $(k/l/m)*s$	\$5,200	(baseline)
Factored productivity cost (annual)	\$36,995	\$0
Total annual TCO Estimate	\$187,481	\$142,306

#### Additional information:

- Reductions in infrastructure support and power costs are driven by assumption (u).
- Unforeseen incremental costs (v) stem from the cost of addressing functional requirements that open-source software may not readily support. These costs may take the form of software purchases, consulting time, or time and effort to find, integrate, deploy and support additional layered solutions
- This analysis assumes a 220 day working year—for example 22 days of effort (a little over four working weeks) would be considered to be 10 percent of an individual's time when accounting for cost
- There are costs related to administration and user-productivity regardless of whether open-source or commercial software is chosen. Therefore the model only accounts for incremental costs that may manifest themselves in open source deployments. Similar costs associated with Platform HPC to be baseline costs.
- This analysis is simplified for this whitepaper, but is based on a TCO calculator developed by IBM Platform Computing.

#### For more information

To learn more about IBM Platform Computing, please contact your IBM marketing representative or IBM Business Partner, or visit the following website: [ibm.com/platformcomputing](http://ibm.com/platformcomputing)



© Copyright IBM Corporation 2012

IBM Corporation  
Systems and Technology Group  
Route 100  
Somers, NY 10589

Produced in the United States of America  
May 2012

IBM, the IBM logo, [ibm.com](http://ibm.com), Platform Computing, Platform Symphony, and Platform LSF are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at [ibm.com/legal/copytrade.shtml](http://ibm.com/legal/copytrade.shtml)

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

The content in this document (including currency OR pricing references which exclude applicable taxes) is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary. THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

<sup>1</sup> [insidehpc.com](http://insidehpc.com), April 12 2010, Microsoft and the Linux-friendly Windows HPC server

<sup>2</sup> <http://top500.org/stats/list/36/osfam>

<sup>3</sup> The cluster file manager allows node changes to be made "on the fly" without the need for a re-provisioning operation that results in downtime on other cluster managers.

<sup>4</sup> Because users can deploy a commercially supported management solution on an open-source operating system or deploy an open-source cluster management solution on a commercially supported Linux, for purposes of this analysis OS licensing and support costs are not included as differentiating factor.



Please Recycle