A COST-BENEFIT LOOK AT OPEN-SOURCE VS. COMMERCIAL HPC WORKLOAD MANAGERS

Introduction

High-performance computing (HPC) fuels scientific discovery and innovation across multiple industries. The combination of large datasets, advanced simulation techniques, and machine learning helps organizations generate insights that would not be possible without modern HPC infrastructure.

Given HPC's outsized role in driving business results, selecting the right management software is critical. This is especially true in commercial organizations where time is money. In this paper, we discuss the pros and cons of open-source software in HPC and make the case for commercial workload management. While open-source workload managers are fine in some situations, they can present disadvantages in production environments:

- **In HPC, performance is key** – If open-source software even slightly degrades the performance, productivity, or reliability of HPC infrastructure, it will be viewed as a poor decision. Even small impacts on HPC productivity can have large impacts on critical measures such as revenue, profitability, and time to market.

- **The hidden costs of free** – Open-source software often trades up-front costs for added risks and ongoing expenses. Organizations can find themselves facing higher integration and support costs maintaining custom solutions, becoming overly reliant on expensive integrators or consultants, and having less strategic flexibility.

- **Costs of downtime** – Organizations should carefully consider their availability requirements and think through the impact of critical bugs, outages, and security vulnerabilities and their potential impacts on HPC operations.

Not only can Altair commercial workload managers enable better performance, productivity, and reliability, they can reduce costs while substantially boosting return on investment (ROI). We present simple ROI calculations in this paper that illustrate how small improvements in efficiency can yield dramatic financial returns. We also present a total cost of ownership (TCO) model that HPC operators can use to estimate some of the hidden costs that often accompany open-source deployments. Armed with this information, HPC users can make more informed decisions based on their unique environment and business requirements.

The Critical Role of HPC in Industry

In business, HPC is a force multiplier. Investments in HPC can have a dramatic impact on the bottom line. According to Hyperion Research, organizations realize an average revenue increase of $463 per dollar invested in HPC. Similarly, each dollar invested generates approximately $44 of profits. To put this in perspective, a $100,000 HPC investment by a private corporation will typically yield $4.4 million in incremental earnings.[1]

Figure 1 shows average ROIs by industry based on Hyperion's analysis of 150 use cases. While amounts vary, most industries realize a significant return on their HPC investments.

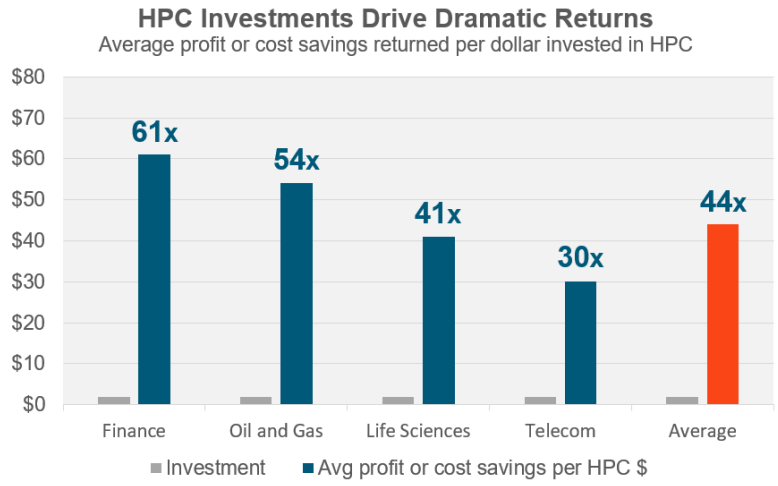

High performance computing generates on average **$44** in profit for every dollar invested in HPC systems.

Source: Hyperion Research, 2020

Figure 1 - ROI associated with HPC investments

Hyperion's analysis illustrates HPC's importance in driving competitiveness and profitability. It is also a reminder that while reducing costs is a good thing, any savings are typically dwarfed by HPC's outsized impact on the bottom line. Underinvesting in HPC may actually be detrimental if cost-saving initiatives impact productivity. In other words, scrimping on HPC can be "penny wise and pound foolish."

Given the importance of HPC, operators need to focus on productivity and ensure that HPC investments deliver a solid financial return. HPC operators should focus on:

- Maximizing the effectiveness and efficiency of the HPC environment (most critical)
- Minimizing TCO, as long as it does not impact productivity
- Anticipating and avoiding downstream risks that may impact future productivity or lead to future expenses

Open-source Software in HPC

While definitions vary, open-source software refers to software projects where source code is made freely available. Organizations can use, modify, or redistribute software subject to terms governed by an open-source license. Open source celebrates principles such as open exchange, collaborative participation, transparency, and meritocracy. At their best, open-source projects are agile, well supported, and responsive to changing user requirements.

[1] HPCwire, Hyperion Research 2020 – The ROI on HPC? $44 in profit for every $1 in HPC

Weighing the Pros and Cons – While open source promises rapid innovation, high quality, and freedom from license fees, this is not guaranteed. In some cases open-source software can be stale, lack features, and be led by just a handful of developers rather than a broad community. Simply allowing users to fork, change, and redistribute software doesn't necessarily make it better. Open-source software carries both benefits and disadvantages, as illustrated in Figure 2.

Figure 2 – Open-source software carries both benefits and disadvantages

Open-source development models and pricing strategies are two different things. Commercial entities may offer open-source software but monetize their efforts in other ways. These include services, paid add-ons, or commercially supported versions that offer features unavailable in a free community edition. Organizations need to have a realistic, clear-eyed view of open-source projects and form a realistic assessment of their benefits, potential risks, and actual costs.

Potential Open-source Disadvantages – Some of the potential risks and disadvantages associated with open-source projects are listed below:

1. **Project May Wither and Die** – Sometimes, open-source projects can simply stop. Developers lose interest, stop adding enhancements, and new releases grind to a halt. Others may pick up the pieces, but there is typically a significant learning curve. New project leads may have different development priorities than their predecessors and take projects in different directions. Examples of projects that have stopped (or are on life support) include ReiserFS, CentOS, and Ganglia.[2]

2. **Excessive Fragmentation** – Another risk is that projects can sometimes fork too often, leading to multiple parallel efforts, but none with critical mass. Fragmentation can arise because of disagreements between developers or because users in different industries have different requirements and prioritize different feature sets.

3. **Instability** – A typical challenge with open-source projects is that they change too quickly with too many releases, leading to overall instability. Developers often focus on "cool new features" at the expense of testing, ensuring backward compatibility, and devising clean upgrade procedures. This leads to buggy, unreliable software that is difficult to support in production environments.

4. **Infringement of Intellectual Property** – In some cases, open-source developers may infringe on someone else's intellectual property or violate another open-source license. This is a more common problem than many people realize. Stewards of open-source projects often lack the resources to undertake proper legal reviews, leaving contributors and community users in potential jeopardy.

5. **Lagging Capabilities** – Without a commercial motive to continually enhance the software, functionality may lag in new features, security, and compliance. Open-source users can find themselves living with inefficient, non-competitive offerings or needing to invest in developing new features such as cloud bursting capabilities or container support to advance AI initiatives.

[2] GitHub – Ganglia lacks active maintainers.

6. **Commercial Risks** – Sometimes, open-source projects are led by a single commercial entity. The practical result can be like a closed-source model. Business priorities can change, and commercial interests sometimes trumps open-source altruism. CentOS is a good example. The free OS, widely used in HPC, was developed downstream of Red Hat Enterprise Linux (RHEL). Red Hat recently announced the EOL of open-source CentOS 8 effective December 2021, leaving many HPC sites looking for alternatives.[3] Similarly, Elastic, the popular open-source ELK stack provider, recently decided to change licensing away from an open-source Apache license to the more restrictive Server Side Public License (SSPL) used by MongoDB. They did this to dissuade third-party cloud providers from monetizing Elastic's open-source offerings.[4] This move will have significant ramifications for users who rely on solutions downstream of Elastic's core components including a variety of monitoring and alerting applications.[5]

It is expensive to sustain a software development effort. Developers like to focus on new features, but this is just the tip of the iceberg in terms of development effort. Development teams need to worry about security patches, maintenance releases, QA, supporting new hardware and software environments, performing regression testing, and more — which means open-source projects tend to have commercial backers who will monetize their investment either by selling services directly or creating dependencies for users which will push them to buying other products.

How Free Software Can Lead to Proprietary Lock-in – Another underappreciated risk is that small or medium-sized organizations frequently have just one or two people running the HPC environment. Their job satisfaction and security are bolstered by the fact that they are critical to the organization. These well-meaning individuals frequently have an incentive to leverage open source. Selecting free software avoids up-front costs and demonstrates that they are frugal with their employer's money. Also, deploying and maintaining open-source software requires unique skills and expertise, helping solidify their value to the organization. This is not only a hidden cost in terms of the effort required to maintain the software, it is also a significant business risk when these key individuals leave the organization.

> **"Open source is not free – it simply trades up-front costs for ongoing costs and added risks."**

This dynamic becomes even more challenging when third-party integrators, consultants, or hardware manufacturers become involved, bundling and providing services around open-source software. This supposedly "open approach" can ultimately lead to higher costs as support requirements and site-specific customizations mount. Organizations often find themselves entirely dependent on their service provider. While service providers may have valuable expertise, they are usually not directly involved in the open-source effort. This means that they cannot directly solve technical problems with the software and instead resort to workarounds that ultimately add downstream costs.

Workload Management Is Key to HPC Effectiveness

Workload management is at the heart of most HPC cluster deployments. Unlike other components in the HPC software stack, workload management directly affects drivers of cost and productivity, including:

- **Scalability** – enabling larger, more thorough simulations to complete faster
- **Throughput** – improving productivity by reducing user wait times

[3] Phoenix NAP – CentOS 8 Early EOL in 2021 and CentOS Stream {What Now?}
[4] ZDNet – Elastic changes open-source license to monetize cloud-service use | ZDNet
[5] Wikipedia – In response to these changes, organizations in the Elasticsearch ecosystem including Amazon, Logz.io, CrateDB and Aiven have announced plans to fork projects, leading to a more fragmented Elasticsearch ecosystem.

- **Utilization** – helping enterprises spend less on infrastructure and software licenses
- **Personnel** – boosting productivity with intuitive interfaces and collaborative toolsets

Organizations have a choice of several open-source and commercial workload managers. Open-source solutions include OpenPBS, Slurm™, HTCondor™, and others. Commercial offerings include Altair® PBS Professional®, Altair® Grid Engine®, and Altair® Accelerator™.

Hidden Costs in HPC Workload Management – Figure 3 illustrates eight factors that can impact the TCO of workload management in HPC environments. While open-source schedulers are free, their use can lead to additional infrastructure, personnel, integration, and support costs.

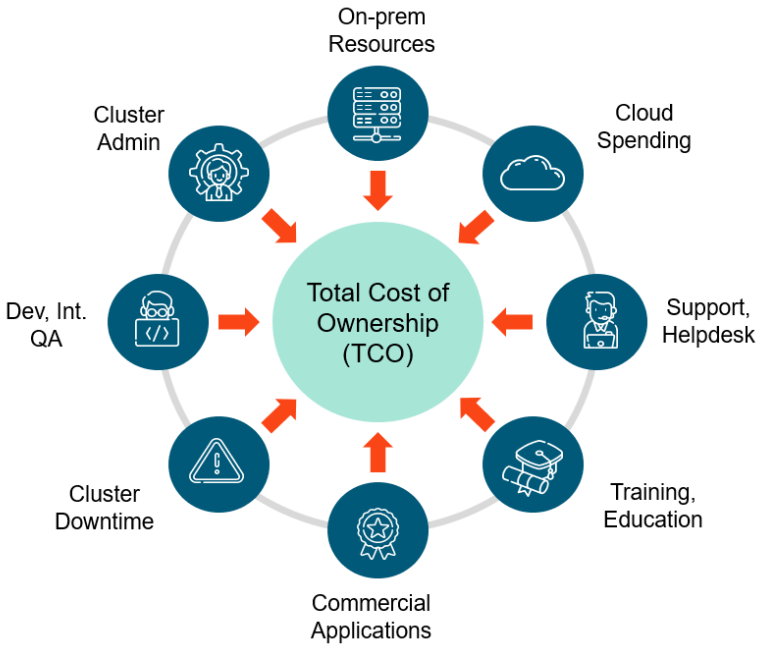

Figure 3 - Factors affecting the TCO of HPC workload managers

Some of these factors that lead to increased TCO are described below:

- **Cluster Administration** – Open-source workload managers are often more challenging to manage than their commercial counterparts. Administrators need to work around missing or incomplete features, limited diagnostics, poor documentation, and a lack of technical support. As a specific example, tasks such as adjusting scheduling policies or determining why jobs are pending can be tedious and time-consuming. Administrators can find themselves experimenting with configuration settings on live clusters and workloads, risking adverse effects. In a commercial environment, administrators would simply use a built-in simulator to predict the impact of scheduling changes and quickly get to the optimal configuration.[6]

- **On-premises and Cloud Resource Use Efficiency** – HPC infrastructure, including on-premises servers, cloud instances, and commercial software licenses, is expensive. Ideally, all of these resources should be fully utilized to maximize productivity and minimize costs. Commercial workload managers tend to have more sophisticated scheduling capabilities than their

[6] The Simulate feature in Altair PBS Professional 2020.1 and later allows administrators to take a snapshot from a PBS complex to determine the order in which jobs will run, and whether jobs will run, given different configuration settings.

open-source counterparts. Administrators can take advantage of advanced resource sharing, high-throughput hierarchical scheduling, GPU sharing, and license-aware scheduling to boost asset utilization.[7] Administrators can also leverage commercial-grade cloud-bursting capabilities. They can leverage cost-efficient spot instances, right-size instance types, and quickly marshall and shut down idle cloud resources to reduce instance charges. Even a 5-10% gain in resource use efficiency can drive substantial savings given the high cost of HPC infrastructure. This is in addition to productivity gains.

- **Development, Integration, and QA Costs** – Open-source workload managers provide core workload scheduling functionality. However, they often lack other needed capabilities. Cluster administrators may be on their own to integrate components such as web interfaces, remote visualization tools, workflow solutions, reporting and monitoring, and cloud adapters. The same is true for application integrations. Each component may have different prerequisites, and even where documented integrations exist, particular combinations of software releases are often untested. These environments can be exceptionally challenging to support and maintain.

- **In-house Feature Development** – HPC operators may find that their selected open-source software lacks a feature. They may work around this limitation by developing a custom add-on or workaround in-house. While this may be necessary in some cases, it comes at a high cost. HPC customers will quickly find themselves in the software business. They will need to integrate, troubleshoot, and revalidate their in-house software with each new software update to ensure that functionality has not regressed from earlier versions. In-house projects have a long tail in terms of maintenance costs. Customers who develop in-house customizations can also find themselves exposed when key employees or consultants leave the organization.

- **Support and Help Desk Costs** – While community support is available for open-source workload managers, it depends on the availability and goodwill of community experts. In commercial settings where time is money, relying on volunteers is a risky strategy. Problems in production settings are often complex, time-critical, and may only show up at scale or under particular circumstances. HPC administrators in commercial settings need guaranteed response times based on SLAs. Also, the support organization needs the capacity to replicate problems (at scale), quickly identify solutions, and work with developers to produce patches if required. Organizations sometimes deploy open-source workload managers to reduce cost, only to find themselves paying a premium for additional skilled in-house consultants and support personnel. The irony is that this can cost more than a commercial solution with professional support. Also, retained consultants are often not directly involved in developing the open-source project. They will lack the skills needed to troubleshoot and quickly remediate issues that inevitably arise in production.

- **Planned and Unplanned Downtime** – Another cost frequently overlooked is the cost of downtime. When the workload manager is unavailable, the entire HPC environment grinds to a halt. In a 2019 study, Hyperion Research estimated that the daily cost of HPC downtime ranged from less than USD 100K to over USD 1M per day.[8] In many environments, jobs are long-running. If these jobs are not checkpointable and fail due to issues with the workload manager, even cluster outages of just a few minutes can translate into days of lost productivity.

The Case for Commercial Workload Management

While open-source workload managers have their place, investments in commercial workload managers typically more than justify their costs given the high ROI multiples associated with performance and efficiency gains. Data from Statista shows that HPC customers are voting with their wallets and are willing to pay for improved performance and productivity. The commercial HPC middleware and application software market is expected to grow from $6.03B in 2021 to $8.43B in 2024 despite abundant open-source offerings.[9]

[7] Altair newsroom – Latest Release of Altair PBS Professional Offers Up to Tenfold Faster Performance for More Diverse and Dynamic HPC Workloads
[8] NextPlatform.com study citing estimated costs of HPC downtime due to storage failures – Unveiling the Hidden Costs of HPC Storage
[9] Statista – HPC revenue forecast by segment

Altair Commercial Workload Managers

Altair provides commercial HPC workload management solutions tailored to a variety of industry and application requirements. This means that customers have the flexibility to choose the solution that is right for their business. Altair commercial workload management solutions include the following:

Altair PBS Professional – PBS Professional is a commercially supported, industry-leading HPC workload manager. Enterprises prefer PBS Professional for its commercial licensing, support, indemnification, and the additional reliability that comes with hardened, enterprise-quality packages. PBS Professional provides advanced capabilities not available in most open-source schedulers including hierarchical scheduling, cloud bursting, forecasting, and simulation capabilities. It also provides robust security, performance, administrative, and usability-related features[10].

Altair Grid Engine – Altair also provides Altair Grid Engine, a commercial scheduler based on the successful Grid Engine open-source scheduler. Altair Grid Engine, formerly Univa Grid Engine, is an ideal choice for customers already familiar with Grid Engine variants, including Sun Grid Engine (SGE), Oracle Grid Engine (OGE), Open Grid Scheduler (OGS), or Son of Grid Engine (SoGE). Altair Grid Engine provides several advanced features, including enhanced scalability and throughput, container support, reliability and diagnosability improvements, advanced GPU scheduling, and a modern RESTful API[11].

Altair Accelerator – Accelerator is a high-throughput, enterprise-grade commercial scheduler designed to meet the complex demands of semiconductor and electronic design automation (EDA) users[12]. It provides exceptionally high throughput and low latency, enabling semiconductor firms to accelerate simulations and regression tests. Electronics manufacturers can also use commercial EDA licenses more efficiently and allocate resources among design teams. Altair Accelerator includes comprehensive policy management and a full-featured GUI providing visibility to jobs by user, workflow, and project.

Altair Open Source

Altair provides an open-source version of its popular PBS Professional workload manager, OpenPBS, for customers that prefer an open-source workload manager. OpenPBS can be freely downloaded from http://openpbs.org. This provides customers with the flexibility to start with open-source software and easily upgrade to a commercially supported offering later, depending on requirements.

A Complete Solution for HPC

HPC workload management is about more than just the core scheduler. Customers need a variety of capabilities, from web portals to workflow managers to cloud connectors to reporting and analytics solutions, as illustrated in Figure 4.

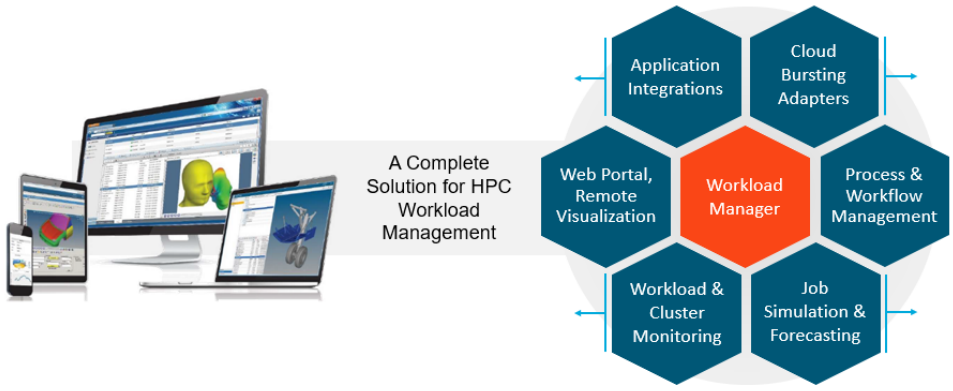

Figure 4 - Altair provides a complete solution for HPC workload management

[10] InsideHPC – Altair CTO Interview: PBS Professional Gets a Major Upgrade
[11] Altair newsroom – Altair Grid Engine – Building a Modern HPC Scheduler
[12] Altair newsroom – Altair Acquires Runtime Design Automation, Broadens Software Portfolio for High Performance Computing

Altair provides a full suite of capabilities around its commercial workload managers. Examples include Altair® FlowTracer™ (dependency management), Altair® Access™ (usability, remote visualization), Altair SAO (software asset management), Altair® Control™ (administration), and Altair® NavOps® (cloud integration).

Altair's HPC and cloud tools are all fully supported, pre-integrated, and designed to work together. Customers benefit from ease of integration, faster implementation, and capabilities not found in some open-source alternatives. They also realize savings related to development, integration, life cycle management, and support.

Multiple Sources of Savings

While commercial software helps drive significant efficiency gains leading to substantial ROI improvements of the complete HPC infrastructure, it can often be cost-justified based on TCO savings alone. Some specific areas where Altair commercial workload management solutions can demonstrate quantifiable savings are as follows:

- **Reduced Solution Integration Costs** – Using a complete, pre-integrated solution, customers can avoid the time, effort, and expenses associated with sourcing, integrating, and maintaining open-source components. Examples include web portals, monitoring tools, cloud adapters, and workflow managers.

- **Avoid In-house Development Costs** – Open-source schedulers sometimes offer only bare-bones functionality in critical areas such as containers, GPU scheduling, and commercial license management. This often forces customers to develop and maintain in-house solutions at an extra cost.

- **Improved Performance and Resource Utilization** – Commercial workload managers often provide better throughput and resource utilization than open-source alternatives. For example, Altair Grid Engine achieves between 2x and 9.5x the performance of open-source alternatives in published benchmarks.[13] With higher throughput and utilization, organizations can reduce infrastructure costs on-premises and in the cloud while also reducing user wait times.

- **Application Integrations** – Building, scripting, and optimizing application integrations can be time-consuming for cluster administrators. Activities include queue configuration, development of customized hooks, user interface templates, and more. Commercial schedulers often provide pre-tested, documented integration templates, removing this as a source of complexity.

- **Commercial Services and Support** – Commercial software providers typically offer training, implementation and configuration assistance, and support services. By taking advantage of these services, organizations can reduce administrator workload and helpdesk costs and avoid becoming overly reliant on external consultants and contractors.

- **Reduce Planned and Unplanned Downtime** – Finally, with higher-quality, more thoroughly tested software, users can reduce the risk of downtime and its associated costs in terms of lost productivity.

A Financial Model to Estimate TCO Reduction

While every HPC environment is different, it is possible to model the cost savings opportunities described above. Using a financial model, HPC operators can quickly assess the pros and cons of open source vs. commercial workload management solutions in their own environment.

Altair has developed a simple TCO model in consultation with customers to help estimate some of the hidden (and not so hidden) costs associated with open-source deployments. While results will vary by organization, the model provided in Appendix A provides a useful starting point. HPC operators can plug in their own estimates and assumptions and tailor the model to their own environment.

[13] Altair Grid Engine Benchmark Technical Brief

A sample calculation detailed in Appendix A shows that **a typical customer can achieve savings or cost avoidance in the range of $300-400K per year using a commercial workload manager[14].**

These savings are typically more than sufficient to cost-justify a commercial workload manager. Key areas of savings are shown in Figure 5. Areas of savings include infrastructure and application-related costs, administration, and costs related to development, integration, and QA.

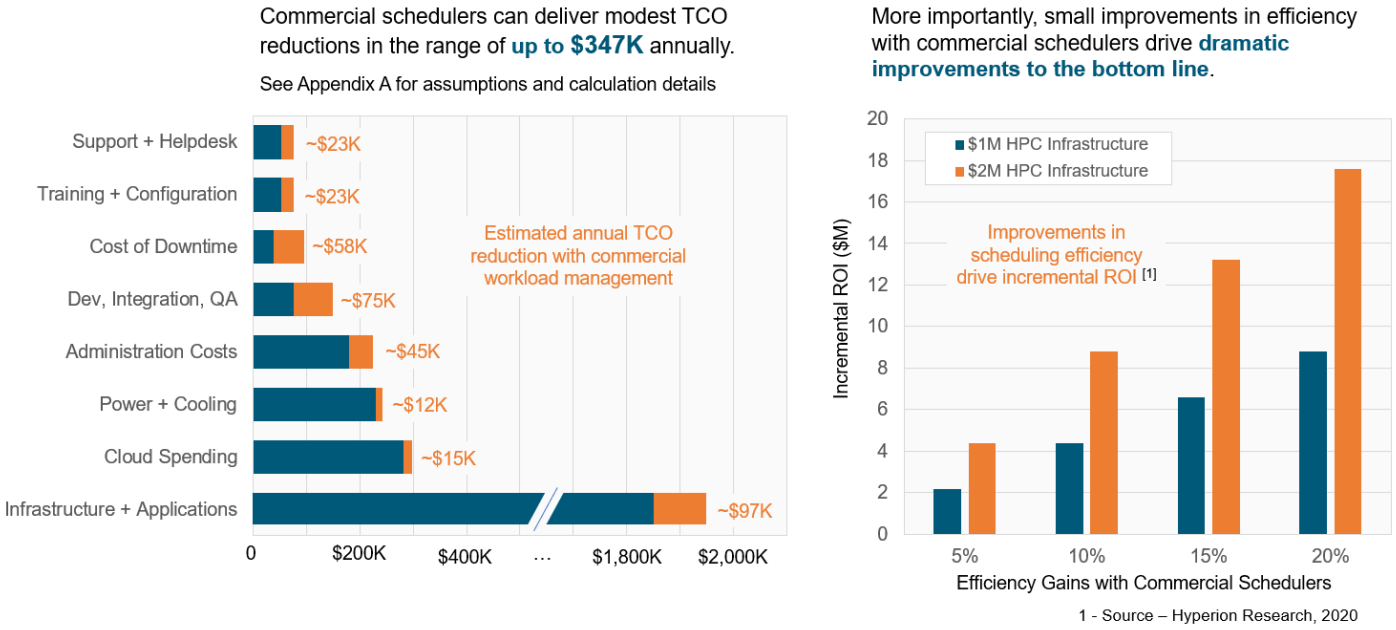

Figure 5 - Commercial workload managers help reduce TCO and improve ROI

While reducing TCO is beneficial, it is important not to lose sight of the main advantage of commercial software. Commercial solutions can often help HPC users drive higher productivity and innovate faster than competitors while avoiding risks related to support, performance, reliability, and continuity.

According to Hyperion Research, **investments in HPC on average deliver a 44x ROI**. This means that a 10% improvement in performance or productivity represents much more than just a 10% savings – it can be the difference between being first or late to market and can drive millions of dollars of bottom-line impact**. For a site that invests $2M in HPC infrastructure annually, an efficiency gain of just 10% can yield approximately $8.8M in financial returns** as shown in Figure 5[15]. Given the high level of ROI associated with HPC investments, it only makes sense to use workload management software that maximizes the effectiveness of the environment and avoids unnecessary risks.

Selecting the Right Strategy

Customers have a variety of choices when it comes to HPC workload management. Whether to use open source or commercially supported products will depend on an organization's own cost-benefit analysis.

[14] This is before factoring annual license and support costs for a commercial workload manager, which will vary. Savings opportunities will depend on the HPC environment. Assumptions and calculation details are in Appendix A.
[15] See Appendix A, Section I for ROI calculations

Enterprises doing production work (e.g., manufacturing, automotive, aerospace, oil and gas, and other Fortune 5000 businesses) will often prefer workload managers with commercial support. Public-sector organizations less focused on production work (e.g., academia and research facilities) may opt for solutions such as OpenPBS offered under an open-source license.

Customers should consider their goals, carefully evaluate costs and potential risks, and choose the solution that makes the most sense in their environment.

When evaluating solutions, it is helpful to keep the following points in mind:

- Keep in mind that open-source software supported by integrators or consultants can represent a significant source of expense and can lead to “lock-in” as site-specific customizations mount.
- Understand the tradeoffs between the “pay now vs. pay later” trap. Think in terms of TCO, ROI, and the overall effectiveness of the HPC environment.
- Take a holistic view of your infrastructure. HPC is multi-dimensional – CPUs, GPUs, memory, I/O, licenses, electricity, cloud spend, etc. – solutions that address only a few dimensions will run slower and cost more.
- Finally, carefully consider your availability requirements. Think through the impact of critical bugs, outages, and security vulnerabilities and their potential impacts.

Conclusion

HPC plays a critical role in research, design, and manufacturing. While open-source software helps accelerate innovation in many HPC disciplines, it is not without disadvantages. This is particularly true for critical software components such as workload managers that directly impact the performance, reliability, and efficiency of the HPC environment.

Given the high cost of hardware, software, and personnel, the HPC environment must operate efficiently. In production environments, hidden costs related to open-source administration, support, and lifecycle management costs can quickly mount. Ironically, open-source users can find themselves “locked in,” dependent on third parties to maintain customized environments that are ultimately less capable than commercial alternatives. The simple method of estimating TCO presented in this paper can help organizations weigh the pros and cons of open source and make an informed decision based on their own environment and requirements.

Finally, HPC operators need to remember that TCO is only part of the equation. Given the critical role of HPC in business, even small gains in workload management performance and reliability can drive productivity gains that have an outsized bottom-line impact.

To learn more about Altair’s industry-leading HPC tools and workload management solutions, visit https://www.altair.com/hpc-cloud-applications/.

Appendix A – Sample ROI and Cost Calculations

Section I. ROI Calculations Based on Scheduling Efficiency Gains

Based on Hyperion Research 2020 findings that each dollar invested in HPC yields 44 dollars of incremental profits on average.

Annual HPC Investment ($)	Efficiency Gain with Commercial Scheduler (%)	Equivalent Gain in HPC Capacity ($)	Estimated ROI on HPC ($) - 44x
$1M	5%	$50K	$2.2M
	10%	$100K	$4.4M
	15%	$150K	$6.6M
	20%	$200K	$8.8M
$2M	5%	$100K	$4.4M
	10%	$200K	$8.8M
	15%	$300K	$13.2M
	20%	$400K	$17.6M

Section II. Baseline assumptions about the HPC environment

	How many users in the environment?	
1	Number of users	40 users
2	Number of full-time HPC administrators (FTEs)	1.5 FTEs
3	FTE or consultant time spent on scripting, integration, QA	1 FTE
4	FTE or consultant time spent on internal help desk/support	.5 FTE
5	FTE or consultant time spent on training, installs, config, etc.	.5 FTE

	Cost of personnel	
6	Annual cost of full-time employee or equivalent	$150,000 per FTE

	On-premises infrastructure	
7	Capital cost per HPC rack mount server	$18,000
8	Additional component costs per server (switches, interconnects, HBAs, hardware, PDUs, etc.)	$3,000
9	Asset depreciation period (years)	3 years
10	Number of physical servers	200 servers
11	Number of servers installed per rack	20 servers/rack
12	Number of racks in the data center	10 racks
13	Capital cost per server + supporting HW infrastructure	$21,000
14	Cost per server per year	$7,000/year
15	Total annual cost of server infrastructure	$1,400,000/year

	Power and cooling cost estimates	
16	Total power consumption for HPC server	800 W
17	Power per rack	16 kW
18	Estimated production load (%)	80%
19	Total power consumption (kilowatts)	160 kW
20	PUE (power use efficiency)	1.8
21	Total power and cooling (factoring PUE)	288 kW
22	Adjusted power consumption based on load	230.4 kW
23	Average cost per kWh	$0.12/kWh
24	Total annual costs for power and cooling	$242,196

	Commercial software licenses	
25	Annual cost of commercial software licenses	$300,000

	Cloud spending assumptions	
26	Typical cloud instance type (on demand)	cgad.16xlarge (AWS)
27	On-demand hourly cost	$0.69/h
28	Estimated instance hours/month	36,000 h/m
29	Estimated instance hours/year	432,000 h/y
30	Total annual cloud spending	$297,216

	Assumptions related to downtime	
31	Estimated cost of downtime per user per hour	$200/hour
32	Planned and unplanned outages per year	6 outages
33	Average duration per outage	2 hours

Section III. Anticipated improvements with a commercial workload manager

	High-level estimates (productivity gain, time savings, etc.)	
34	How much more efficient are administrators with commercial software?	20%
35	Efficiency/throughput gains due to more efficient scheduling	5%
36	Reduction in requirement for in-house dev, integration, QA	50%
37	Reduction in requirement for in-house support/helpdesk	30%
38	Reduction for in house training, education, installation services	30%
39	Opportunity to reduce planned/unplanned downtime	60%

Section IV. Calculated annual operating expenses

	Description	Open-source costs – baseline ($)	Commercial scheduler costs ($)	Potential savings opportunity ($)
a	Annual operating cost of on-premises infrastructure (including racks, switches, software licenses, etc.) Calculation: (16)+(26)	$1,942,196	$1,845,087	**$97,110**
b	Annual power and cooling costs for on-premises infrastructure Calculation: (25)	$242,196	$230,087	**$12,110**
c	Annual cloud spending Calculation: (31)	$297,216	$282,355	**$14,861**
d	Cluster administration costs Calculation: (2) * (6)	$225,000	$180,000	**$45,000**
e	Development, integration, and QA-related costs Calculation: (3) * (6)	$150,000	$75,000	**$75,000**
f	Cost of in-house helpdesk/support functions Calculation: (4) * (6)	$75,000	$52,500	**$22,500**
g	Internal costs related to training, education, installation and configuration services Calculation: (5) * (6)	$75,000	$52,500	**$22,500**
h	Total cost of downtime (planned and unplanned) Calculation: (1) * (32) * (33) *(34)	$96,000	$38,400	**$57,600**
	Totals:	$3,102,609	$2,755,929	**$346,680**