INCREASING HPC CLUSTER PRODUCTIVITY THROUGH SYSTEM RESOURCE TRACKING

Bright Workload Accounting and Reporting

Bright Computing

Easily deploy and manage your clustered HPC infrastructure

Bright Cluster Manager for HPC lets you deploy complete clusters over bare metal and manage them effectively. Bright Cluster Manager provides single-pane-of-glass management for the hardware, the operating system, the HPC software, and users. With Bright Cluster Manager, your system administrators can quickly get clusters up and running and keep them running reliably throughout their lifecycle—all with the ease and elegance of a full-featured, enterprise-grade cluster manager. As the most powerful, vendor-independent cluster management software solution available today, Bright's intuitive management console, scalable architecture, and breadth of integration are unparalleled in the market. Bright Cluster Manager is:

- **Easy to deploy**
  - Answer a few questions, and the wizard does the rest
  - Install directly on bare metal—nothing to pre-install
  - Take the pain out of deployment
- **Easy to monitor**
  - Powerful GUI provides comprehensive metrics and alerts you when there is trouble
  - Automatic health checks keep an eye on the cluster for you
- **Easy to manage**
  - Choose between graphical or command line UI to get your work done
  - Optimize the use of your IT resources
  - Use included HPC tools and libraries

GETTING MORE VALUE FROM HPC CLUSTERS

Today, more and more companies rely on high-performance computing (HPC) to drive innovation in their products and their business. Industries such as consumer electronics, food and drug, seismology, automotive, financial services, aviation, construction, and others use HPC to complete their most complicated simulations and find answers to their most complex questions. Given this growing demand on HPC, data center administrators are under relentless pressure to extract more value from the computing resources of their HPC systems.

There are several steps that can be taken to increase the productivity of HPC systems.

First, administrators that can readily determine how users' jobs consume resources such as memory, CPUs, and GPUs have a distinct advantage by identifying users and jobs that are wasting resources that can be reclaimed for use by others. It is not uncommon for users to over-request resources for their jobs, either because they don't know exactly how much their job needs, or intentionally to hold resources until they need them. In either case, valuable resources are allocated to jobs but go unused, which prevents them from being used by other jobs. These cases need to be identified so they can be eliminated.

Another effective way to increase HPC system productivity is through aggregation. Many system administrators are now consolidating siloed HPC systems into a single shared system that aggregates unused resources on a broader scale to service more workloads. While consolidating silos into a single, shared HPC cluster generally benefits all stakeholders, it is usually beneficial to provide users with "showback" reports that quantifies usage and the benefits provided by aggregating.

the most powerful, vendor-independent cluster management software solution available today

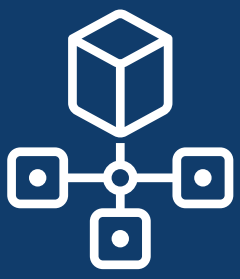

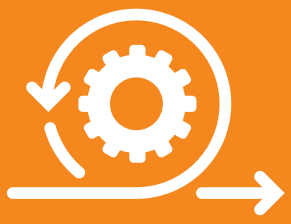

And finally, maximizing the effective use of an HPC system requires users to be good stewards of the system's resources. In most cases, users consider these resources to be "free" and accessible without any accounting of their cost. In these cases, a "chargeback" report that accounts for system use by job, individual, or department provides visibility and awareness that supports good stewardship.

Bright Cluster Manager's Workload Accounting and Reporting (WAR) capability was created to address the steps discussed above: To give administrators, managers, and users the information they need to use HPC system resources effectively, to maximize system productivity, to enable effective resource sharing, to identify waste and to provide chargeback capability.

IDENTIFYING RESOURCES BEING WASTED

The first step toward making your HPC environment more productive is "looking under the hood" to determine which workloads are wasting resources. Using compute resources to run workloads is certainly not wasting them. But what if a workload uses only a fraction of the resource, and by doing so, prevents other workloads from running?

Identifying those situations allows administrators to change system parameters to enable lighter workloads to run on less powerful compute nodes, and heavier workloads to run on more powerful compute nodes. This way, both types of workloads can run simultaneously, and the cluster becomes more productive.

Table 1.0 on the following page illustrates how Bright Workload Accounting and Reporting identifies workloads that are occupying resources such as GPUs, but not actually using them (or using them inefficiently).

to maximize system productivity,
to enable effective resource sharing,
to identify waste

Table 1.0
Bright WAR illustrates wasted HPC compute resources

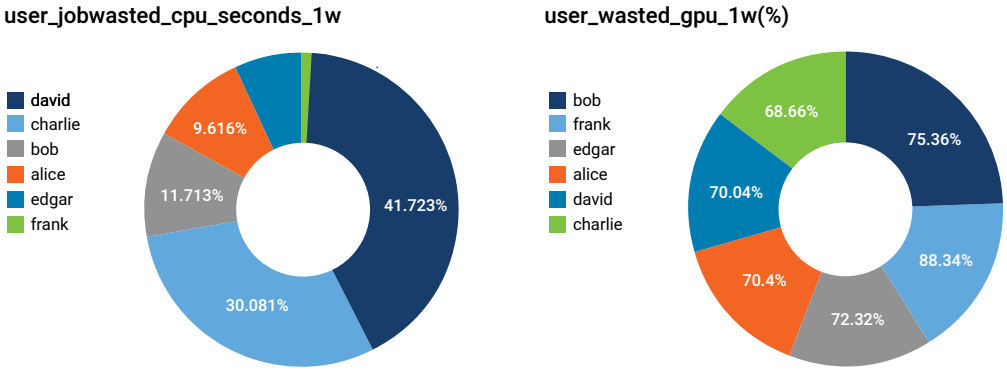

The report shown above illustrates that jobs run by David use about 42% of the allocated CPUs, while jobs run by Charlie use about 30%. Perhaps, more importantly, none of the jobs are using the GPUs efficiently, but Bob is wasting significantly more resources than anyone else.

Investigation using Bright's integrated job monitoring capabilities reveals the problem. Bob's Convolutional MNIST jobs are running, but they are occupying a GPU without actually using it. The matrix in Figure 1 (below) shows one of Bob's Convolutional MNIST jobs on the left using resources effectively, while the right side shows a more recent job not using a requested GPU.

Figure 1.
Bright job metrics shown in Bright View

PROMOTING RESOURCE SHARING

Over the years, computing requirements have changed, and data centers have evolved to meet those requirements. While this evolution has offered many positive benefits, there is at least one negative side effect—siloed compute clusters. Siloed infrastructures emerge as companies purchase new systems rather than expand the capacity and capabilities of existing systems (in this case, HPC clusters). This often occurs when new departments procure their own systems for a project, or when a new generation of technology is required to serve more demanding workloads. While these are valid reasons to purchase new systems, there are consequences of having these systems operate as siloed (stand-alone) clusters:

- **Unrealized productivity gains:** Siloed systems prevent idle capacity from being utilized by other users with pending jobs that don't have access to the system.
- **Higher administrative expense and burden:** Siloed systems often have separate designated administrators, whereas aggregated systems can be managed by a single administrator. In addition, siloed systems often use different management tools, which increases complexity and effort for the organization.
- **Inability to concentrate compute:** Siloed systems prevent organizations from aggregating and maximizing the total capacity that can be brought to bear on important and time-sensitive jobs when needed.

This is why Bright Computing advocates creating a single shared infrastructure, and why Bright strives to ensure that customers can run all types of workloads on that infrastructure.

When using Bright Cluster Manager to create a single shared computing infrastructure, WAR reports promote efficient resource usage by showing users and administrators how resources are being allocated and how effectively they are being used.

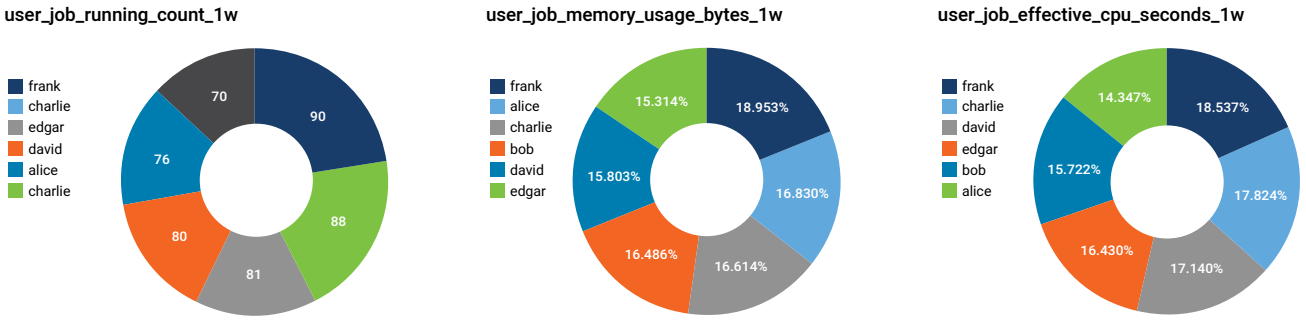

Table 2
Bright WAR illustrates shared HPC resources across the cluster

ENCOURAGING PRODUCTIVE USE OF RESOURCES

When siloed clusters are combined into a more productive shared computing cluster, users benefit from greater availability of computing resources, and they are able to run more jobs. Running more jobs faster accelerates time to results and time to market for the business.

While running more jobs is a good thing, there is a point of diminishing returns, where each subsequent job produces increasingly smaller benefits. When resources are viewed as "free," users might use them even if they receive only minimal benefits. In this case, the cost to run a job might be more than it is worth.

Bright WAR encourages users to utilize HPC resources productively. By generating a charge-back (also called show-back) report and assigning a monetary value to each unit resource (e.g., MB*sec, or CPU*sec), each stakeholder can be charged for their team's use of the shared resources.

Bright WAR charge-back reports result in good stewardship of the company's HPC resources. Even if stakeholders are not actually charged, the reports provide insight into how much the work completed by each group costs. Stakeholders can use this information to make well-informed value judgments.

In addition to charge-back reporting, other measures—such as configuring fair share in the workload management systems—can, and generally should, be used to encourage efficient resource usage.

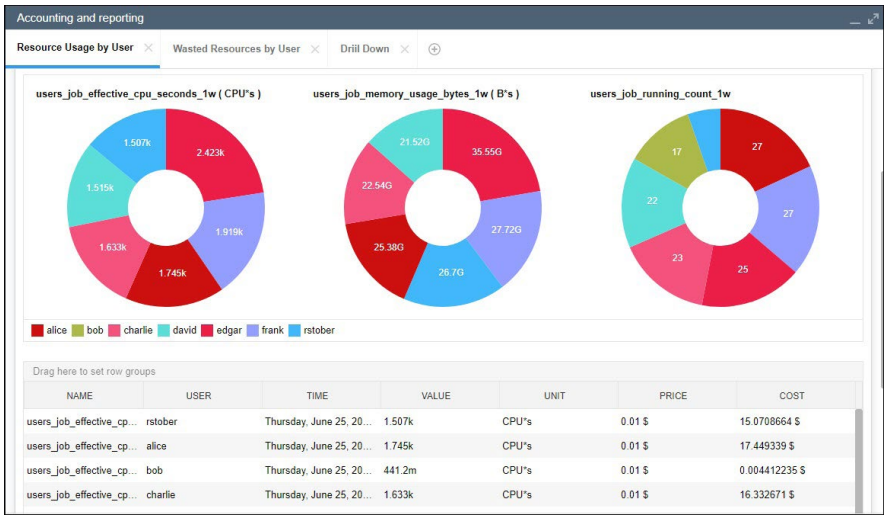

Table 3.
Bright WAR shows charges for resources in-line with the report

A charge-back report allows organizations to assign costs to each compute unit in two ways: The report can be used to (1) charge groups for their use of resources, or (2) track the costs of projects.

CONCLUSION

Companies that use HPC to drive innovation in their products and business know that computing clusters are critical to their success. When appropriately managed, HPC clusters can give today's businesses a significant competitive advantage.

Bright WAR provides the information you need to maximize the productivity of your HPC compute clusters. Bright WAR:

- Helps you identify resources being allocated, but not used efficiently
- Shows you who is using HPC resources and how much they are using
- Allows you to (optionally) charge groups for their resource usage

If you already use Bright Cluster Manager to manage your compute clusters, Bright WAR is included with your Bright software subscription. You can begin using Bright WAR today to extract more work and value from your existing HPC clusters.

If you want to try Bright Cluster Manager for yourself, a full-featured version is available for free for clusters up to 8 nodes through Bright's Easy8 program.

For more information about Bright Workload Accounting and Reporting or to learn more about Bright Cluster Manager, please contact us directly at info@brightcomputing.com

Tel: +1 408 300 9448
info@BrightComputing.com
www.brightcomputing.com

Bright Computing