



HIGH-PERFORMANCE OPEN COMPUTING: A NEW GENERATION OF BIG DATA-DRIVEN INSIGHTS FOR THE FEDERAL GOVERNMENT

EXECUTIVE SUMMARY

Federal government agencies are demanding big data insights at unprecedented levels. Even after decades of steady improvements in processor technology, data demands are now outpacing the performance capabilities of traditional servers. But high-performance open computing servers offer a promising alternative.

This white paper presents the value of high-performance open computing for government organizations. Groups such as the OpenPOWER Foundation, formed by IBM and dozens of other companies, are working to surpass the boundaries of traditional chip technologies. Innovations from this collaborative group, combined with complementary advancements from the open source community, are enabling new levels of performance, scalability, and security.

The paper explains where high-performance open computing fits into today's government organizations. It provides examples of the impact this emerging technology can have on everything from weather forecasting to national security. The paper also highlights the technical advantages of high-performance open servers, such as IBM Power Systems, and best practice recommendations for deploying and managing these systems.

Ultimately, this paper will help guide federal government organizations on the journey to transform their infrastructures with high-performance open computing. These transformations will usher in a new age of big data-driven applications that enable real-time decision-making that was never before possible.

INTRODUCTION

For decades, federal government agencies have enjoyed steady advances in high performance computing. Like clockwork, processor performance doubled every 18- 24 months, feeding a nearly insatiable appetite for big data insights. But the limits of physics have slowed that consistent pace, and with it, the ability to achieve predictable performance gains. A new approach is crucial to continue the pace of advancement.

To complicate the issue, demand for insights out of ever-expanding data is now out-pacing the ability of technology to deliver those insights. Traditional system architectures have reached a saturation point, and incremental performance gains are harder and harder to achieve. Yet, data continues to explode – especially unstructured data, whether it is from Twitter feeds, video capture from UAVs, or facial recognition technologies.

To keep up, agencies typically have thrown more hardware at the problem, spreading workloads across multiple systems. This has led to escalating overhead and management costs. The fact is, more physical machines mean more maintenance, updates, bug fixes, and security patches—essentially more of everything.

There is an alternative to these high-end applications. Many federal agencies are turning to high-performance open source computing, including IBM Power Systems running Linux. This approach is proving to be an economical solution for gaining the extreme speed, high availability, scalability, and uncompromising security needed in today's big data-driven government environments.

To help overcome the physical limitations of current chip technology, IBM formed the OpenPOWER Foundation, which at the time of this writing has grown to over 200-plus companies working collaboratively in an open forum to contribute to technology advances at a much faster rate than any one company could alone. Member companies are continuously innovating on POWER CPU processors and developing custom solutions for workload acceleration through specialized co-processors and advanced I/O techniques. (Figure 1)

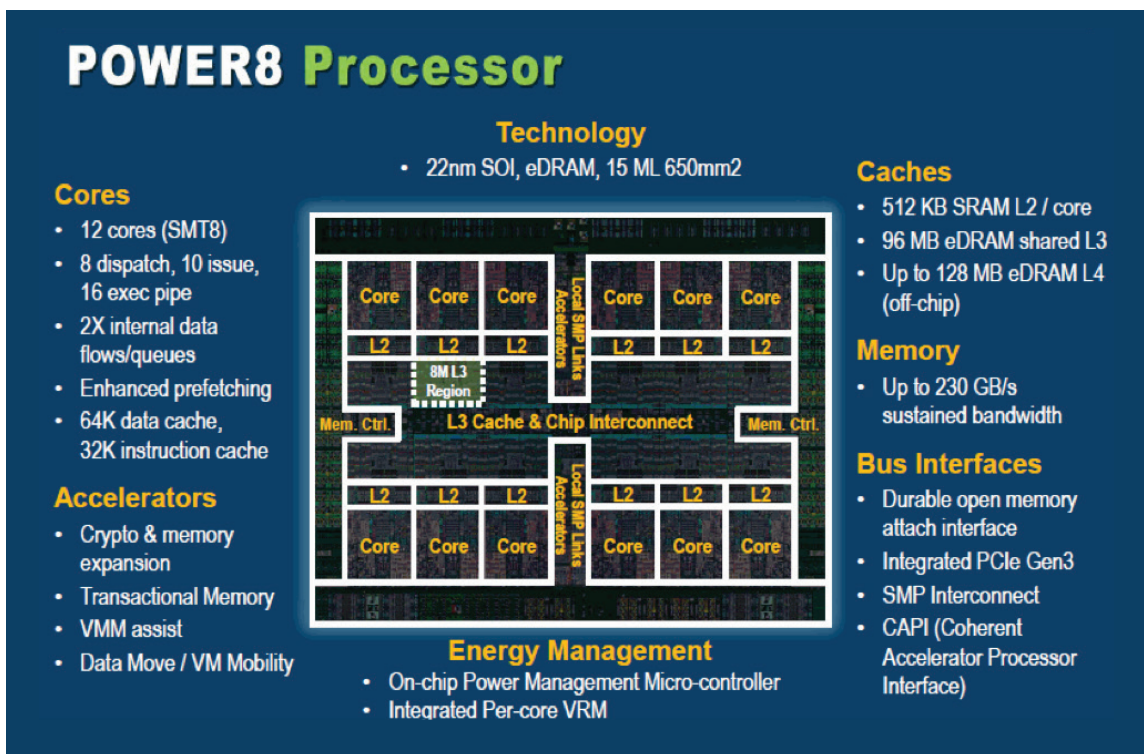


Figure 1 – IBM POWER Collaborative Processor Architecture. Copyright IBM.

This open ecosystem is producing rapid innovations that enable high-performance open computing servers to keep pace with big data growth. Like the open source community, the OpenPOWER Foundation operates efficiently, so the resulting technologies are open and affordable. When combined with other open technologies such as Linux, these systems not only deliver the robust capabilities government agencies require, they also fit within already strained budgets.

HIGH-DEMAND ENTERPRISE-CLASS APPLICATIONS

High-performance open computing is an ideal fit for enterprise-class applications with seriously demanding workloads. That's because it can process massive amounts of structured and unstructured data almost instantaneously.

Agencies traditionally consolidated data from multiple sources into large data warehouses and then tried to extract meaningful intelligence. The old methodology could take anywhere from a few hours to days—far too long for results that could impact public health and safety, or even national security.

In contrast, with high-performance open computing, data can stream continuously into an in-memory database such as SAP HANA, MongoDB, or GPUdb to produce real-time, actionable insights. There are many popular applications suited for these environments, including the entire IBM portfolio (IBM Cognos business intelligence, IBM BigInsights for Apache Hadoop, IBM WebSphere service-oriented architecture, and others).

The human impact of deploying these solutions within the federal government is profound. For example, consider predictive earthquake analysis. Imagine being able to continuously analyze geological and atmospheric conditions to determine with great precision when and where an earthquake might be about to occur. Alerting communities hours, or even days in advance could save countless lives.

UNBOUNDED SCALABILITY FOR UNPREDICTABLE WORKLOADS

In many cases, it's difficult to anticipate future compute requirements. Many agencies face uncertain data growth, which means any new additions—however small—must be capable of on-demand growth. This necessitates highly scalable and flexible infrastructures that can take on additional demands without major revisions to their fundamental design or capabilities.

"...high-performance open computing systems enable agencies to perform more work with fewer physical machines. It's more cost effective to carve virtual environments from a single large asset than to run multiple smaller assets. This provides more operational compute environments using fewer sockets, which often means lower licensing costs. Plus, a consolidated environment is easier to manage, leading to additional savings in time and money."

Data collection sources are constantly expanding with virtually no limit. For the federal government, these sources can range from buoys in the ocean, to cell towers on land, to satellites in the sky. With the rise of wearables and the Internet of Things, every person, place, or thing may be a source of streaming information. Traditional server technologies are not designed to handle this kind of exponential data explosion, and consequently, they will be unable to perform as needed without huge escalations in complexity and cost.

Alternatively, high-performance open computing systems are specifically designed to meet rapid scalability demands.

Because of their speed and flexibility, high-performance open computing solutions like Linux on IBM Power Systems are gaining significant traction in federal government. From creating new weather models to simulating nuclear stockpiles, the U.S. government is leading the next wave of exascale computing. In the process, it often turns to IBM Power Systems on Linux.

NEW HEIGHTS IN PERFORMANCE, ECONOMY, AND SECURITY

There are strong incentives for government agencies to adopt high-performance open computing. Among the most compelling reasons is that solutions like Linux on IBM Power Systems deliver higher performance per dollar than traditional server infrastructures. In fact, federal agencies not only get better value out of the box, but they also can expect lower total cost of ownership over the long term.

That's because high-performance open computing systems enable agencies to perform more work with fewer physical machines. It's more cost effective to carve virtual environments from a single large asset than to run multiple smaller assets. This provides more operational compute environments using fewer sockets, which often means lower licensing costs. Plus, a consolidated environment is easier to manage, leading to additional savings in time and money.

Most important, high-performance open computing provides a stable, reliable, and secure environment. IBM Power Systems, for example, are widely recognized for their highly secure architecture, which meets the most stringent security certifications of federal agencies (ibm.com and cve.mitre.org). IBM also provides a security and compliance solution designed for virtualized environments on Power Systems. It should be no surprise that the U.S. Computer Emergency Readiness Team (US-CERT) issues far fewer advisories for Power Systems than for any other server solution, especially when coupled with the IBM PowerVM hypervisor.

On top of this secure architecture, Linux brings another layer of security and stability. Because Linux is continually monitored and refined in the global open source community, vulnerabilities are discovered and addressed more quickly than they are in proprietary software developed by a single company. Such close, 24x7 scrutiny of the code has made Linux highly robust and trusted for many of the most demanding environments in the U.S. government.

THE VALUE OF WORKING WITH A SEASONED SYSTEMS INTEGRATOR

Embarking on a transformation to high-performance open computing may seem daunting at first. It is a different mindset and new technology. That's why it pays to work with a seasoned systems integrator with deep technical knowledge and broad federal government experience. Having the right consulting expertise

to guide your journey is the best way to ensure a smooth, high-quality deployment that is aligned with your specific mission and organizational requirements.

Jeskell brings extensive experience integrating best-of-breed technologies from top vendors, including IBM, Cisco, Dell, Brocade, Lenovo, Dell, VMware, and Red Hat. Our solution architects and project managers are experts in big data and high-performance computing. Moreover, our staff is widely populated with former federal government engineers. So we offer a valuable combination of government perspective and in-depth technical expertise that allows us to capably lead projects from concept to production while avoiding pitfalls that can derail transformation projects.

Our company has had proven success in many system integration projects across the U.S. federal government. Our government clients trust us to provide sound judgment and advice on solving vexing problems with the right solutions. If Jeskell recommends IBM Power Systems on Linux, you

“ Because Linux is continually monitored and refined in the global open source community, vulnerabilities are discovered and addressed more quickly than they are in proprietary software developed by a single company. Such close, 24x7 scrutiny of the code has made Linux highly robust and trusted for many of the most demanding environments in the U.S. government”

can be assured it’s the best fit for your needs. You can also count on Jeskell to deploy the solution quickly and effectively, along with knowledge transfer and responsive technical support to ensure your ongoing success.

BEST PRACTICE RECOMMENDATIONS

Drawing on extensive experience working with government agencies, Jeskell has developed a set of best practice recommendations for government organizations moving to a high-performance open computing environment. By adhering to these recommendations, agencies can gain greater assurance of achieving a successful transformation:

Take a holistic approach

It’s important to consider all components of the infrastructure—including compute, storage, networking, security, and applications—and how they are integrated. This is increasingly critical as more elements are connected inside and outside an individual agency. A holistic approach provides a clearer understanding of system interdependencies, which will determine precise I/O requirements and guide the most appropriate and secure solution configurations.

Ensure appropriate staff knowledge and experience

Staff involved in planning, design, deployment, and administration of high-performance open computing environments should have knowledge and experience that spans servers, storage, and networking. This is critical to building and maintaining a properly balanced infrastructure that delivers the required performance, stability, and cost effectiveness.

Design to application requirements

The high-performance open infrastructure should be configured to match the technical requirements of the application and characteristics of the application workloads. Ultimately, the needs of the application will dictate whether to run it on multiple compute nodes across a geographically dispersed region or on a single scale-up system.

When addressing these best practice recommendations, consider working with a systems integrator that can bring them to the table on day one. This is a reliable way to streamline solution delivery, reduce risk, and save money over the long term.

CONCLUSION

Demand for big data insights will only accelerate as big data continues its exponential growth. High-performance open computing is the key to getting the most out of what big data has to offer. As compute power grows to process increasingly greater volumes of data in shorter periods of time, federal agencies will gain unprecedented capabilities to analyze ever-more granular sets of data.

Imagine being able to analyze tens of millions of temperature, pressure, and humidity readings continuously collected from around the world. This could have enormous implications for the accuracy of large-scale meteorological patterns, from hurricanes to global climate change. It also could enable pinpoint accuracy of localized weather conditions to determine, for example, exactly when and where a tornado is most likely to touch down.

Weather is just one example. Simply put, big data and high-performance open computing hold the potential to uncover new insights that were never before possible. Already, technologies such as IBM Power Systems on Linux are proving to be enablers of advanced big data projects. These systems are collecting and processing data at a steadily increasing rate, resulting in faster decision making. As open technology continues to offer advancements in performance, scalability, and security, it presents an unprecedented opportunity for federal agencies to improve quality of services and safety for U.S. citizens.

DOES YOUR AGENCY HAVE QUESTIONS ABOUT
HIGH PERFORMANCE COMPUTING?

Contact Jeskell today for a free consultation.
info@jeskell.com or 1-877-JESKELL