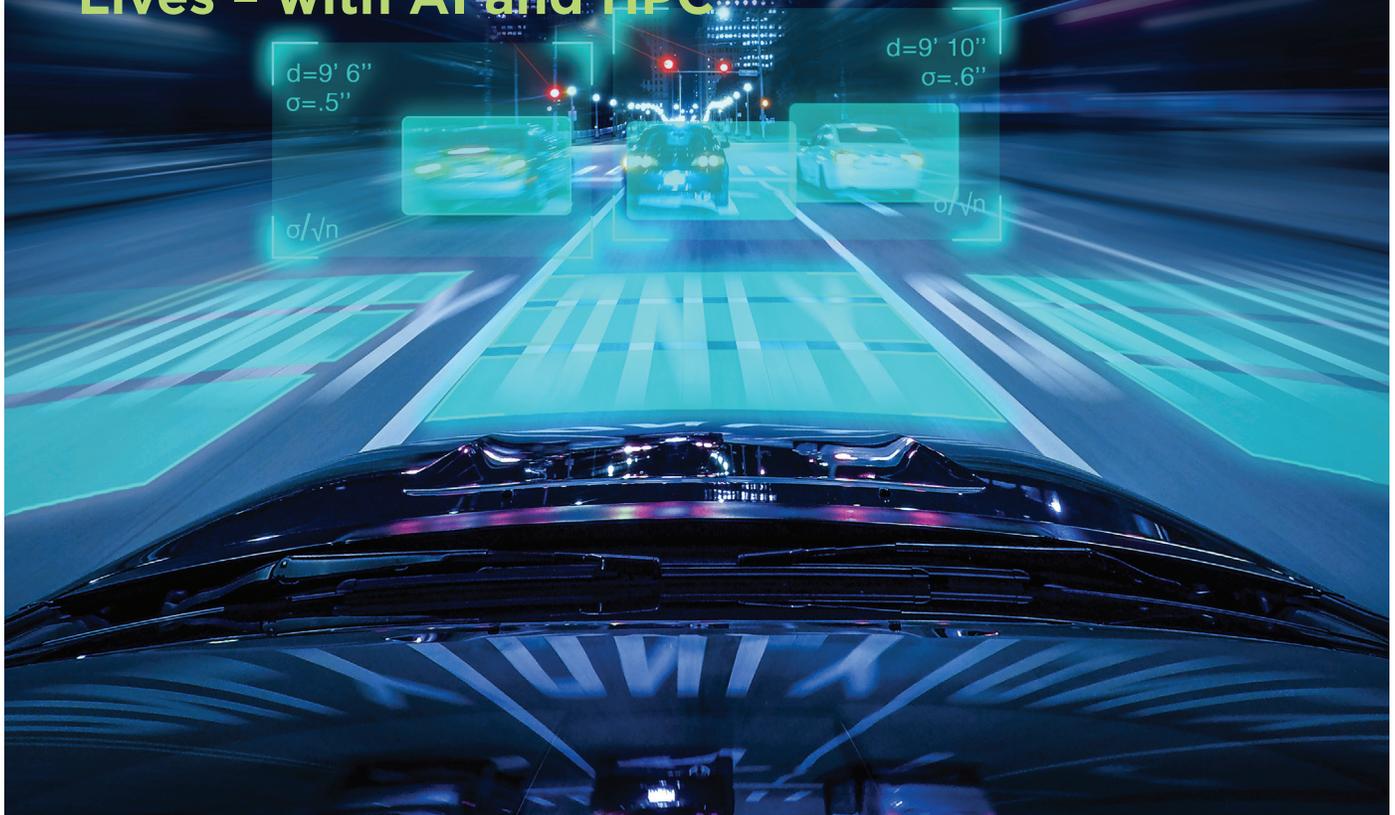
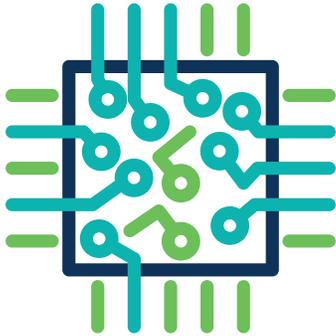


Powering Artificial Intelligence with High-Performance Computing

Transforming Businesses – and Lives – with AI and HPC





Disruptive technologies such as cognitive computing, the Internet of Things, and smart cities are powered by high-performance computing and fueled by advanced data analytics. Businesses around the world today are recognizing that an HPC infrastructure is vital to supporting the AI and analytics applications of tomorrow.

It's hard to find an IT topic hotter than artificial intelligence (AI). The huge volumes of data that enterprises can collect today offer enormous potential, and AI applications are helping to tap into that potential. That is why these apps have already become vital across industries. But as data grows and enterprises rely on AI more and more, underlying systems are put to the test. To keep up, enterprises are starting to adopt the techniques and technologies of high-performance computing (HPC).

HPC is no longer a world of custom supercomputers. An HPC infrastructure is within reach of any enterprise. If you're wondering how to support your evolving AI applications, HPC may hold the answer. Let's look at how you can use software-defined infrastructure based on HPC to power your AI workloads and turn your data into value.

What We Mean by Artificial Intelligence

For the sake of this paper, we're broadly defining AI as any system that is trying to solve the kind of problem—such as recognizing patterns—that human thinking does. AI encompasses machine learning (ML), which is a process in which examples are used to train a computer to recognize patterns, and deep learning (DL). Deep learning is a branch of ML that uses digital neural networks to create systems that learn on their own.

Artificial Intelligence Is Everywhere

The lack of sophistication about AI in news reports belies the truth: AI is already an important tool for many organizations. Organizations are eager to take advantage of data from the Internet of Things, smart cities and more, and AI offers one of the best ways to perform data analysis on large and unstructured data sets.

In finance, AI trading systems have already proved successful enough that firms—including BlackRock, the largest fund company in the world—are handing over funds to these systems.¹ The finance industry is also using AI for automated credit scoring, loan analysis and real-time fraud detection.

AI trading systems have already proved successful enough that firms—including BlackRock, the largest fund company in the world—are handing over funds to these systems.

AI is also remaking healthcare, a field with an enormous amount of data that can benefit from the strength of ML and DL approaches. SigTuple is using AI to analyze medical samples² and the FDA has even cleared, for the first time, a cloud computing and DL solution for use in a clinical setting.³

From self-driving vehicles to online chat bots, AI is everywhere. But as these examples show, it's not just ubiquitous: It's also remaking the very world we live in.

1 <https://www.nytimes.com/2017/03/28/business/dealbook/blackrock-actively-managed-funds-computer-models.html>

2 <https://www.forbes.com/sites/chynes/2017/08/31/four-companies-using-ai-to-transform-the-world/#623064264038>

3 <https://www.prnewswire.com/news-releases/arterys-receives-fda-clearance-for-the-first-zero-footprint-medical-imaging-analytics-cloud-software-with-deep-learning-for-cardiac-mri-300387880.html>

Powering Artificial Intelligence and the Future of Business

The future of your enterprise likely depends on how you can embrace disruptive technologies such as AI. To transform enterprises, however, AI solutions must run reliably and smoothly. That is easier said than done.

AI systems require a lot of computing power. As volumes of data grow, so do training times and computational requirements. According to Dileep George, co-founder of the machine-learning startup Vicarious, at least 80 percent of the recent advances in AI can be attributed to the availability of more computer power.⁴ Budgeting these computational resources can become one of the primary challenges for AI systems.⁵

Enterprises have already experienced minor setbacks in attempts to embrace AI. In healthcare imaging, demand for infrastructure has pushed some organizations to look to advanced computational offerings in the cloud.⁶ The truth is that the future of AI requires an HPC platform designed for parallel computing. An HPC infrastructure allows your organization to scale computationally and build ML algorithms that can take advantage of huge amounts of data.

The truth is that the future of AI requires HPC.

Why High-Performance Computing?

HPC systems excel at tackling data-intensive tasks, while their ability to run processes in parallel can drastically reduce the time it takes to perform functions. Features such as advanced multi-pathing and I/O capabilities can help speed large and demanding workloads.

Because the HPC community pioneered many of the hardware, software and programming advances that will help power AI, experts see an increase in new interest from nontraditional HPC customers.⁷ This is one reason the HPC market is projected to experience a 6.2 percent compound annual growth rate between 2016 and 2021, to a total value of \$30 billion dollars.⁸

How to Bring HPC to Your Organization

HPC today no longer implies bespoke supercomputers built for government and academic research institutions. In fact, enterprises can build high-performance infrastructures while dealing with far fewer vendors and far less customization than most realize.



Today's HPC systems are based on Linux clusters running on industry-standard x86 or Aarch64 hardware—the same kind of Linux clusters enterprises are likely already using for their big data and cloud architectures.

Linux runs on 100 percent of the world's top 500 supercomputers.⁹

The HPC community is making it easier than ever to implement a software-defined HPC infrastructure on the hardware you already have. The OpenHPC group and its members work to simplify the task of adopting HPC techniques and technologies and applying them to enterprise tasks. Below are seven steps you can follow to start your HPC journey.

1. Plan Your Workload

Your obvious first step is to decide what you are going to run on your HPC cluster. Consider not only what you will start with, but also how you expect to scale the application or applications.

2. Repurpose Hardware

Your next step is to determine your hardware needs. As mentioned above, HPC can be implemented today as a software-defined infrastructure on top of existing hardware. That means you likely don't need to make any changes to your hardware, nor to the physical location in which it resides.

3. Maintain Your Networking

If you already have a Linux cluster, your current, existing top-of-rack switching and Ethernet networking will likely suffice for your HPC infrastructure. While more advanced HPC interconnects are available from multiple vendors, you can likely avoid those investments for now.

4 <https://www.technologyreview.com/s/513696/deep-learning/>
 5 <http://serialmetrics.com/blog/common-machine-learning-challenges/>
 6 <https://www.nextplatform.com/2017/11/30/hospitals-untangling-infrastructure-deep-learning-projects/>
 7 <https://www.nextplatform.com/2017/11/14/samsung-invests-cray-supercomputer-deep-learning-initiatives/>
 8 Trends in the Worldwide HPC Market, Hyperion, Nov 2017.
 9 <https://www.top500.org/statistics/details/osfam/1>

4. Take Advantage of Your Existing Storage

You probably do not need new storage to support your HPC system. In fact, for many organizations, it will be the other way around. HPC systems can support applications already working on the data in your current storage environments.

If you are currently planning changes or expansion to your storage, consider a software-defined storage environment, which would allow a single system to support HPC and other requirements in your data center.

5. Select the Right Middleware

Much of the magic in HPC happens at the middleware layer, where you enable parallel computing through tools such as workload schedulers and Message Passing Interfaces (MPIs). Luckily, many of the tools that once came from various vendors are now available as open source.

With SUSE, those tools are even packaged together. SUSE, Linux Enterprise Server for High Performance Computing includes the HPC software packages you need (such as the parallel file system) and supports multiple MPIs.

6. Manage Your Cluster

To keep your HPC cluster performing, it's vital to have a parallel cluster management solution. The leading open-source option is SLURM, which is also included in SUSE Linux Enterprise Server for High Performance Computing. Container management is also becoming important in some HPC environments. A popular option, Singularity, is available in the SUSE package hub.

7. Plan for Expansion

The faster and better your AI systems run, the more your enterprise will come to rely on them. At the same time, you can only expect data volumes to grow. That means one thing is sure: Your enterprise will soon be calling for an expansion of your HPC infrastructure.

“We have relied on SUSE Linux Enterprise Server in high-performance computing for approximately 15 years, and have always been very satisfied with the operating system.”

Herbert Huber
Division Head of High Performance Systems
Leibniz Supercomputing Centre

Some Additional Considerations

The Cloud

If your organization is committed to the cloud—perhaps with a cloud-first strategy—you don't have to pass up the benefits of HPC. HPC-as-a-service offerings give you the benefits of HPC with the convenience of the cloud.

Building a Business Case

You will likely need to build a business case for your move to HPC. This should at minimum include two things: the value of your AI workloads to the business and the significant cost savings that you will gain by reusing existing hardware. This limiting of capital expenses will go a long way in helping to justify your infrastructure.

Reducing Risk

One of the best things you can do to reduce the risks involved in transitioning your infrastructure to HPC is to work with a trusted vendor. SUSE has been a leader in enterprise Linux since its founding in 1992. SUSE Linux Enterprise Server currently runs on half of the world's top 50 supercomputers and SUSE is the commercial Linux leader in supercomputing's prestigious Top 500.

Changing the World through High-Performance Computing and Artificial Intelligence

What could an HPC infrastructure mean for your AI goals and your business? It could help you join the ranks of companies changing the world. Behind the growth of innovative AI systems are HPC infrastructures. They have moved beyond the realm of government and academia and will transform big data analytics and generate disruptive solutions that change our very way of life. The only question is whether your enterprise will be along for the ride.



**For more information,
contact your local SUSE
Solutions Provider, visit us
online or call SUSE at:**

1-800-796-3700 (U.S. and Canada)
1-801-861-4500 (Worldwide)

SUSE
1800 South Novell Place
Provo, UT 84606

SUSE
Maxfeldstrasse 5
90409 Nuremberg
Germany

www.suse.com