

# Red Hat, NVIDIA Align on Open Source Solutions to Fuel Emerging Workloads

OCT 23, 2018

*Red Hat's leading Linux and Kubernetes platforms now available on NVIDIA® DGX-1™ systems, companies extend open source community collaboration to help power new workloads like artificial intelligence*

RALEIGH, N.C.--(BUSINESS WIRE)-- Red Hat, Inc. (NYSE: RHT), the world's leading provider of open source solutions, today announced it is collaborating with NVIDIA to bring a new wave of open innovation around emerging workloads like artificial intelligence (AI), deep learning and data science to enterprise datacenters around the world. Driving this effort is the certification of the world's leading enterprise Linux platform, Red Hat Enterprise Linux, on NVIDIA® DGX-1™ systems. This certification provides a foundation for the rest of the Red Hat portfolio, including Red Hat OpenShift Container Platform, to be deployed and jointly supported on NVIDIA's AI supercomputers.

AI is often looked at as an enabling technology for enterprises seeking to gain a competitive edge and market differentiation through digital transformation. Research firm Ovum highlights the growing interest in AI, stating that, "The proportion of enterprises with active interest in AI initiatives (i.e., planning, trialing, or deploying) shifted from 62% last year to 77% this year [2018]."

As NVIDIA DGX-1 systems enter the datacenter to provide powerful AI platforms, IT teams will often want to manage and maintain these systems within their existing operations, which frequently are Red Hat Enterprise Linux-based. The ability to use Red Hat Enterprise Linux on NVIDIA DGX-1 hardware provides not only a familiar control plane and SELinux-secured user environment for IT teams to weave into their existing workflows, but can also serve as a catalyst to drive further innovation with AI and deep learning.

Enterprises across many industries, including those in the public sector, oil and gas, financial services and more, will now be able to more effectively integrate NVIDIA architecture into existing infrastructure environments with Red Hat Enterprise Linux acting as the common "bridge." In support of greater flexibility across enterprise environments, current Red Hat Enterprise Linux customers will be able to use existing subscriptions on NVIDIA DGX-1 systems, backed by both Red Hat and NVIDIA. Additionally, organizations using Red Hat Enterprise Linux on NVIDIA DGX-1 systems will also have a foundation for the industry's most comprehensive enterprise Kubernetes offering in Red Hat OpenShift Container Platform, extending cloud-native support to emerging workloads.

ISVs can move existing Red Hat Enterprise Linux-certified applications to NVIDIA's new class of systems with little-to-no changes, gaining easier access to advanced AI and deep learning capabilities to support more intelligent workloads.

In the realm of high-performance computing, Red Hat and NVIDIA deliver technologies and expertise to fuel two of the world's fastest supercomputers, [Summit](#) and [Sierra](#). As part of a broader group of industry leaders, Red Hat helped bring both of these systems online providing a model not only for next-generation high-performance computing (HPC) in performance-sensitive environments but also for the enterprise datacenter of the future.

Beyond Red Hat Enterprise Linux certified and supported on NVIDIA DGX systems, the companies intend to collaborate on broader open source initiatives, including:

NVIDIA GPU Cloud (NGC) containers on Red Hat OpenShift - Red Hat and NVIDIA plan to deliver NGC containers that provide users with GPU-optimized software tools for AI and HPC based on Red Hat technologies, enabling mutual customers to take full advantage of NVIDIA GPUs.

Heterogeneous memory management (HMM) - Red Hat and NVIDIA plan to continue upstream development of the heterogeneous memory management (HMM) feature. This kernel feature allows devices to access and mirror the content of a system's memory into their own, enabling significant performance improvements for applications using GPUs.

Red Hat at GTC DC

To learn more about how Red Hat and NVIDIA align on open source solutions to fuel emerging workloads, visit Red Hat (booth #44) at the GPU Technology Conference (GTC) in Washington, D.C., from October 23-24, 2018. At the event, Red Hat and NVIDIA will also jointly present a "Best practices for deploying Red Hat platforms on DGX systems in datacenters" on Tuesday, October 23, 3:30PM - 4:20PM.

Supporting Quotes

*Chris Wright, CTO, Red Hat*

"The growing interest in performance-sensitive workloads, like AI and machine learning, requires a different approach to enterprise computing, a need that NVIDIA is already helping to address at the architectural level with NVIDIA DGX-1. With Red Hat Enterprise Linux and Red Hat OpenShift Container Platform, Red Hat is adding enterprise-ready software innovation to NVIDIA's powerful hardware, helping enterprises to drive emerging workloads while retaining the stability, reliability and familiarity that they have come to expect of their production systems."

*Charlie Boyle, senior director, DGX Systems, NVIDIA*

"NVIDIA DGX systems combined with NVIDIA-optimized AI software provide a powerful combination of performance, productivity and value to data science workflows from desk-side to datacenter. With the certification of Red Hat Enterprise Linux on DGX-1, we're enabling enterprise IT to meet the growing demand of customers who expect enterprise support at all layers of the software stack,

without inhibiting effortless AI exploration.”

*Bronis Supinski, chief technology officer, Livermore Computing; head of Lawrence Livermore National Laboratory's Advanced Technology systems*

“Red Hat Enterprise Linux’s enablement of NVIDIA GPUs on our Sierra supercomputer provides commonality across our systems, greatly facilitating our users’ ability to exploit the world’s third fastest computer. Since we use Red Hat Enterprise Linux as part of our software stack on our commodity technology systems, that GPU enablement reduces the effort required to move applications to Sierra, a capability that we anticipate will help benefit other GPU-based systems, such as DGX-1.”

*Charles Onstott, CTO, SAIC*

“Leveraging our Innovation Factory, and our key collaboration with Red Hat and NVIDIA, SAIC has demonstrated agility and collaboration with rapid testing and integration of elements of the stack deployed on GPU to optimize the performance of our analytic engine. Our time to design, deploy and demonstrate was reduced to weeks and will continue to iterate in a true DevOps fashion.”

#### Additional Resources

Read more about how Red Hat and NVIDIA are driving the future of GPU-accelerated workloads

Learn more about [Red Hat OpenShift Container Platform](#)

Learn more about [Red Hat Enterprise Linux](#)

#### Connect with Red Hat

Learn more about [Red Hat](#)

Get more news in the [Red Hat newsroom](#)

Read the [Red Hat blog](#)

Follow [Red Hat on Twitter](#)

Join [Red Hat on Facebook](#)

Watch [Red Hat videos on YouTube](#)

Join [Red Hat on Google+](#)

Follow [Red Hat on LinkedIn](#)

#### About Red Hat, Inc.

Red Hat is the world's leading provider of open source software solutions, using a community-powered approach to provide reliable and high-performing cloud, Linux, middleware, storage and virtualization technologies. Red Hat also offers award-winning support, training, and consulting services. As a connective hub in a global network of enterprises, partners, and open source communities, Red Hat helps create relevant, innovative technologies that liberate resources for growth and prepare customers for the future of IT. Learn more at <http://www.redhat.com>.

#### Forward-Looking Statements

Certain statements contained in this press release may constitute “forward-looking statements” within the meaning of the Private Securities Litigation Reform Act of 1995. Forward-looking statements provide current expectations of future events based on certain assumptions and include any statement that does not directly relate to any historical or current fact. Actual results may differ materially from those indicated by such forward-looking statements as a result of various important factors, including: risks related to the ability of the Company to compete effectively; the ability to deliver and stimulate demand for new products and technological innovations on a timely basis; delays or reductions in information technology spending; the integration of acquisitions and the ability to market successfully acquired technologies and products; risks related to errors or defects in our offerings and third-party products upon which our offerings depend; risks related to the security of our offerings and other data security vulnerabilities; fluctuations in exchange rates; changes in and a dependence on key personnel; the effects of industry consolidation; uncertainty and adverse results in litigation and related settlements; the inability to adequately protect Company intellectual property and the potential for infringement or breach of license claims of or relating to third party intellectual property; the ability to meet financial and operational challenges encountered in our international operations; and ineffective management of, and control over, the Company’s growth and international operations, as well as other factors contained in our most recent Quarterly Report on Form 10-Q (copies of which may be accessed through the Securities and Exchange Commission’s website at <http://www.sec.gov>), including those found therein under the captions “Risk Factors” and “Management’s Discussion and Analysis of Financial Condition and Results of Operations”. In addition to these factors, actual future performance, outcomes, and results may differ materially because of more general factors including (without limitation) general industry and market conditions and growth rates, economic and political conditions, governmental and public policy changes and the impact of natural disasters such as earthquakes and floods. The forward-looking statements included in this press release represent the Company’s views as of the date of this press release and these views could change. However, while the Company may elect to update these forward-looking statements at some point in the future, the Company specifically disclaims any obligation to do so. These forward-looking statements should not be relied upon as representing the Company’s views as of any date subsequent to the date of this press release.

*Red Hat, Red Hat Enterprise Linux, the Shadowman logo, and OpenShift are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the U.S. and other countries. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.*

View source version on businesswire.com: <https://www.businesswire.com/news/home/20181023005605/en/>

Red Hat, Inc.  
John Terrill  
[jterrill@redhat.com](mailto:jterrill@redhat.com)  
+1-570-772-3286

