Whitepaper

Sylabs.io

Singularity Community and SingularityPRO on HPE high-performance servers

The power of open source for enterprise performance computing

Hewlett Packard Enterprise

Table of Contents

Executive summary

Singularity has become an attractive container technology for running batch-style jobs because it was designed specifically to encapsulate reproducible application stacks into a single file. Its simplicity allows for seamless integration with GPUs and interconnects that are common to high-performance computing (HPC) environments. Running SingularityPRO™ on Hewlett Packard Enterprise (HPE) HPC server platforms expands upon Singularity Community's open source capabilities and includes commercial support and access to a growing value-added container ecosystem.

This white paper illustrates the business drivers for adopting container-based software models and the capabilities built into the SingularityPRO commercial offering. It will also address how Singularity container technology running on HPE platforms solve the challenges of parallelized AI, deep-learning/machine-learning, and data analytics workloads on large clusters—all without compromising security or privacy.

Three key takeaways from this whitepaper include:

- Typical use cases for Singularity Community/SingularityPRO running on HPE server platforms
- The unique value of SingularityPRO for today's enterprises
- The benefits of SingularityPRO compared to other container offerings

Introduction

Containers are a hot topic in every facet of high-performance computing. Applied use cases are seen in a variety of industries, including academia, finance, enterprise, and pharmaceuticals. 451 Research expects more than 250% growth in the container market from 2016 to 2020[1]. Containers combine speed and density with the mobility of traditional virtual machines (VMs) while requiring far fewer components to remain portable and run anywhere.

Containers are made possible by a set of facilities in the Linux® kernel that allow lightweight partitioning of a host operating system into isolated spaces where applications can safely run. Using containers presents lower overhead in terms of a smaller memory footprint and higher efficiency because they share the kernel with the host operating system—which means containers can achieve higher density. In short, containers enable more productivity.

Not only are containers orders of magnitude faster in provisioning, and lighter weight, they also enable applications to work in the same way on developers' workstations, on-premises servers, and any public or private cloud.

Proven open source container solution

Released in 2016, Singularity Community is an open source-based container platform designed for scientific and HPC environments. For HPC, Singularity makes what was previously impossible, possible.

With Singularity, the entire execution environment is contained within a single file that starts with a base Linux distribution, augmented by applications, libraries, data, and scripts—all for a containerized application workflow. Singularity containers easily integrate into standard HPC workflows and can be deployed and started on tens of thousands of nodes with minimal effort.

By moving away from the microservices architecture embraced by other container platforms, Singularity's unique design meets HPC users' needs for a container solution that not only offers high performance, but also supports mobility, reproducibility, and seamless integration with host-provided resources. In addition to enabling greater control over the application environments, Singularity also supports a bring-your-own-environment (BYOE) model—transporting a configuration from a scientist's workstation to the data center.

High-performance enterprise-class container platform

SingularityPRO builds on the success of the open source Singularity Community version, leveraging the open source code base to provide a container platform designed for Enterprise Performance Computing (EPC), including deep learning, IoT, and predictive analytics workloads.

SingularityPRO includes all of the functionality of the open source version, plus enterprise-grade enhancements that make the platform stronger, highly secure, and more feature-rich (described below). Where the open source version of Singularity is subject to rolling code changes from the open source community at large, SingularityPRO is curated and supported by Sylabs, the company behind Singularity.

Use cases

SingularityPRO running on HPE platforms (including HPE Superdome Flex, HPE Integrity Superdome X, HPE Integrity MC990 X, and HPE Apollo systems) delivers high-performance computing to enterprises. This is done by providing a secure and repeatable method to package applications and their dependencies into a single file that is cryptographically verifiable to ensure reproducibility. These features are critically important in the following enterprise use cases.

A Major milestone in Memory-Driven Computing

To help enterprises embrace the possibilities of a world transformed by exponential data growth, HPE offers Superdome Flex—the industry's only in-memory computing solution with a unique modular design that scales easily and economically for businesses of any size. A significant milestone in the Memory-Driven Computing innovation roadmap, this platform will help enterprises stay ahead of the competition by turning critical data into real-time business insights. Built to handle the most demanding applications, HPE Superdome Flex delivers an unprecedented combination of scale, modularity, flexibility, and reliability so that enterprises can turn these insights into action, and action into success—knowing that the business will remain always on.

Cluster Multi-tenancy

In an HPC environment, users are not allowed full, unrestricted administrative/root access to shared production systems[2]. Instead, users often receive credentials with limited access to reduce the threat surface areas. While limited-user credentials satisfy security, compliance, and audit requirements, users must be able to have enough environment privilege to develop, modify, and test their application containers.

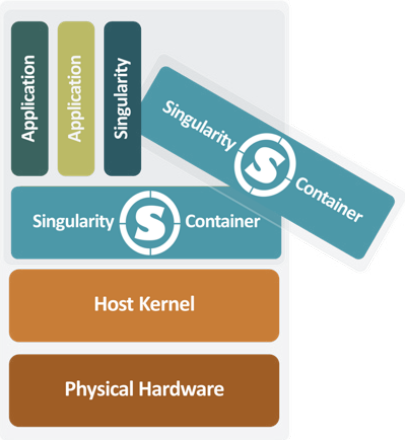

Figure 1: Singularity adds a new layer of isolation

Unlike other platforms, Singularity does not require a user to have root privileges within a container, and it does not require users to be added to a special group with advanced privileges to start the container runtime. Singularity's unique security model ensures that *untrusted users* can run *untrusted containers* without impacting the security of the underlying host system.

Enabling users to deploy Singularity containers on a cluster provides the flexibility they need, while also maintaining the security posture of the cluster.

Mobility of compute

Enterprise workloads are evolving. Jobs now consist of artificial intelligence (AI), machine learning (ML), and deep learning (DL) workloads that were solely within the domain of the scientific research community. Supporting the demanding EPC use cases found in today's life sciences, defense, financial technology, oil and gas, manufacturing, and many other types of workloads require a container platform that delivers high levels of performance, portability, and security.

Singularity running on HPE server platforms delivers such a platform—enabling users to create an application environment for running HPC workloads and applications without the performance penalties or complexities of accessing GPU and network interconnects. SingularityPRO simplifies the deployment of applications across different clusters and supercomputers (HPE Superdome Flex, HPE Integrity Superdome X, and HPE Integrity MC990 X systems) by avoiding the laborious process of re-hosting the applications for each distinct environment—without requiring a virtualized hardware layer. Singularity containers are just single files. If you can move a file from one host to another, you can deploy a Singularity container.

Mission-critical innovations

For enterprises running mission-critical applications on costly proprietary systems, HPE Integrity Superdome X sets new high standards for x86 availability, scalability, and performance. The ideal platform for critical Linux and Windows® workloads, HPE Superdome X blends x86 efficiencies with proven HPE mission-critical innovations for a superior uptime experience and groundbreaking performance. Breakthrough scalability of up to 16 sockets and 48 TB of memory handle in-memory databases and large scale-up x86 workloads. Through the unique HPE nPars technology, Superdome X adds agility and delivers 20x greater reliability than platforms relying on soft partitions alone. For maximizing application uptime, standardizing, or consolidating, HPE Integrity Superdome X helps transform today's mission-critical environments.

The Singularity Image Format (SIF) is a conduit for transporting entire application environments, as well as providing users and administrators with a means of protection. With Singularity single-file containers, users benefit from extreme mobility, enhanced reproducibility, and compliance control.

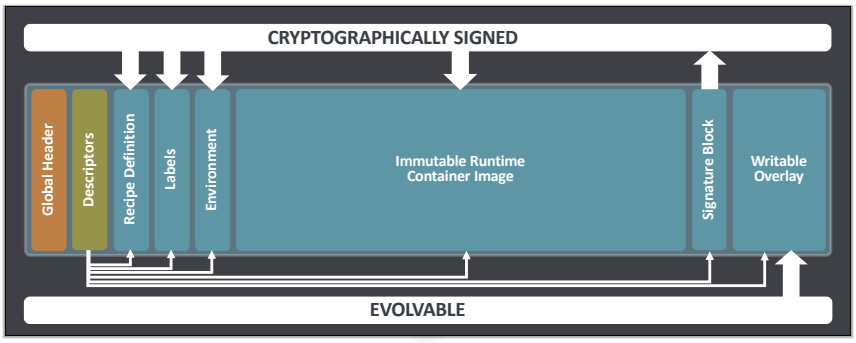

Figure 2: Singularity Image Format file structure and usage

SingularityPRO and associated Single Image Format (SIF) containers can have cryptographically signed and evolvable overlays to enable a controls-compliant workflow, which creates trusted containers. Unlike other container platforms, SingularityPRO has a mechanism to validate a runtime image and all data regions through a self-signing mechanism. By signing and verifying containers, distributors and users establish a level of trust unavailable to other container formats.

Drop-in replacement for standalone processes

Singularity integrates with all batch resource managers—with zero modifications—by calling the Singularity command directly.

One of Singularity's architecturally defined features is the ability to execute containers as if they were native programs or scripts on a host computer. As a result, integration with schedulers such as Univa Grid Engine, Torque, SLURM, SGE, and many others is as simple as running any other command. All standard input, output, errors, pipes, IPC, and other communication pathways used by locally running programs are synchronized with the applications running locally within the container.

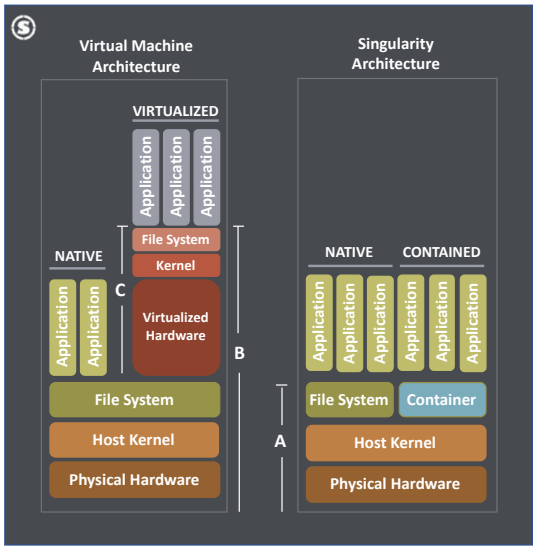

Figure 3: Positioning of Singularity in a Linux system

High-performance interconnects such as InfiniBand and Intel® Omni-Path Architecture (Intel OPA) are prevalent in the HPC/enterprise performance computing (EPC) space. Deep-learning workloads/applications also benefit from the high-bandwidth and low-latency characteristics of these interconnect technologies.

Singularity offers native support for OpenMPI by utilizing a hybrid MPI container approach, where OpenMPI exists both inside and outside the container. Similar to the support for InfiniBand and Intel OPA devices, Singularity natively supports any PCIe-attached device within the compute node, such as accelerators (GPUs).

Enabling the data-driven organization

The HPE Apollo Family is designed to deliver efficient rack-scale solutions for Big Data, analytics, object storage, and high-performance computing workloads. With rack-scale efficiency, the HPE Apollo Systems Family:

- *Delivers just the right amount of performance and efficiency with systems optimized for specific workloads*
- *Accelerates time to value by reducing implementation time*
- *Provides architectural flexibility with both scale-up and scale-out solutions*
- *Helps reduce capital and operating expenditures (CAPEX and OPEX)*

Architectural differences between Singularity and other container platforms

Security is a common concern for enterprises considering the adoption of containers in a shared computing environment. This is due in large part to other container platforms requiring elevated privilege daemons or configurations where the locking capabilities are limited and challenging to implement[3]. Another fundamental difference between Singularity and other containers is the image format itself. A Singularity container is a single file that can be moved around, the same as any other file. Other container runtimes contain layers, which are assembled during runtime and do not offer the same mobility and reproducibility as a Singularity container.

And finally, unlike other container platforms, Singularity favors integration over isolation, allowing it to work with common HPC technologies such as high-speed interconnects, batch schedulers, resource managers, MPIs, and GPUs with little or no additional configuration.

SingularityPRO on HPE infrastructure, ready for the enterprise

SingularityPRO is a certified binary release of Singularity built entirely from the open source code base—augmented with the licensing, support, and expert professional services requested by leading organizations, universities, and laboratories.

Unparalleled scale for data-intensive workloads

HPE Integrity MC990 X Server delivers in-memory computing performance for Linux-based applications at an unparalleled scale with mission-critical reliability and modular flexibility. An advanced symmetric multiprocessing (SMP) system designed for data-intensive workloads, the HPE MC990 X Server features enterprise-class Intel Xeon® E7-8800/4800 v4 processors and robust reliability, availability, and serviceability. The 5U modular chassis contains 4 sockets with up to 192 threads. By adding chassis and leveraging high-bandwidth NUMAlink technology, the HPE MC990 X Server can scale as a single system from 4 to 32 sockets and from 1 to 48 TB of cache-coherent shared memory.

Stronger platform, better support

While many components of an enterprise computing environment (local or cloud) consist of essential open software components, administration and support of the software need to come from somewhere. In short, "free" software is not really free.

Building on the success of Singularity Community—an open source container development platform used by over 25,000 top academic, government and enterprise users, that's installed on over 3 million cores and running over a million jobs per day—SingularityPRO includes numerous enterprise-grade support features:

- Long-term support, where security patches and bug-fixes are backported into SingularityPRO versions. This way, administrators are released from the burden of continually updating the Singularity code base to the latest open-source version.
- Early releases of security patches, delivered to SingularityPRO customers before propagation into the source community release.
- Stability, by providing long-term support, along with bug and security fixes.
- Customized service/support options, enabling SingularityPRO users to choose the tiered service/support option that best meets their needs.
- Access to a vast ecosystem of resources, including a container Remote Builder, Container Library, and Key-signing service (described on right).

Features	Singularity Community	SingularityPRO
SIF: Single File Container Format	■	■
Cryptographically Verifiable	■	■
No Persistent Global Daemon Process	■	■
Support for Non-root Users	■	■
Running Containers	■	■
Blocking Privilege Escalation within a Container	■	■
"Bring Your Own Environment" Usage Model	■	■
Support for AI/HPC Workflows and Architectures	■	■
Support for GPUs Natively	■	■
Code Curation		■
Streamlined Security Updates		■
Sylabs Cloud Features		■
Signed Packages and Repositories		■
Additional Self-Service Help		■
Container Build Services		■
Cryptographic Key Service		■
Container Library		■

Figure 4: Subscription provides access to SingularityPRO and a vast ecosystem of services. Compare features and choose the right version for your organization.

SingularityPRO add-on services

In 2018, Sylabs is making available multiple value-add services for SingularityPRO. Access to these services will be available for demonstration purposes to open source Singularity Community users. The services will also be offered under various tiers (trial, SMB plan, and Enterprise plan) to SingularityPRO customers.

Remote Builder

Building a container requires elevated privilege. In many HPC and EPC environments, however, elevated privileges are not possible because:

- Regular users cannot have administrative access to any cluster resource
- Using an external workstation to build a container breaks the chain of trust

The Remote Builder addresses these challenges by moving the build process to a secure, controlled environment.

During the build process, output streams back to the requester, so the user can monitor the build's progress. Upon completion, the SIF image is transferred back to the user's workstation, from which point it can be executed with Singularity, or sent to the Container Library—with no elevated administrative privileges required. In addition, no workflow modifications are necessary. Adding a single flag enables the build to be completed remotely without elevated privilege.

The Remote Builder implements appropriate levels of isolation between the components performing the builds with elevated privileges, isolating them from a shared cluster. System administrators receive a turnkey solution that empowers users to build Singularity images, as well as provides a centralized auditing and monitoring console for Singularity builds. These services are available in the cloud and on-premises.

Container Library

The Container Library was created and designed for hosting SingularityPRO containers. The full-featured Library can be hosted on-premises in your data center or the Sylabs cloud. Users can upload, download, search, and browse for containers in public and private areas, as well as share private containers with other users or via a generated link. Security and privacy in the Container Library are based on a user-owner of library objects, and the concept of public or private collections.

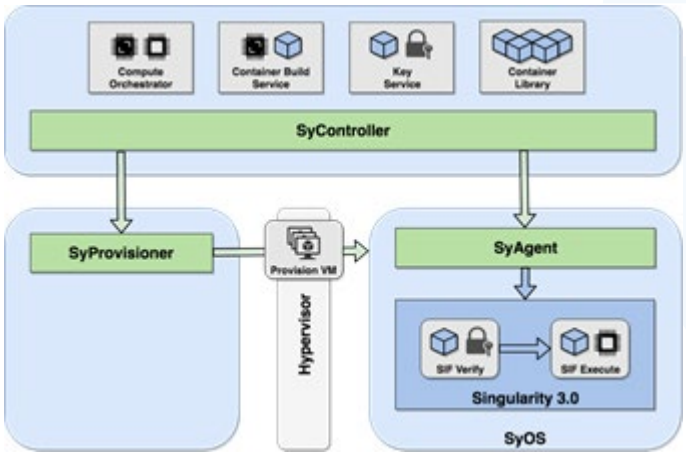

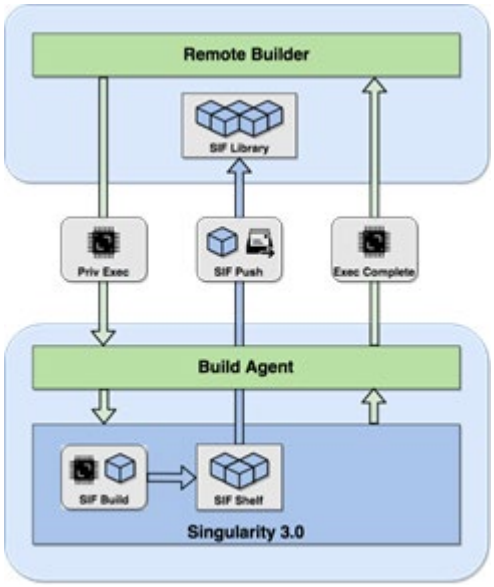

Key-signing and verification services

With Singularity 3.0, the new Singularity Image Format (SIF) will deliver container signing and validation services to Singularity and the Container Library. These key-signing and verification services eliminate the risk of unknowingly downloading and running compromised or rogue containers.

The ability to quickly identify containers signed by trusted sources—both internal and external—enhances an organization's auditing capabilities and its ability to enforce policies for restricting the types of containers allowed to run on a cluster.

Conclusions

Containers promise to seamlessly move applications between environments—from development to QA to a 10,000-node cluster. Containers ensure that each application will run the same way and will produce the same result in any environment—only faster.

Singularity running on HPE HPC platforms simplifies the process of moving containers across a single infrastructure or across hybrid environments. This solution also preserves privilege separation to satisfy the security, privacy, and auditing requirements found in all supercomputer and enterprise environments.

Raising the bar for container platforms, SingularityPRO running on HPE HPC platforms leverages the power of AI, machine learning, and deep learning to deliver unique enterprise-level services. SingularityPRO's advanced ecosystem of resources not only extends the overall value of the platform but also extends its ease of use and security.

Hewlett Packard Enterprise

[1] https://451research.com/images/Marketing/press_releases/Application-container-market-will-reach-2-7bn-in-2020_final_graphic

[2] Even though users have limited access to production systems, they can have full administrative access to their own development virtual machine.

[3] Docker daemon attack surface, https://docs.docker.com/engine/security/security/#docker-daemon-attack-surface

sales@sylabs.io
www.sylabs.io/contact
www.sylabs.io

Sylabs.io