

Spectrum Scale Cloud Native / Containers



Disclaimer

This information is provided on an "AS IS" basis without warranty of any kind, express or implied, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Some jurisdictions do not allow disclaimers of express or implied warranties in certain transactions; therefore, this statement may not apply to you.

IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion. Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code, or functionality. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here.

IBM reserves the right to change product specifications and offerings at any time without notice. This publication could include technical inaccuracies or typographical errors. References herein to IBM products and services do not imply that IBM intends to make them available in all countries.

Goal: Deliver High Performance File Services to Containerized Application Workloads

Support Workloads that Require High Performance File Services

- Analytics & Cognitive
- High Performance Computing
- AI Data Pipeline

Support the Workload Ecosystem in the Cloud

- Containerized Applications, Storage
- Ephemeral and Persistent Storage Volumes

Flexible Deployment

- Dynamic Provisioning, Configuration, Upgrade

Support for Multiple Clouds

- Public, Private, Hybrid

Support Hybrid Use Cases

- Cloud Burst – Single Name Space
- Multi Cloud Data Sharing
- Archive
- High Performance Tiering

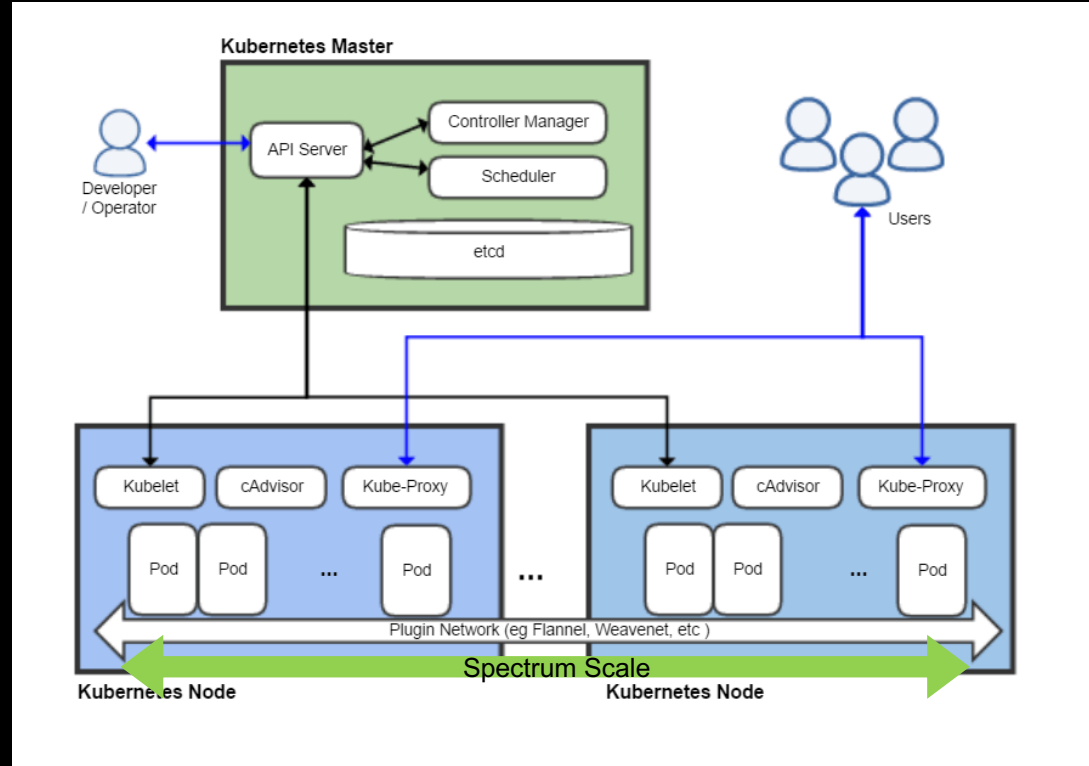
Solution Integration (Partners)

Why Containerized Spectrum Scale?



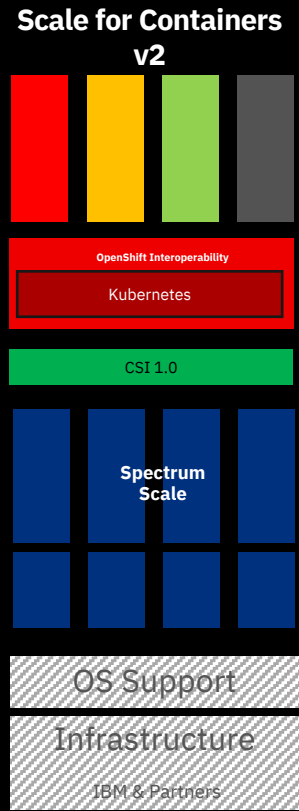


Kubernetes provides a scale-out platform for running application containers.

Spectrum Scale is an ideal fit for this scenario, because it provides all data to all worker nodes directly. Containers can be freely moved between nodes without having to remap/remount volumes on the worker node.

Spectrum Scale scales far beyond the limits of a Kubernetes cluster and offers unmatched performance.



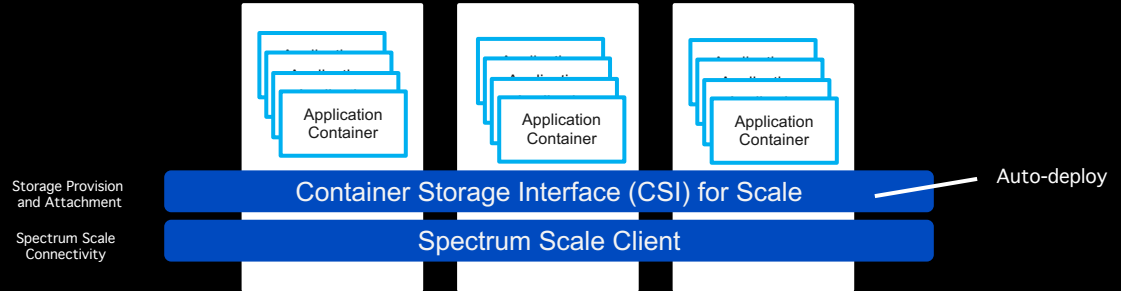
Evolution of IBM Spectrum Scale Containers



Spectrum Scale Containers Models

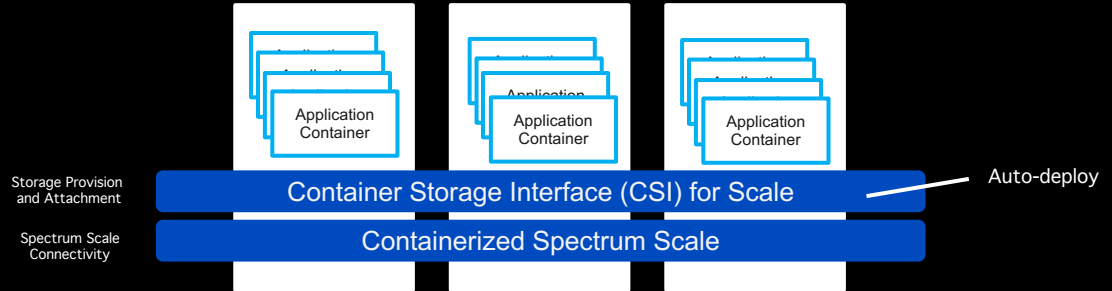
Storage for Containers

- Container Ready Storage
- No support for RHCOS
- Traditional deployment & management



Storage in Containers

- Containerized Storage
- Supports RHCOS
- Deployment & management via operators

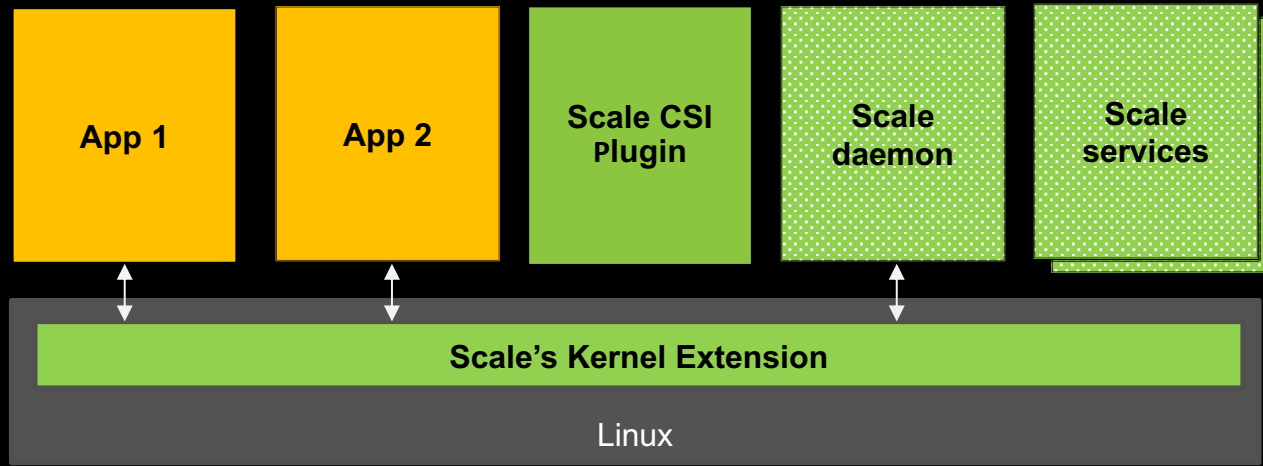


Containerized Spectrum Scale

The core of containerized Scale is a daemonset of containers that run on every worker node, load the kernel extension and start the filesystem daemon. This set of containers needs special privileges to load kernel modules, mount file systems, access block devices and access the host network.

Other Scale services like ReST API, GUI, performance data collection, etc. run in separate containers.

Deployment and operation is implemented via operators.



Demo

