

Scientific Computing with Application Containers: Onedata and Hyperflow Use Cases

Michał Orzechowski^{1,2}, **Bartosz Baliś**¹, **Michał Cwiertnia**², Łukasz Dutka², Renata G. Słota¹, Jacek Kitowski^{1,2} ¹AGH University of Science and Technology, Department of Computer Science, Krakow, Poland ²AGH University of Science and Technology, ACK Cyfronet AGH, Krakow, Poland



Motivation

- Scientific computing management remains challenging
 - Resource management, data management, software deployment, ...
- Virtualization, **application containers**, and infrastructure automation can mitigate deployment and execution complexity
- However, in scientific computing app containers are not used to their full potential
 - Mainly as a means to provide user-defined software



Objectives

- Discuss benefits of application containers in scientific computing
- Demonstrate on two case studies:
 - Onedata: scientific data management
 - HyperFlow: scientific workflow management
- How to improve support for application containers in computing centers?



Containers and container orchestration



An ocean of user containers

Container management software

Kubernetes, Mesos, Docker Swarm **Hosted container management platforms** Amazon ECS/EKS, Google Kubernetes Engine **Serverless container management** Amazon Fargate, Google Knative

Scheduled and packed dynamically onto nodes



Benefits of application containers for (scientific) computing

Development

- Consistent runtime environment (easier testing and debugging)
- Component re-use and built-in versioning

Deployment

- Sandboxing and portable deployment across machines and computing infrastructures
- Runtime environment in the container rather than host machine

• Execution

- Better resource utilization
- Isolation
- Ease to achieve high scalability of scientific application components
- Lightweight virtualization and lower runtime overhead compared to VMs





SCIENTIFIC WORKFLOW MANAGEMENT SYSTEM



Scientific workflow management





Scientific workflow management in HyperFlow





HyperFlow WMS





ONEJATA

ONEDATA EVENTUALLY CONSISTENT VIRTUAL FILESYSTEM FOR MULTI-CLOUD INFRASTRUCTURES CYFRONET AGH









WHO WE ARE?

- Group of developers bringing hybrid cloud open source platform to life
- 5+ years devoted development
- Our main goal is:

CYFRONET

- to deliver data management platform for large scale and distributed problems
- to make the solution decentralized and eventually consistent in order build a mesh of data sources
- to deliver virtual file system for hybrid cloud
- The work is supported by:

T··Svstems···









CAPABILITIES OF ONEDATA

1	Multi-protocol transparent access to data "[] but we want POSIX"
2	Heterogeneity of storage technologies
3	Block data access
4	Easy Data Sharing and publication (DIO)
5	Metadata Management Integrated with Data Management Platform
6	Flexible authentication and authorization
7	Easy integration using API with external services
8	High-throughput data processing



MULTI-PROTOCOL TRANSPARENT ACCESS [...] BUT WE WANT POSIX

- Transparently access and create data in multi-cloud environments
- Care less about data locality, all your • data are accessible wherever you go
- Support for most of the POSIX operations on globally distributed virtual file system
- All data accessible via a unified file system mountable on virtual machines, Grid worker nodes and containers

כ	ONEDATA		total 0 drwxrwx 1 root 1733762 0 Sep 26 19:19 Astronomy Datasets drwxrwx 1 root 1337123 0 Sep 26 19:14 Big Data Experiment drwxrwx 1 root 608582 0 Sep 26 19:18 Cancer Data
]	🛖 astronomy 🗸	FILES	<pre>./Astronomy Datasets: total 0 drwxr_xr_x 1 1124656 1733762 0 Sep 26 19:20 comets</pre>
2	Root directory	📮 jupiter.dat	drwxr-xr-x 1 1124656 1733762 0 Sep 26 19:19 planets
) :es	-	pluto.dat	<pre>./Astronomy Datasets/comets: total 0</pre>
	comets	Venus.dat	-rw-rr 1 1124656 1733762 10000000 Sep 26 19:20 enck.dat -rw-rr 1 1124656 1733762 10000000 Sep 26 19:19 halley.dat
ıps	planets		<pre>./Astronomy Datasets/planets: total 0</pre>
3			-rw-rr 1 1124656 1733762 10000000 Sep 26 19:07 jupiter.dat
ens			-rw-rr 1 1124656 1733762 5000000 Sep 26 19:08 pluto.dat -rw-rr 1 1124656 1733762 2000000 Sep 26 19:08 venus.dat
			<pre>./Big Data Experiment: total 0</pre>
ders			-rw-rr 1 1124656 1337123 10000000 Sep 26 19:08 cats_images.tg
			-rw-rr 1 1124656 1337123 5000000 Sep 26 19:13 galaxies.img
			-rw-rr- 1 1124656 1337123 5000000 Sep 26 19:14 spam_mails.tgz
			./Cancer Data: total 0
			-rw-rr 1 1124656 608582 5000000 Sep 26 19:15 brain_tumor.zip
			-rw-rr 1 1124656 608582 5000000 Sep 26 19:14 duct_cancer.zip
			[root@1f87c053280e oneclient]# 16

[root@1f87c053280e oneclient]# ls

[root@1f87c053280e oneclient]# ls -lR

Astronomy Datasets Big Data Experiment Cancer Data



BLOCK DATA ACCESS

- File distribution between storage locations is underneath the file structure
- Files management on a chunk basis
- Missing chunks delivered on the fly
- API for file management for pre-staging and implementing external data policy management

0)				0			
Data	🛖 S1	File distribution		for file file		SIZ	e Modif
\bigcirc	🔛 Root di	Provider	File blocks	for the the		106	B 2016-0
Spaces		p1	o		10	06 B	
Groups		p2	0		10	06 B	
a23e						_	
Tokens			Close				
Providers						_	



HIGH-THROUGHPUT PROCESSING





CASE STUDIES



User machine Infrastructure description files Infrastructure configuration files Workflow graph and input data Terraform Amazon ECS Scaling rules CloudWatch Security policies AWS resources

Case study: HyperFlow @ Amazon ECS cloud with auto-scaling







How to improve support for containers in computing centers?

- Automated deployment of k8s cluster on Cyfronet Openstack
 Currently done manually it is HARD!
- Better integration of k8s with Openstack:
 - automated storage provisioning with k8s
 - automated Openstack Public IPs assigned to k8s services
 - k8s cluster auto scaling

Unified integration of HPC resources (Prometheus) with k8s: submitting jobs to k8s that actually run on HPC .



Conclusion

Application containers offer a LOT for scientific computing All the benefits of virtualization, but much more lightweight

Onedata and HyperFlow developed using Cyfronet cloud resources – Openstack and PLCloud:

- bamboo cluster for continuous integration
- K8s cluster for large deployments and performance testing

To achieve their full potential, we need container orchestrators Better support in computing centers is desirable





Please visit: <u>www.onedata.org</u> <u>https://github.org/hyperflow-wms</u>