



Singularity Containers

Containers for ...
HPC, EPC, Compute Driven Analytics, Machine Learning, Reproducible Research, and more...

<https://sylabs.io>

SPEAKER INFORMATION



Michael Bauer

Software Engineer, RStor
June 2017 – Jan 2018

Software Engineer, SyLabs
Jan 2018 – ??

bauerm@sylabs.io
@bauerm97 on GitHub

WHAT MAKES SINGULARITY SO SPECIAL



- **Built specifically to support HPC/Science:** Singularity was built by demand, requests, threats and bribes by researchers, scientists, and computational users
- **Single file based container format:** verifiable via checksum and cryptographic signatures enabling reproducible and validated software environments during runtime
- **Extreme mobility:** using standard tools (rsync, scp, GridFTP, NFS, etc.)
- **Controls compliant:** images can be easily archived and managed as any other data
- **Compatible:** with complicated architectures (e.g. HPC, Machine Learning, Cloud, etc.)
- **Security model:** designed to support untrusted users running untrusted containers
(rather than trusted users running trusted containers)



Reproducibility and Portability



Without having access to an identical source, environments must be recreated.


Documentation, recipes, source code and data might be able to get pretty close, but is it identical?

Is it close enough?

IS THIS “GOOD ENOUGH”?

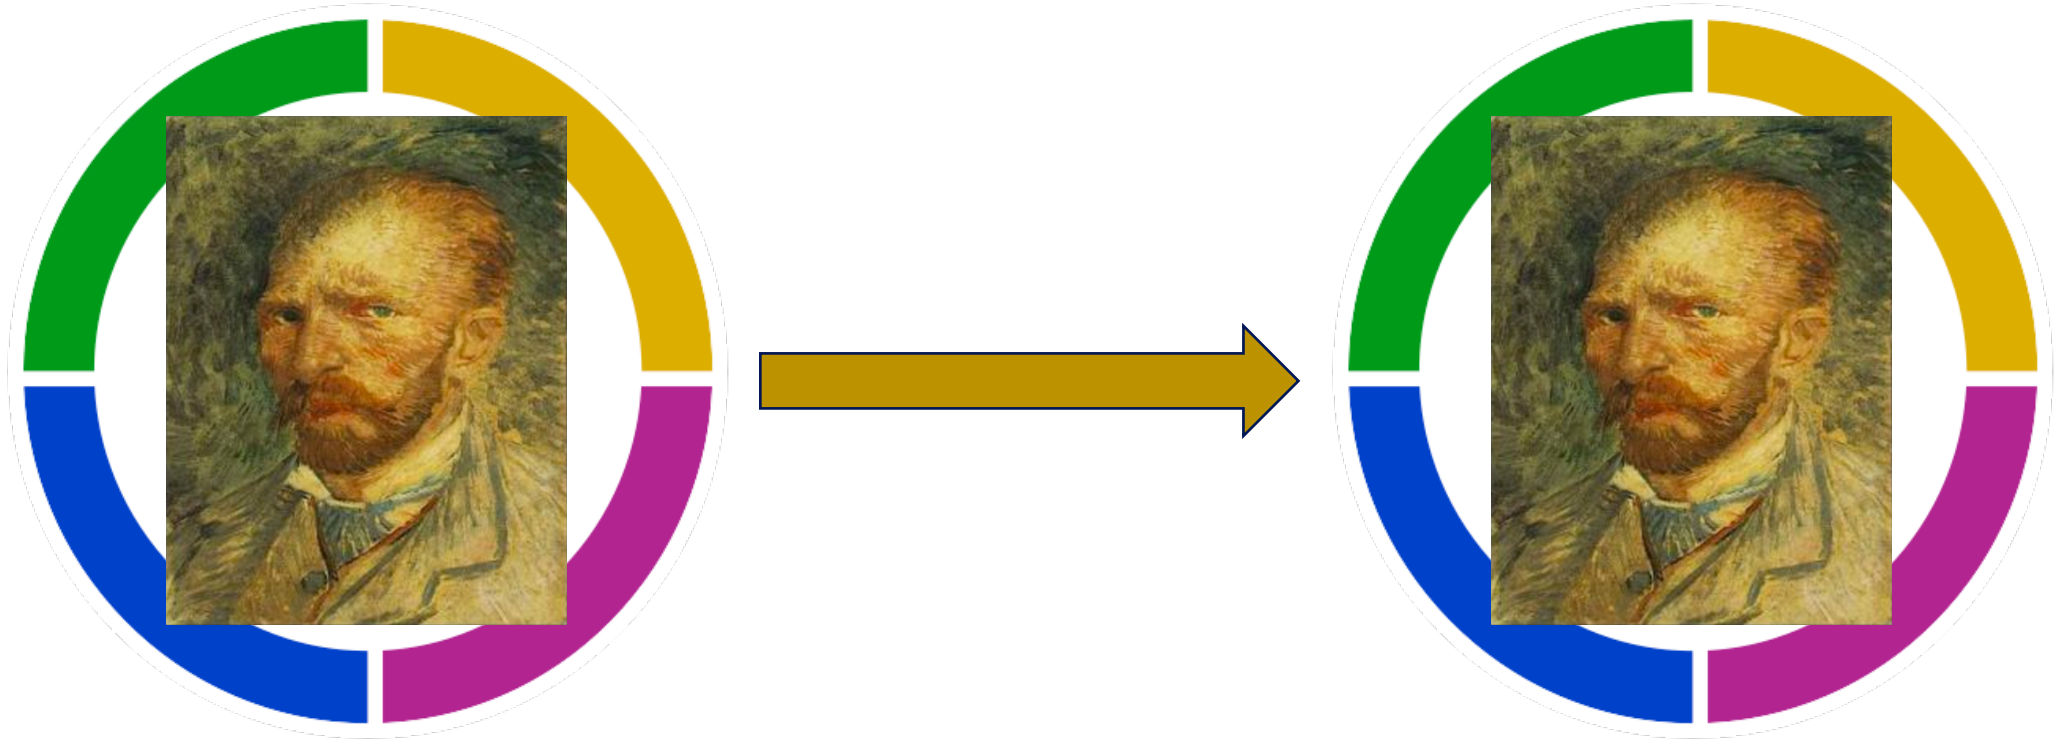


= ?



WHAT IS YOUR REPRODUCIBILITY PLAN?

BIT FOR BIT SOFTWARE REPRODUCIBILITY



SHA:

5f09a35a642a68c467bf230f5e5ea3218e4177a0


SHA:

5f09a35a642a68c467bf230f5e5ea3218e4177a0

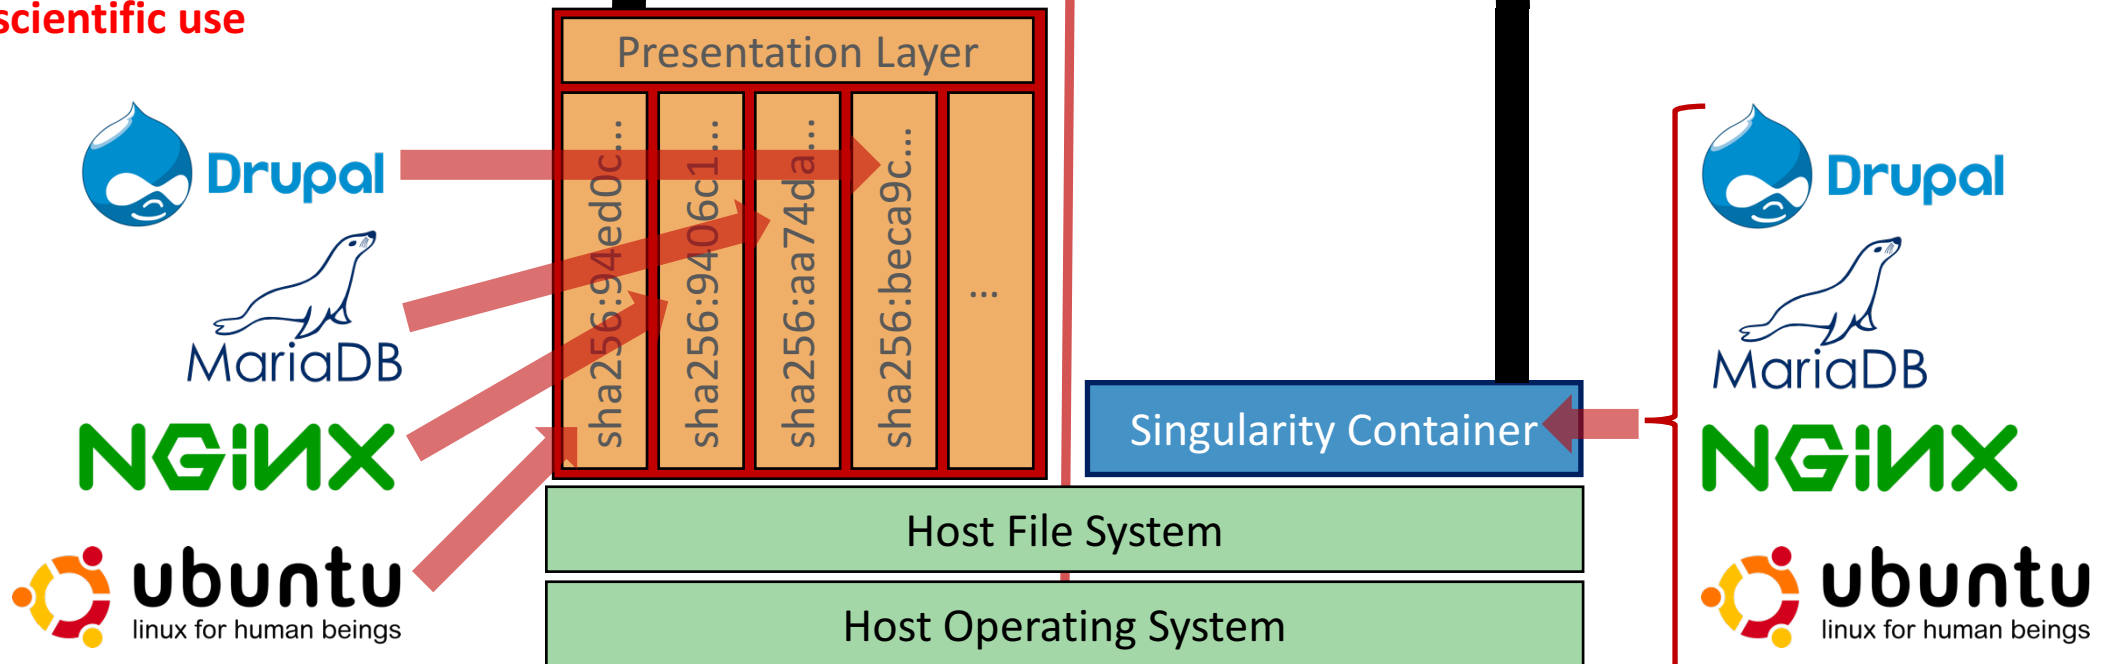
COMPARISON: MOBILITY, VALIDATION AND ARCHIVAL



One must **squash layers**, breaking upstream validation, and/or restore to a compatible registry service to reassemble and use the layers. **Not ideal for many scientific use cases!**



Container is a single file, easy to deal with and use.



ENTERPRISE



SINGULARITY
ENABLES
EXTREME
MOBILITY

SCP, SFTP,
GridFTP/Globus,
Rsync, NFS,
Lustre, Object
Stores, etc..



Single file
containers means
that containers are
easy to manage



COMPUTE



SINGULARITY INVOCATION PERFORMANCE OVER SHARED STORAGE

Objectives:


1. Measure scaling of python startup and import speed with increasing numbers of concurrent python interpreters
2. Compare scaling of a standard python installation with an identical containerized installation ([Singularity](#)).

*Note: Underlying file system is NFS, max jobs was 5120 over 320 nodes, graph is **logarithmic** on both axis.*



DR. WOLFGANG RESCH

[HTTPS://GITHUB.COM/WRESCH/PYTHON_IMPORT_PROBLEM](https://github.com/wresch/python_import_problem)



SINGULARITY 3.0



- Target release date: early summer 2018
- Migration to Go
- Support for custom runtime plugins
- SIF (Singularity Image Format)
- Official native OCI compliance

MIGRATION TO GO



- **Current CLI:**
 - Written in shell scripts and some Python
 - Hard to maintain, debug, expand
- **Solution – Move CLI into Golang:**
 - Modern compiled language for systems programming
 - Standard for other container implementations
 - Rich ecosystem of libraries
- **Proposed Plugin System**

CVMFS WITHIN CONTAINER



- **Issue from yesterday** – Need to access different CVMFS snapshot on a per job bases. Potentially multiple per node
- **Solution** – Mount CVMFS in userspace within container
- **How to do this?** – FUSE needs SUID but Singularity blocks all privilege execution pathways within container

CVMFS WITHIN CONTAINER - SOLUTION

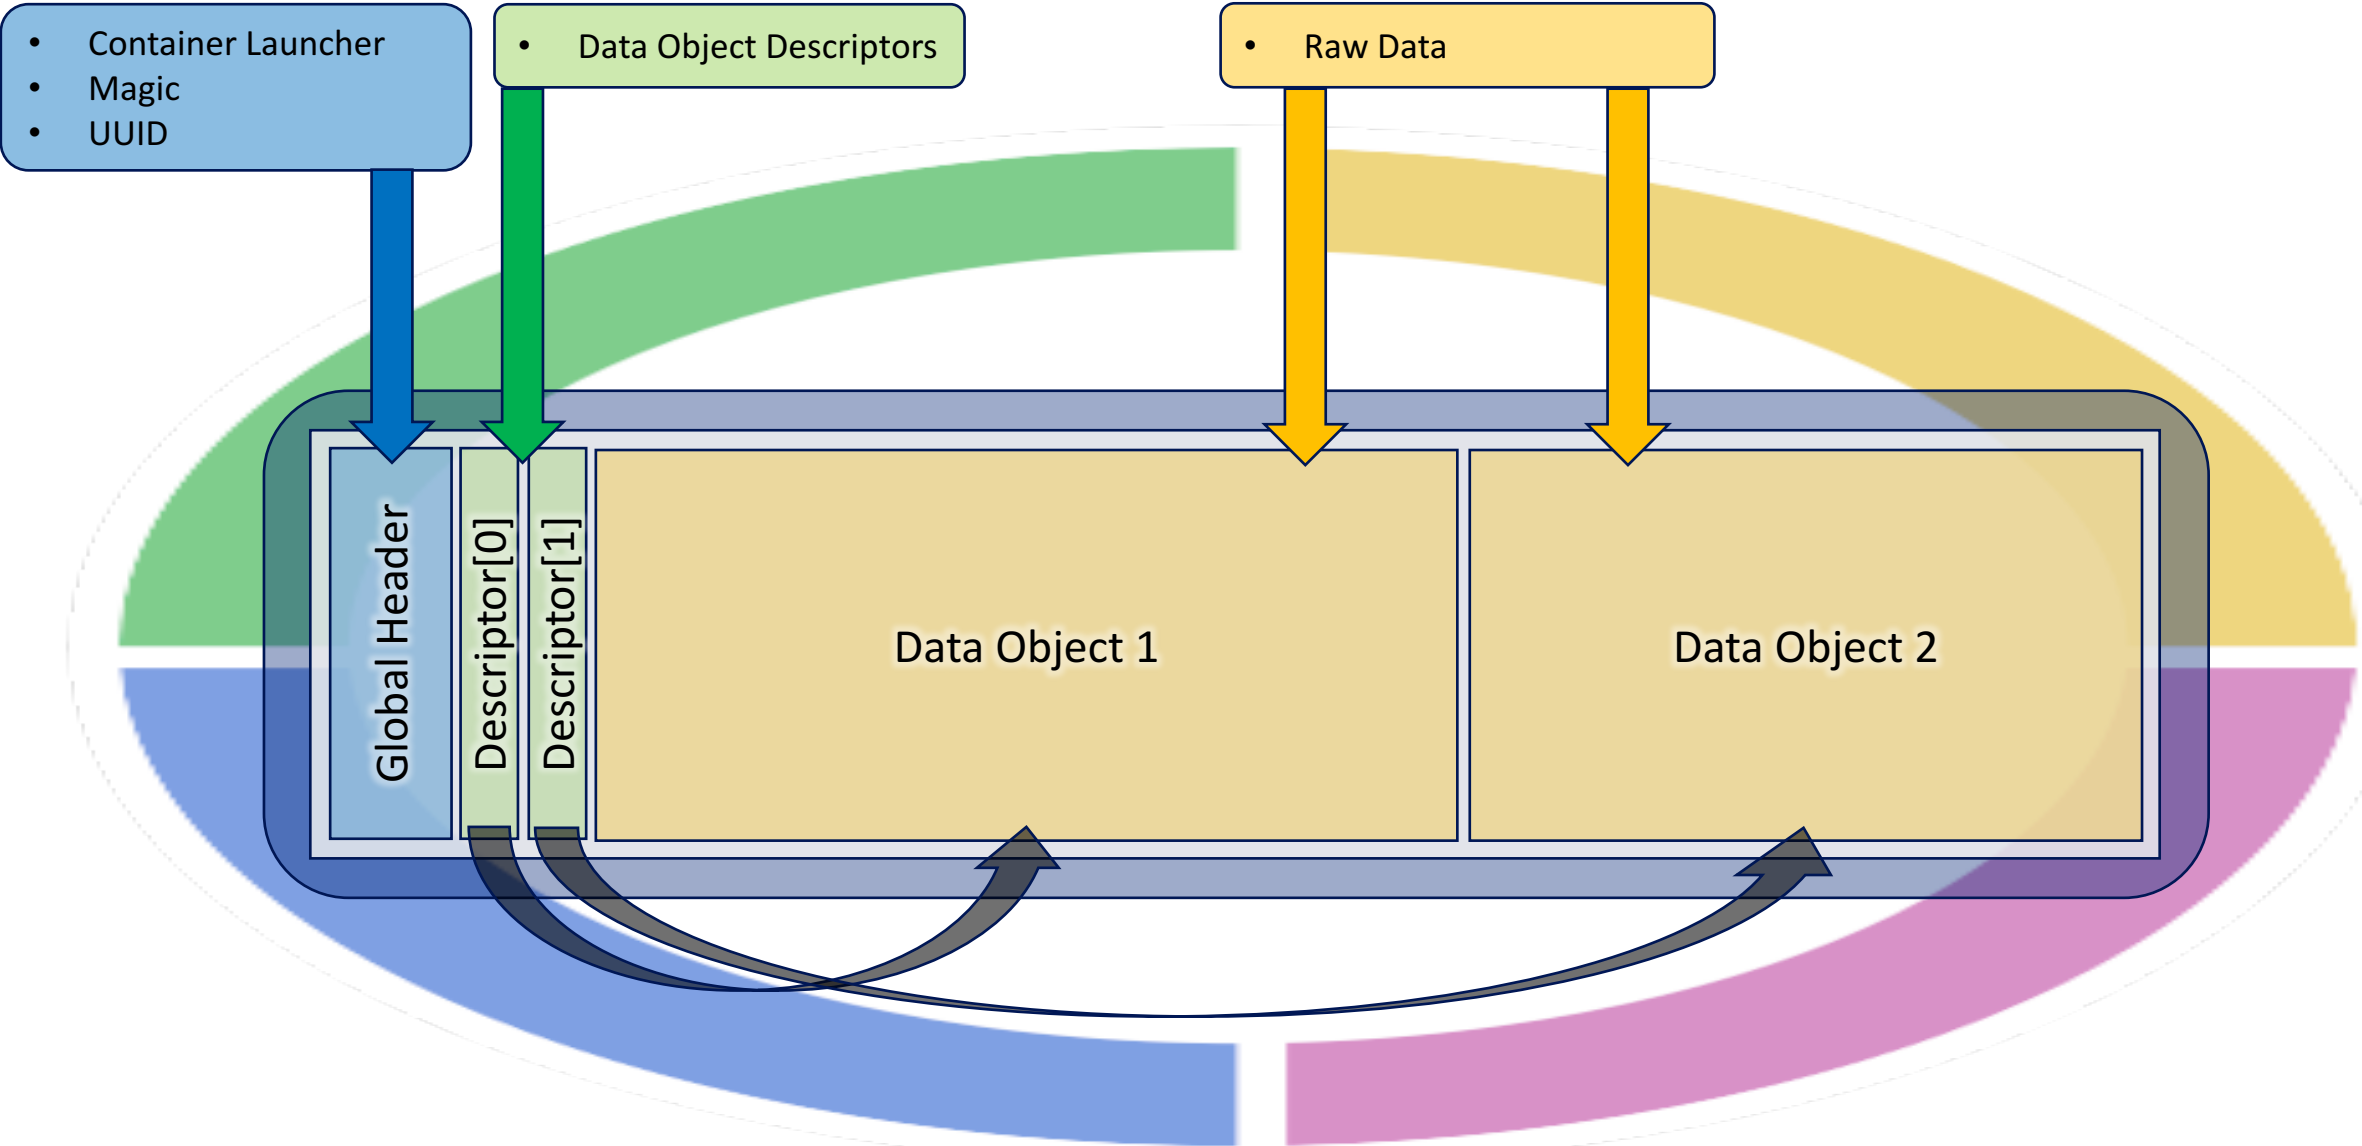


- **Desired result:**
 - Singularity run `--cvmfs [config] ...`
- **Solution Proposed** – Singularity runtime plugin support
 - CVMFS plugin for Singularity
 - Leverage existing SUID from Singularity binary
 - Mount CVMFS in user space before contained proc runs



Supported Image Formats – 3.0

SINGULARITY IMAGE FORMAT (SIF)





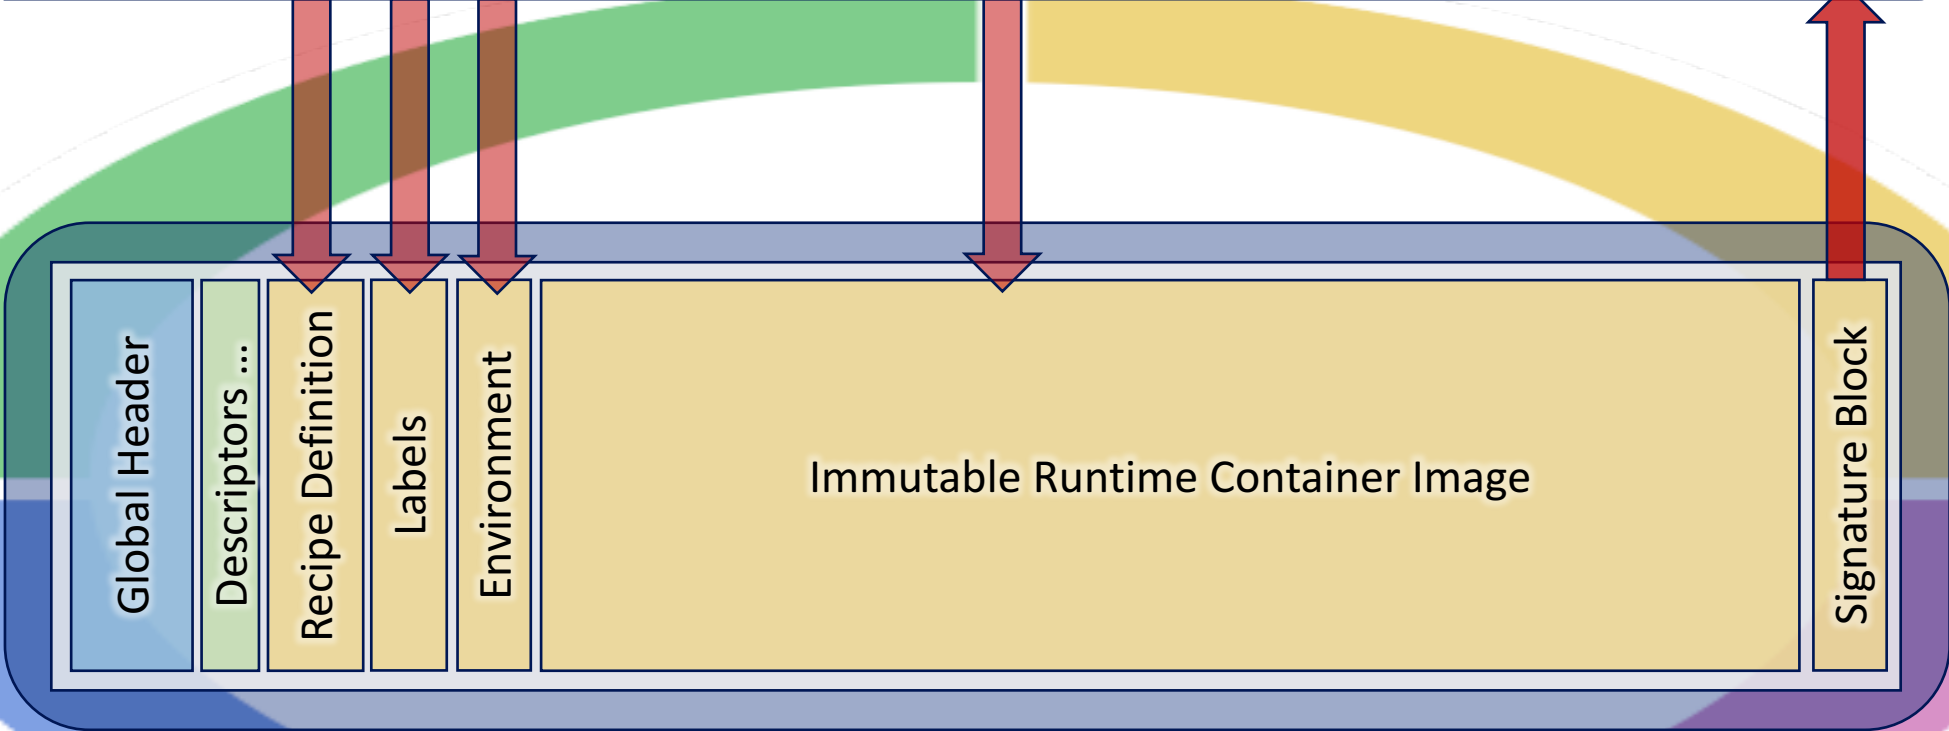
SINGULARITY IMAGE FILE (SIF): FEATURE LIST

- **Cryptographic signatures and verification:** The Singularity image format supports both region checksums as well as signing for your verification pleasures
- **Faster access to container meta-data:** Container meta-data is now part of the image file but outside the container
- **Support for multiple system partitions:** A single container image can *contain* multiple container regions and/or a writable overlay
- **Support for Checkpoint Restart:** Internal support for checkpoint-restarting for mobility of state

SIF EXAMPLE: TRUSTED, SIGNED CONTAINERS



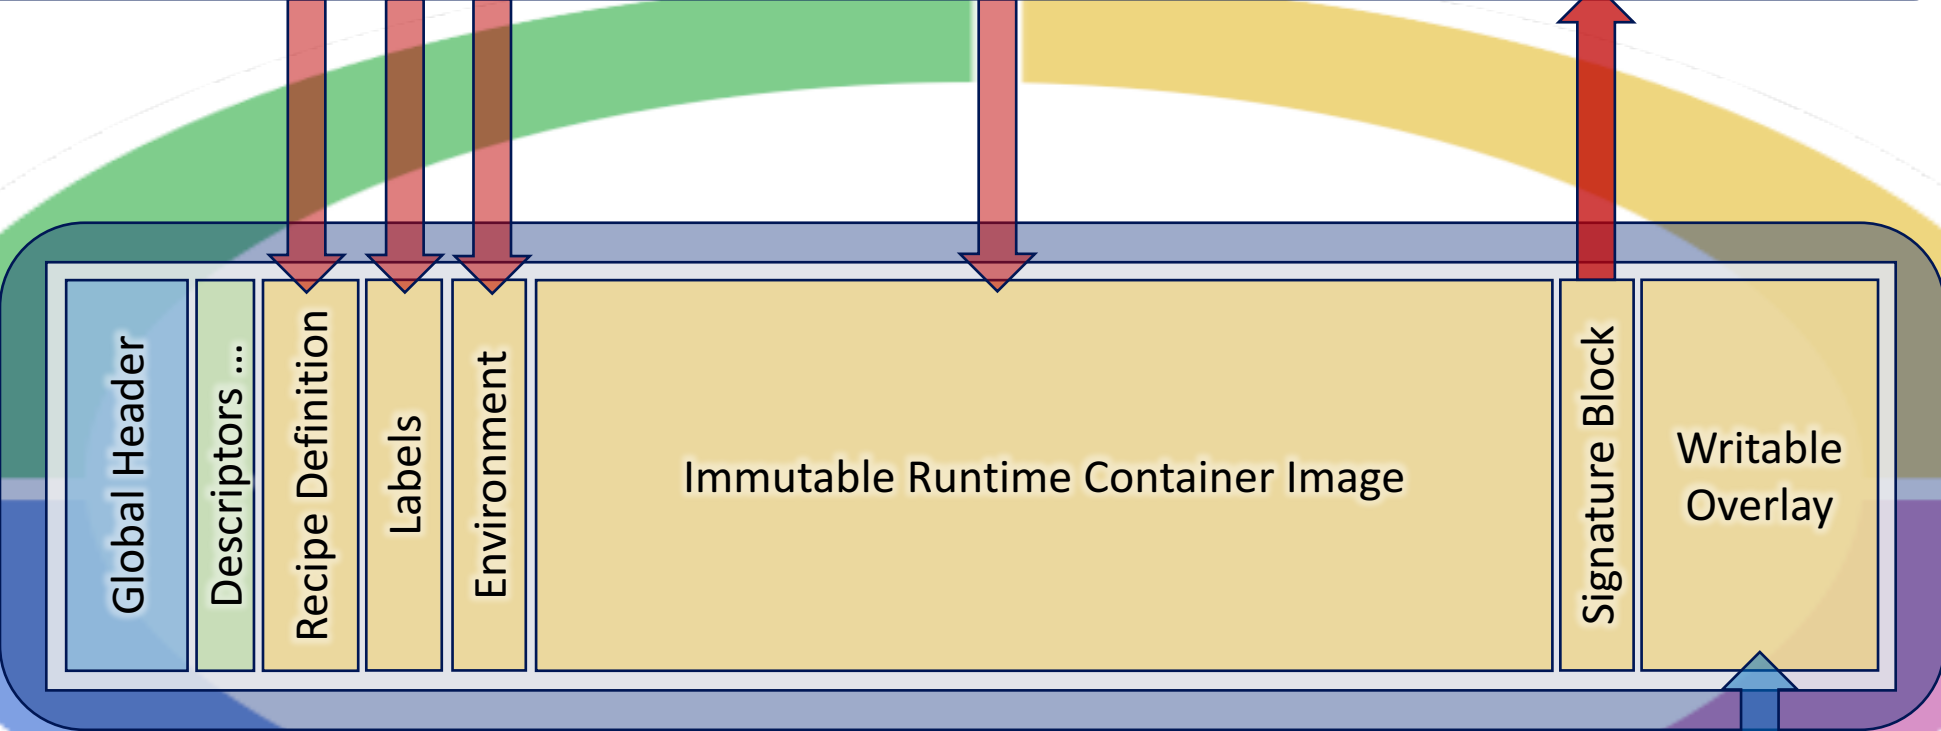
Cryptographically Signed



SIF EXAMPLE: TRUSTED, EVOLVING CONTAINERS



Cryptographically Signed



Modifyable
Coming soon!



OTHER SUPPORTED FORMATS



- **OCI Bundle** – Combination of filesystem on disk + OCI configuration
 - bundle/
 - config.json
 - root/
- **Sandbox** – Filesystem on disk
 - root/
- **Remotely from:**
 - Singularity Hub
 - Docker

OCI COMPLIANCE



- **runtime-spec:**
 - Specify container runtime options as json (config.json in bundle or as option)
 - Support 5 Operations: Create, Start, Kill, Delete Query
- **image-spec:**
 - Bundle described in prev slide
 - Manifest describing how to build container filesystem
- **Native K8s support through CRI-O**



Use Cases

nextflow



- Nextflow is a workflow management language for data-driven computational pipelines
- Nextflow uses Singularity to deploy large-scale distributed scientific workflows
- Commonly used in genomics pipelines
- Supports both HPC cluster and cloud based resources in a portable manner
- Used by:
 - Center for Genomic Regulation (CRG)
 - Pasteur Institute (France)
 - SciLifeLab (Sweden)
 - Sanger Institute (UK)



- The NIH uses Singularity to provide programs like TensorFlow and OpenCV3 which are difficult or impossible to run with their current operating system
- With Singularity they can create "portable reproducible data analysis pipelines"
- Singularity allows the NIH to provide this functionality to their users in a secure environment
- The system-admins found it easy and intuitive to use Singularity
- Some applications have been installed into Singularity containers and used as standalone programs via environment modules for the users (the users don't even know they are running within a container!)



- ALICE jobs are packaged into Singularity containers
- Jobs are executed via Singularity through a modified SLURM script
- At any given moment in time, there are about 2000 Singularity containers active on the system




GSI Green Cube
6 stories tall
30,000 sqft
12 MegaWatts
PUE = 1.07 (world record)



Open Science Grid

- The OSG uses Singularity to provide a consistent runtime environment across heterogeneous resources worldwide
- Container images are distributed via CVMFS to all sites
- About half a million jobs are run through Singularity per day

Instances/Day





OAK RIDGE

National Laboratory

Titan @ Oak Ridge Leadership Computing Facility



Adam Simpson (front) and Matt Belhorn (back), high-performance computing user support specialists at the OLCF, use the Singularity application to develop containers that will allow newer systems to run deep learning packages.



SyLabs, Inc. – Funding Singularity Development

<http://sylabs.io>

SYLABS – WHAT IS IT?




- SyLabs is the entity supporting Singularity development
 - SyLabs is to Singularity Pro as Red Hat is to RHEL
- **Singularity Pro:**
 - Officially maintained/supported package of open source project
 - Priority notification of security critical updates
- Additional cloud based offerings to add value to open source software
- Interested in collaboration? Get in touch – bauerm@sylabs.io



SyLabs – Official Announcement Coming Soon

<https://sylabs.io>

SINGULARITY: CONTRIBUTORS AND THANKS!





Singularity Containers

Containers for ...
HPC, EPC, Compute Driven Analytics, Machine Learning, Reproducible Research, and more...

<https://sylabs.io>

CONTACT INFORMATION



Michael Bauer

bauerm@sylabs.io

@bauerm97 on GitHub

Gregory Kurtzer

Project Lead, Singularity
CEO, SyLabs

g@sylabs.io

@gmkurtzer on GitHub