




CSCS

Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich



Sarus - An OCI-compliant container engine for HPC

HPCW 2019: 5th High Performance Containers Workshop

Lucas Benedicic, CSCS

June 20th, 2019

Comparison with existing solutions

	Suitable for HPC	Pluggable vendor support (standard OCI hooks)	User experience	Admin experience	Maintenance effort
Docker					+
Singula				+	+
Charlie				+	+
Shifter				-	-
LXC				-	+
runc		+++	---	+++	+
Sarus	+				

■ **Suitable for HPC**

- Single squashfs image (parallel filesystem friendly)
- Image loop mount + RAM filesystem (fast image accesses)
- WLM compatible
- Native MPI support
- Native GPU support

Comparison with existing solutions

	Suitable for HPC	Pluggable vendor support (standard OCI hooks)	User experience	Admin experience	Maintenance effort
Docker	--				+
Singularity	+				+
Charliecloud	+				+
Shifter	+				-
LXC	-				+
runc			--	+++	+
Sarus	+	++			

■ **Pluggable vendor support**

- OCI hooks support (runc)
- NVIDIA Container Runtime Hook

Comparison with existing solutions

	Suitable for HPC	Pluggable vendor support (standard OCI hooks)	User experience	Admin experience	Maintenance effort
Docker	--	--			
Singularity	+	-			
Charliecloud	+	--			
Shifter	+	-			
LXC	-	-			
runc		++		++	+
Sarus	+	++	++		

■ User Experience

- Docker-like CLI
- Docker Hub integration
- OverlayFS (writable container filesystem)
- Preserve identity and file permissions

Comparison with existing solutions

	Suitable for HPC	Pluggable vendor support (standard OCI hooks)	User experience	Admin experience	Maintenance effort
Docker	--	-			+
Singularity	+				+
Charliecloud	+	-			+
Shifter	+				-
LXC	-				+
runc		+++	--		+
Sarus	+	+++	+++	+	

■ **Admin experience**

- Single executable binary (easy deployability)
- Customize OCI hooks per system
- Container isolation through **PID** and **runc**

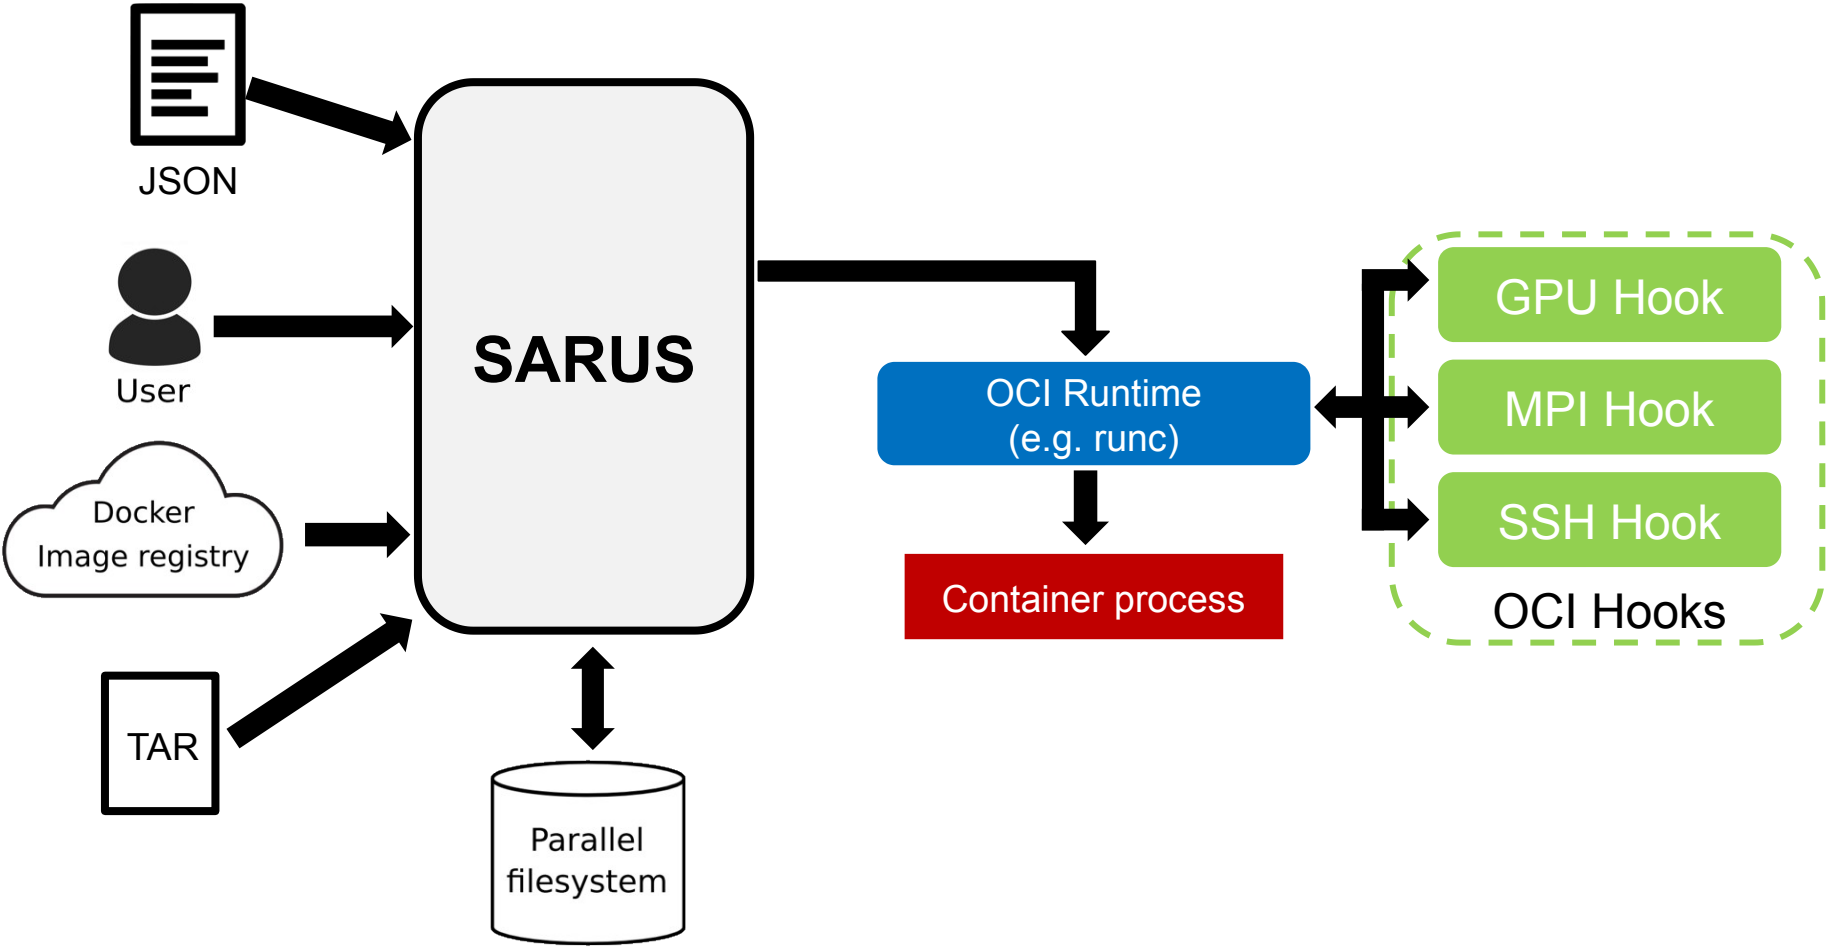
Comparison with existing solutions

	Suitable for HPC	Pluggable vendor support (standard OCI hooks)	User experience	Admin experience	Maintenance effort
Docker	--	--	+		
Singularity	+	-			
Charliecloud	+	--			
Shifter	+	-			
LXC	-	-			
runc		++	--	++	
Sarus	+	++	++	+	+

■ **Maintenance effort**

- Reuse **runc** as the core runtime
- Reuse other OCI-compliant software
- Well tested (unit test coverage 84%)

Architecture overview



Conclusion

Sarus is a container engine for HPC, compliant with open standards, featuring:

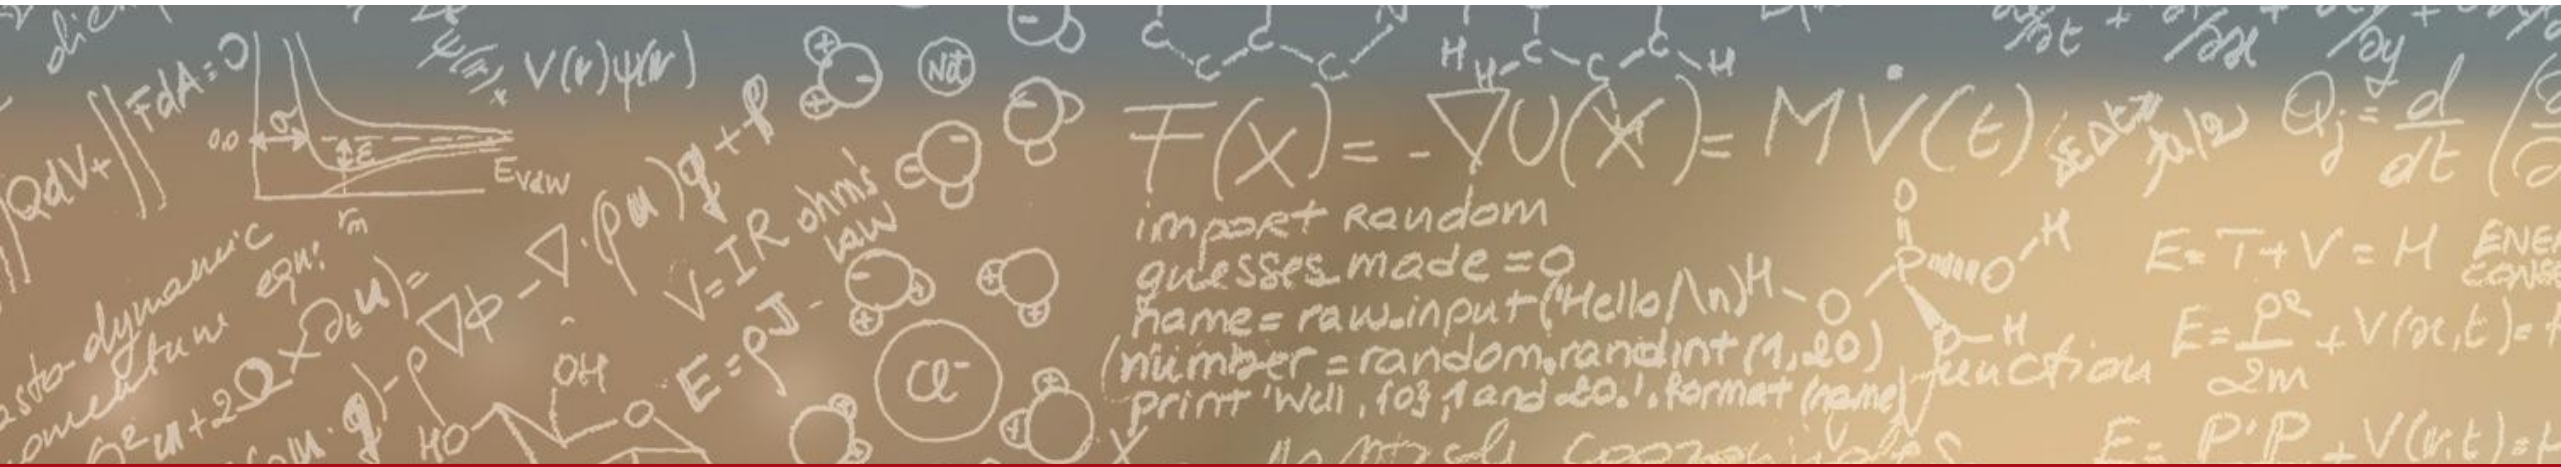
- Transparent native performance through OCI hooks
- Consistent UX with Docker: small learning curve
- Enables use of standard, open, upstream components on HPC systems
- Extensible architecture encourages vendor engagement and improves maintainability



CSCS

Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich



Thank you for your attention.