

CYBERUS

TECHNOLOGY

Automated hypervisor testing and benchmarking on hardware

Markus Napierkowski, Sebastian Eydam
KVM Forum 2024

Who we are



About us

- Self-funded, boot-strapped tech firm
- HQ in Dresden, 25 employees all over Germany

Qualifications

- Team of former Intel and FireEye experts
- Meltdown & Spectre discovery
- Granted several awards



CYBERUS TECHNOLOGY

Fields of Expertise




- Virtualization (Cyberus Hypervisor)
- Test automation (Cidoka)
- Software Lifetime Extension (KronoCore)

Open-Source Interests:

- Cloud Hypervisor / VirtualBox
- Linux KVM
- NixOS / nixpkgs



Our solution for on-hardware test automation: "SoTest"



Our solution for on-hardware test automation: "SoTest"

ID	Priority	Creation Time	User Name	Test Run Name	Progress	Running	Passed	Failed	Disabled	Total
83046	-1	2024-09-07 19:25:06 GMT+2	cyberus/engineering	virtualbox cyberus-release-7-0-20: Native SR-IOV	<div style="width: 100%; height: 10px; background-color: green;"></div>	0	6	0	0	6
83045	-1	2024-09-07 19:20:32 GMT+2	cyberus/engineering	virtualbox cyberus-release-7-0-20: KVM SR-IOV	<div style="width: 100%; height: 10px; background-color: green;"></div>	0	6	0	0	6
83044	-1	2024-09-07 19:15:50 GMT+2	cyberus/engineering	virtualbox cyberus-release-7-0-20: Native Benchmarks	<div style="width: 94%; height: 10px; background-color: green; border: 1px solid red;"></div>	0	17	1	0	18

Name	Tags	amd_m75q_gen2	amd_t16_gen2	coffee_lake_rw14	svp_alder_lake_p360	svp_alder_lake_t16	svp_comet_lake_x13	svp_kaby_lake_r_t480s	svp_raptor_lake_t16	svp_whiskey_lake_t490
Native Benchmark: hdparm	project:virtualbox ram:>=16g	✓	✓	✓	✓	✓	✓	✓	✓	✓
Native Benchmark: kernel-compile	project:virtualbox ram:>=16g	✓	✓	✓	✓	✓	✗	✓	✓	✓



Challenges

- How can we make arbitrary hardware remote controllable?
- How can we integrate on-hardware testing into the development lifecycle?
- What about performance benchmarking?
- How can we handle throughput issues?
- How can we handle flakiness?

Making arbitrary hardware remote-controllable



Power Cycling



Software Configuration



Observe Behavior





Making arbitrary hardware remote-controllable



Power Cycling



Software Configuration



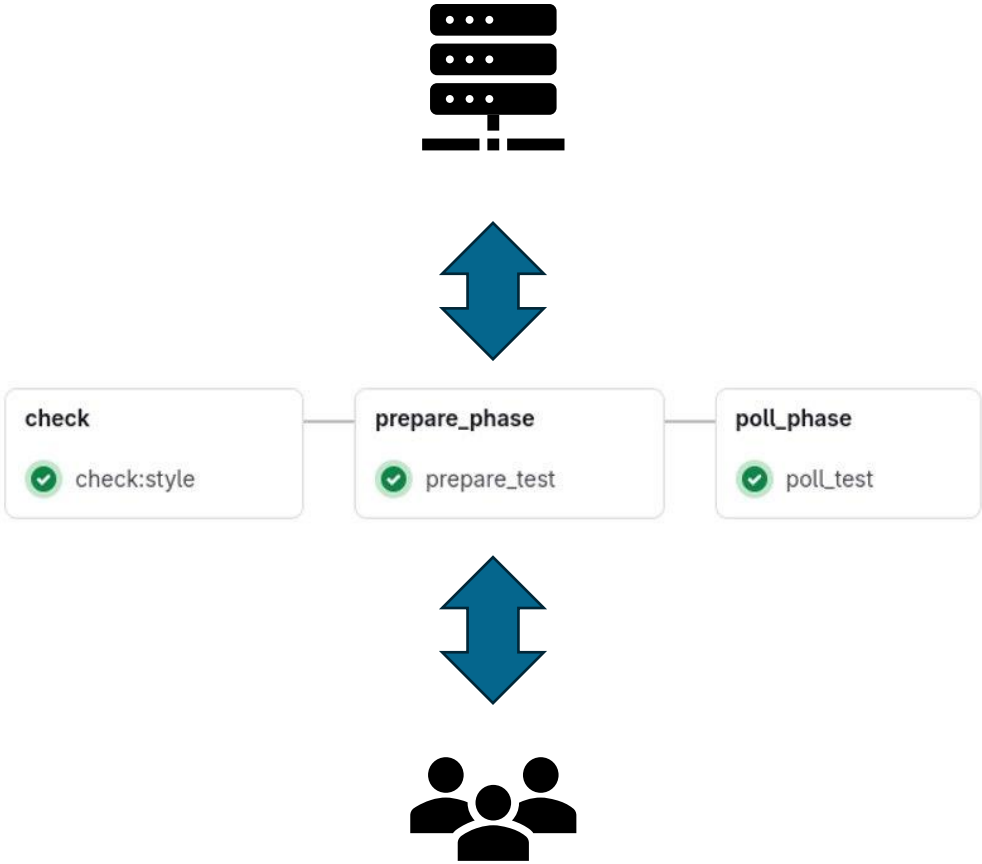
Observe Behavior



Integration into developer workflow



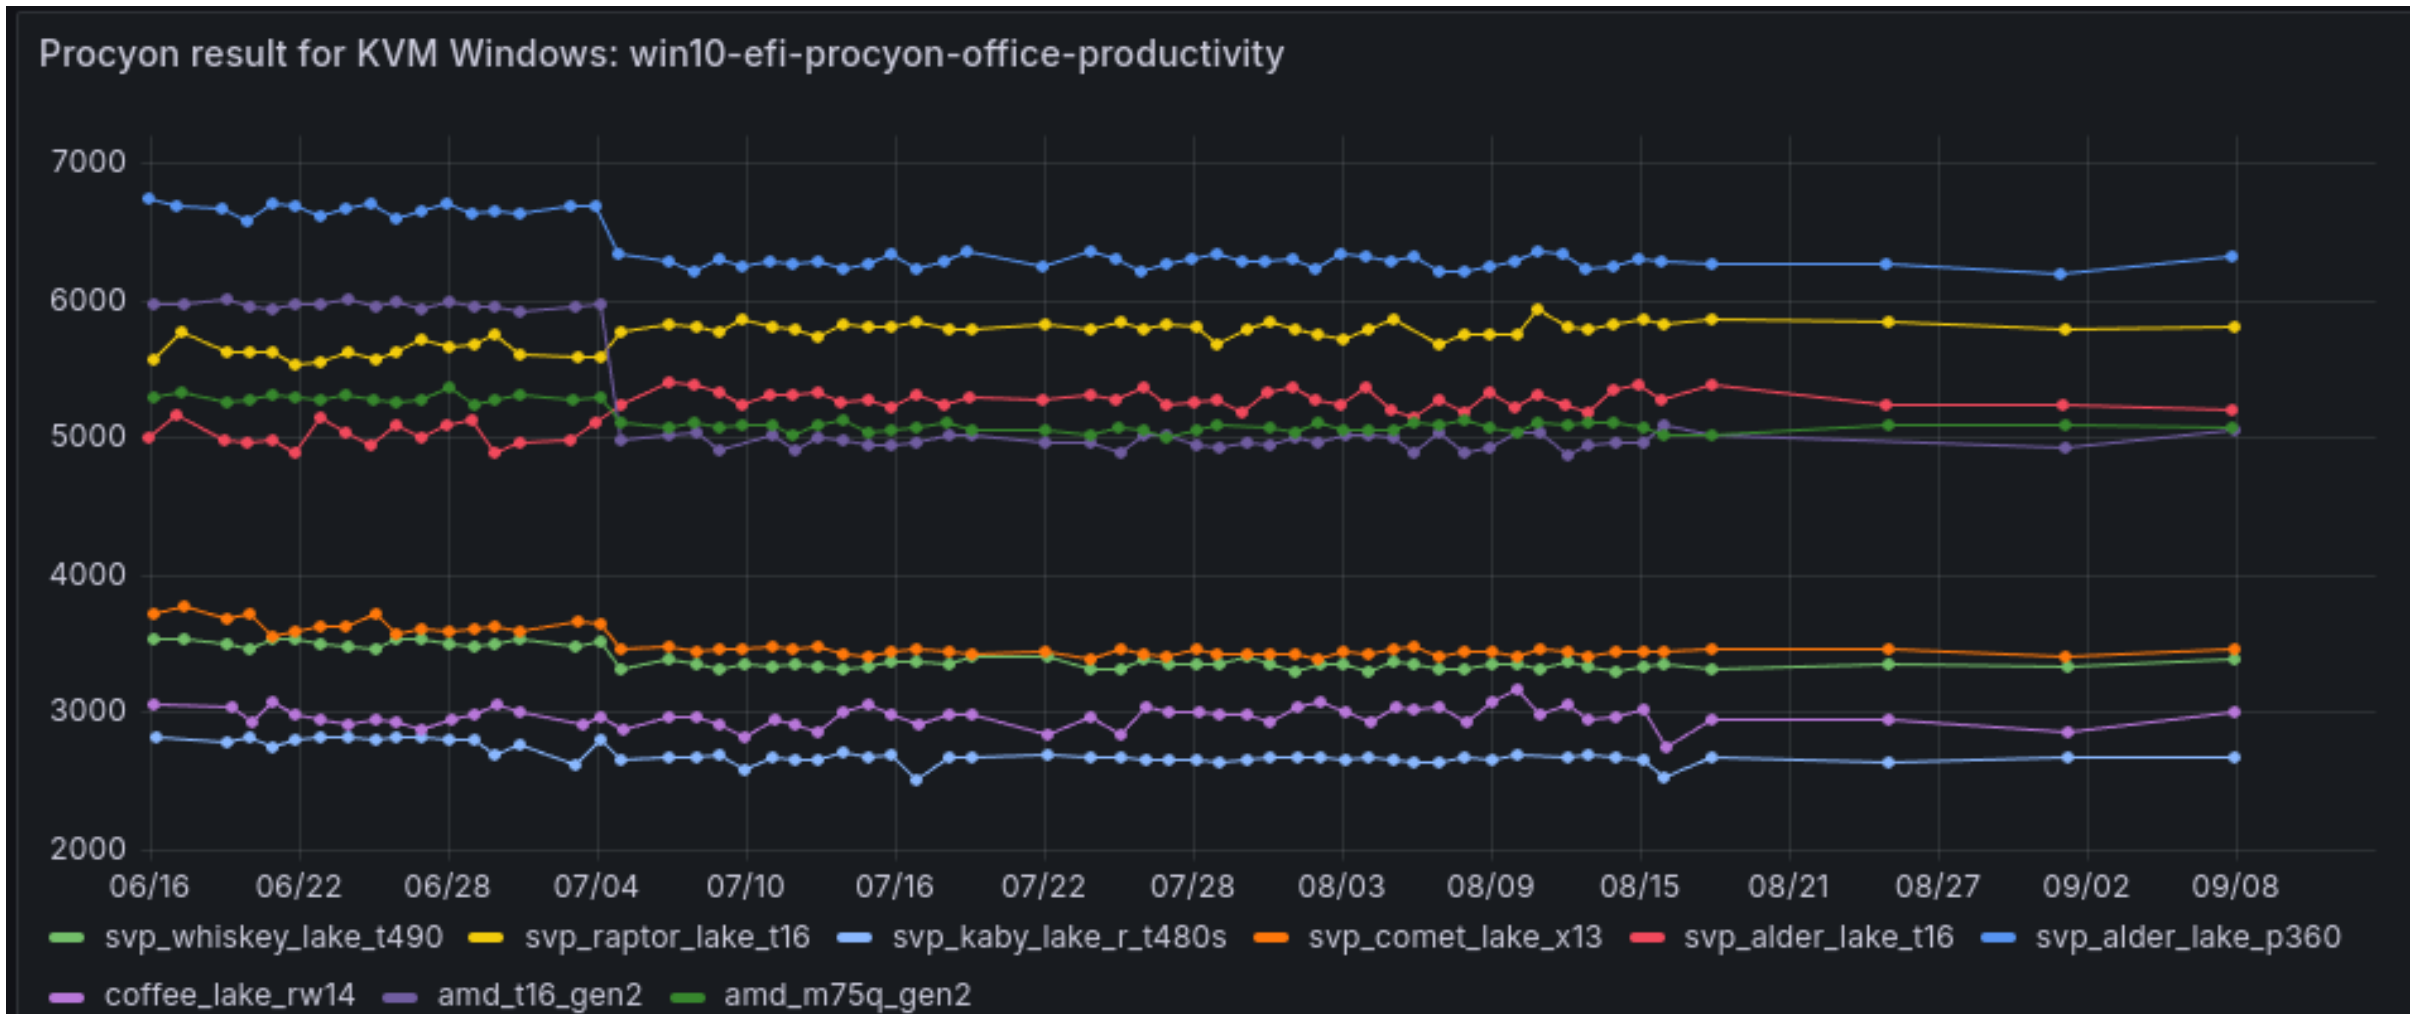
Integration into developer workflow



Integration into developer workflow



Performance Benchmarking




Performance Benchmarking

Disk Speed RND4K Q32T16 WR

Machine Name	Windows 11 SATA	Windows 10 SATA	Windows 11 NVMe	Windows 10 NVMe
intel_whiskey_lake_t490	 1.40	 1.43	 0.963	 1.00
intel_raptor_lake_t16	 1.16	 1.22	 0.835	 0.853
intel_kaby_lake_r_t480s	 1.40	 1.44	 0.990	 1.03
intel_comet_lake_x13	 1.37	 1.44	 0.967	 0.988
intel_alder_lake_t16	 1.14	 1.17	 0.813	 0.836
intel_alder_lake_p360	 1.15	 1.14	 0.999	 0.998
intel_coffee_lake_rw14	 1.34	 1.33	 1.03	 0.959
amd_t16_gen2	 1.68	 1.68	 0.952	 0.960
amd_m75q_gen2	 1.38	 1.39	 0.934	 0.939


Optimizing Throughput

- Don't run all tests every time
- Skip reboots




Optimizing Throughput

- Don't run all tests every time
- Skip reboots
- "Any" Items




Optimizing Throughput

- Don't run all tests every time
- Skip reboots
- "Any" Items
- Prioritized scheduling



Optimizing Throughput

- Don't run all tests every time
- Skip reboots
- "Any" Items
- Prioritized scheduling
- Caching



Handling flaky tests

Temporary infrastructure issues

- Network Problems
- Machines in unexpected state
- Hardware issues

Issues in the test itself

- E.g., Windows driver installation fails

Summary

Automated Testing on Hardware can be hard.
But it pays off!