

Tango Controls Benchmarking Suite

M. Liszcz, P.P. Goryl, S2Innovation, Kraków, Poland

Benchmarking Suite

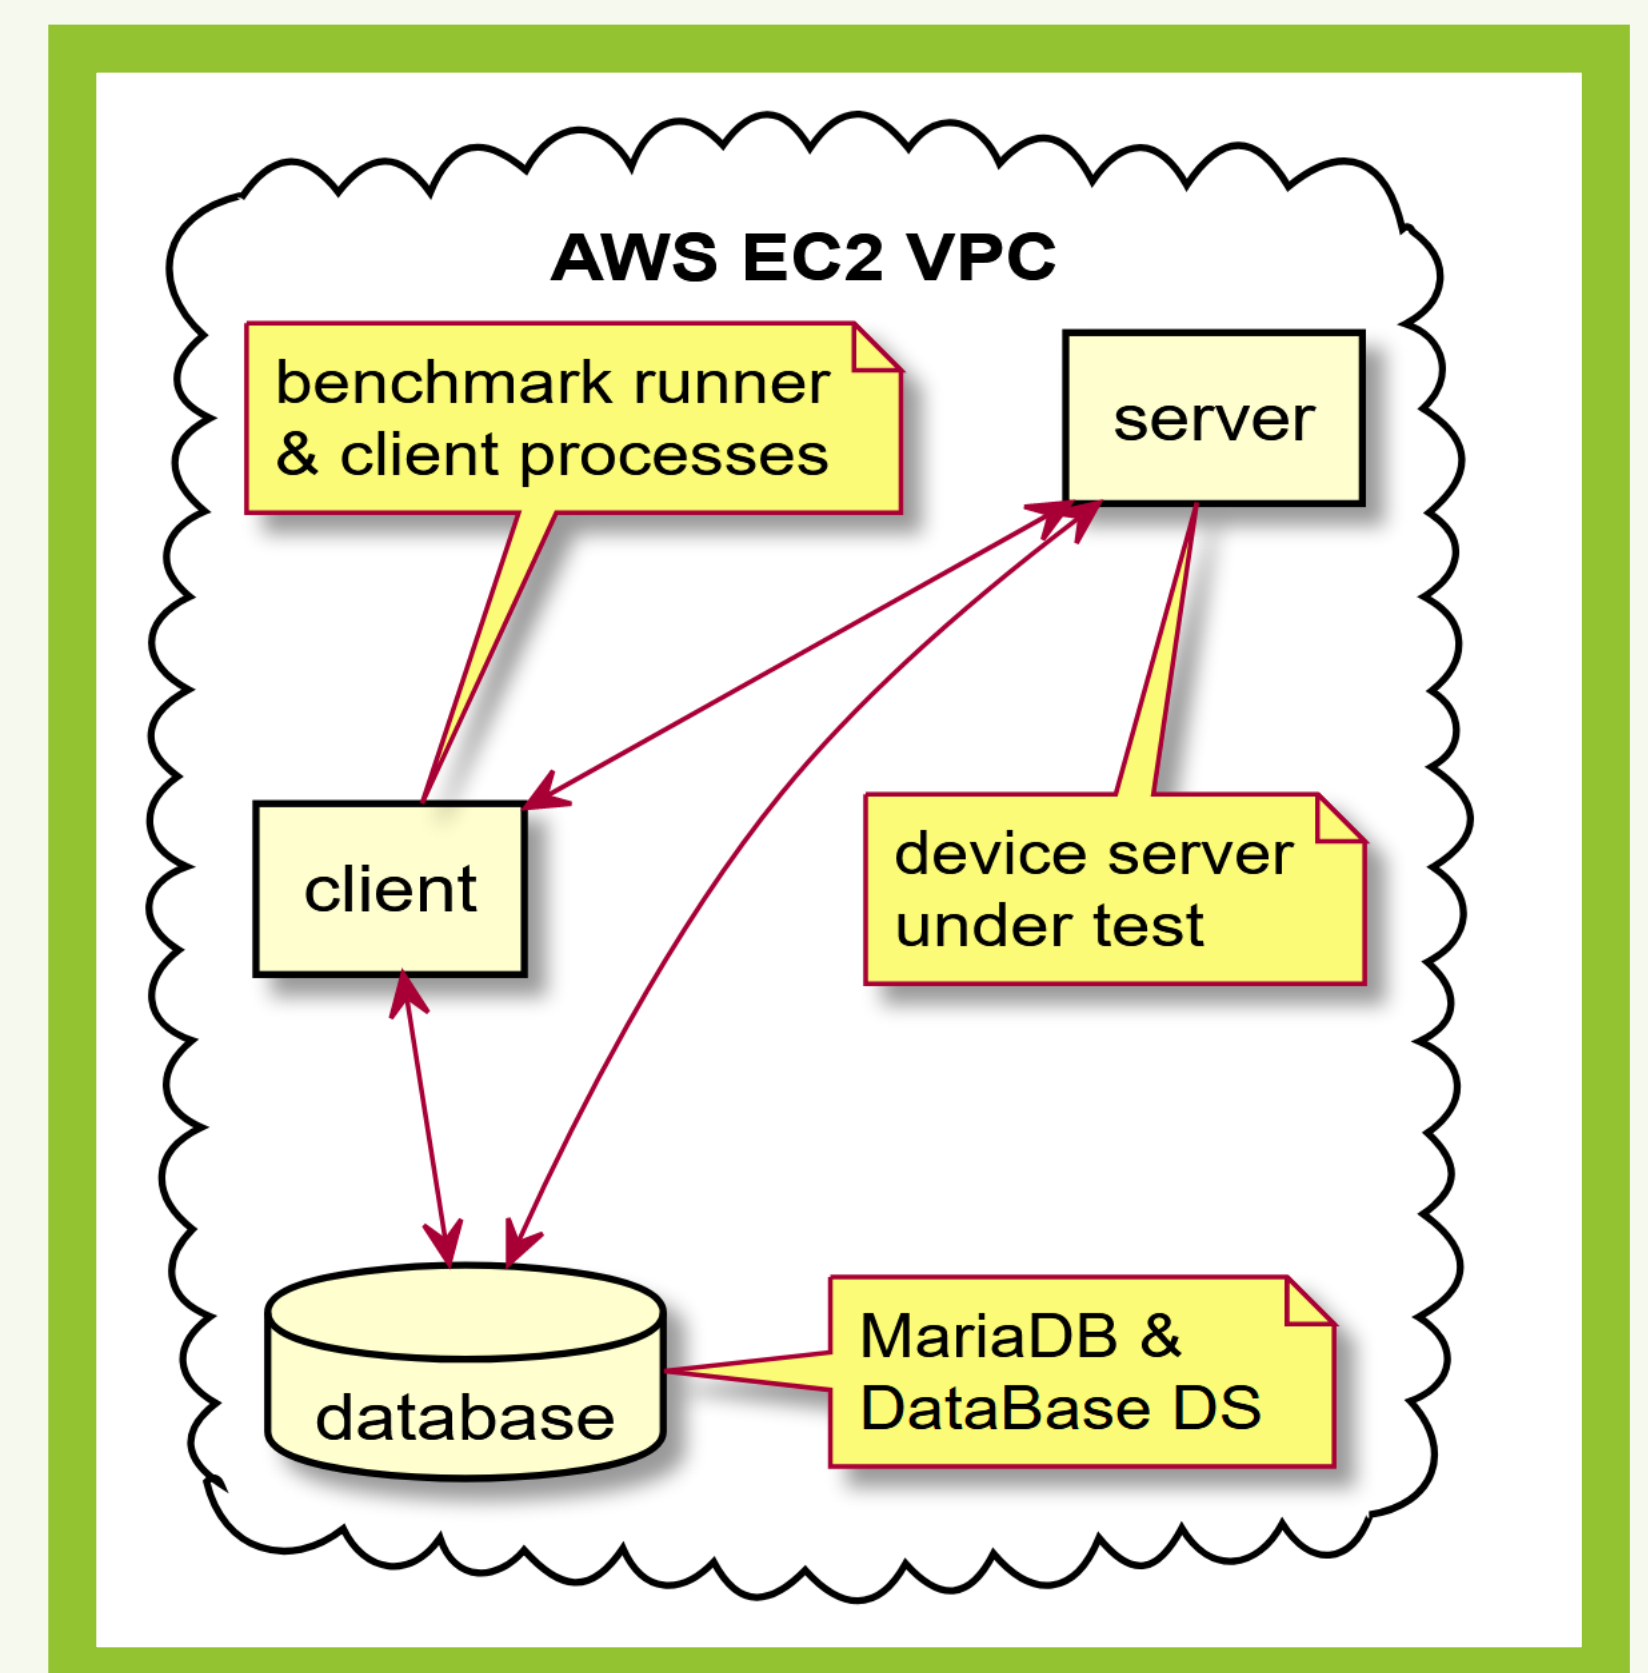
- ▶ A set of tools for measuring efficiency and performance of Tango Controls:
- ▶ **Benchmark scripts:**
 - ▶ Attribute read and write, command execution, event subscription, etc. ...
 - ▶ Start multiple parallel clients,
 - ▶ Produce reports in CSV and ReST formats.
- ▶ **Target device servers:**
 - ▶ Provide interface required by the scripts,
 - ▶ Implemented in Python, C++ and Java.
- ▶ **Benchmark runner:**
 - ▶ Reads configuration from a YAML file,
 - ▶ Starts device servers and runs the tests.

Tango Controls can be used at both small and very large laboratories and scientific facilities. Deploying Tango at large scale requires solutions for **monitoring efficiency, performance and resource utilization**. To address this need the Tango Controls Benchmarking Suite was developed.

Tests on Amazon AWS EC2

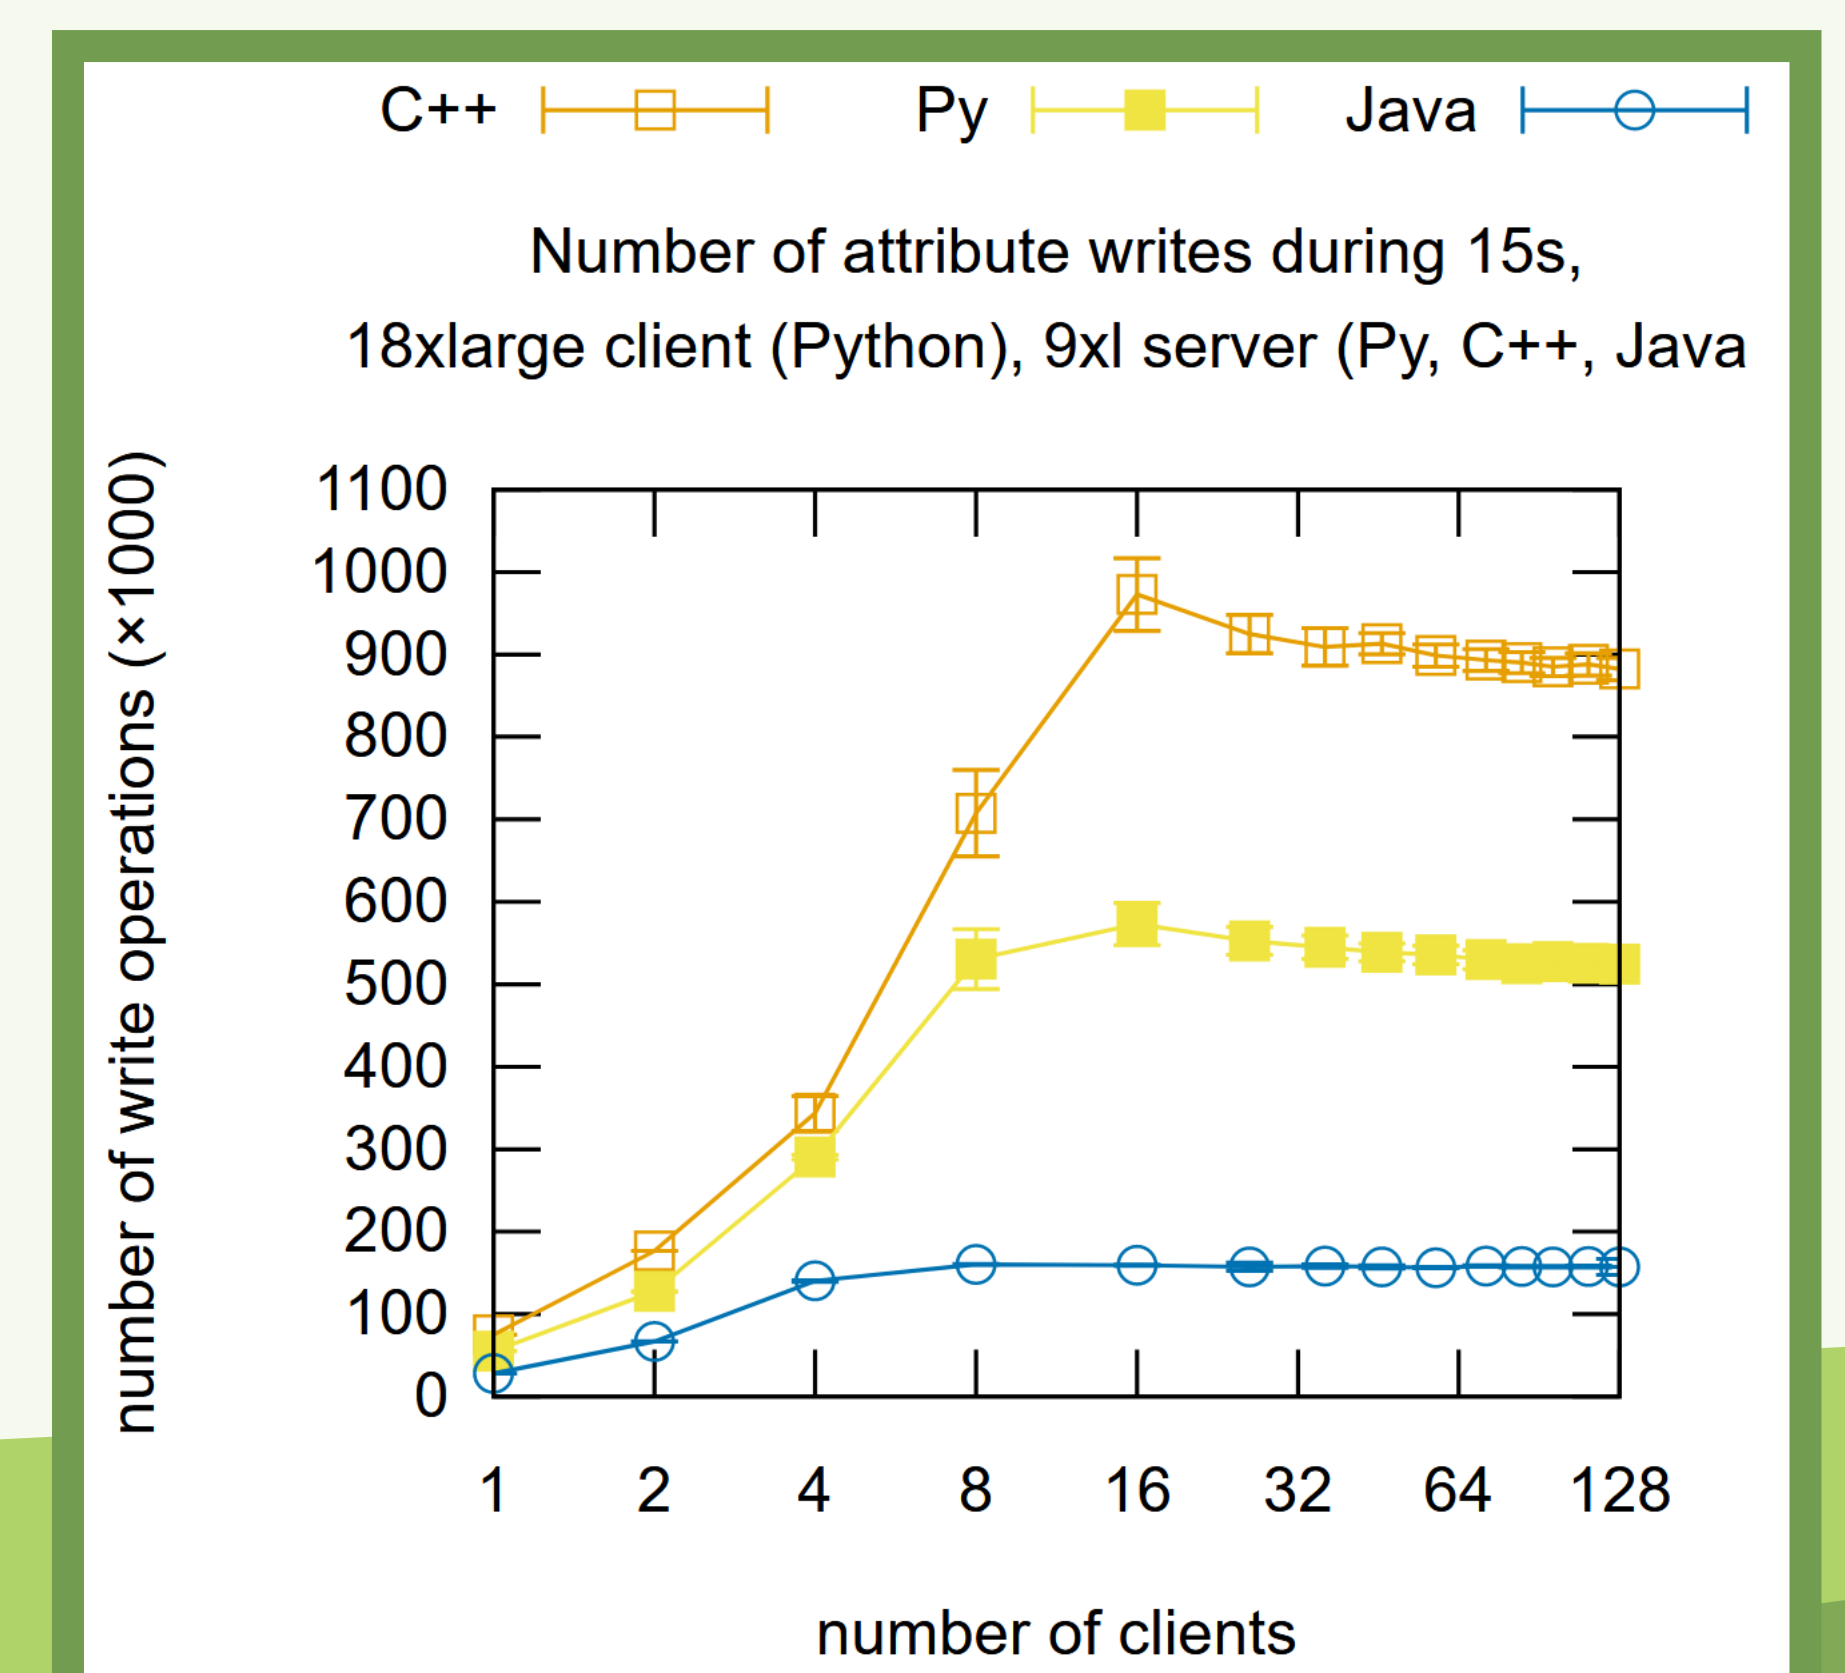
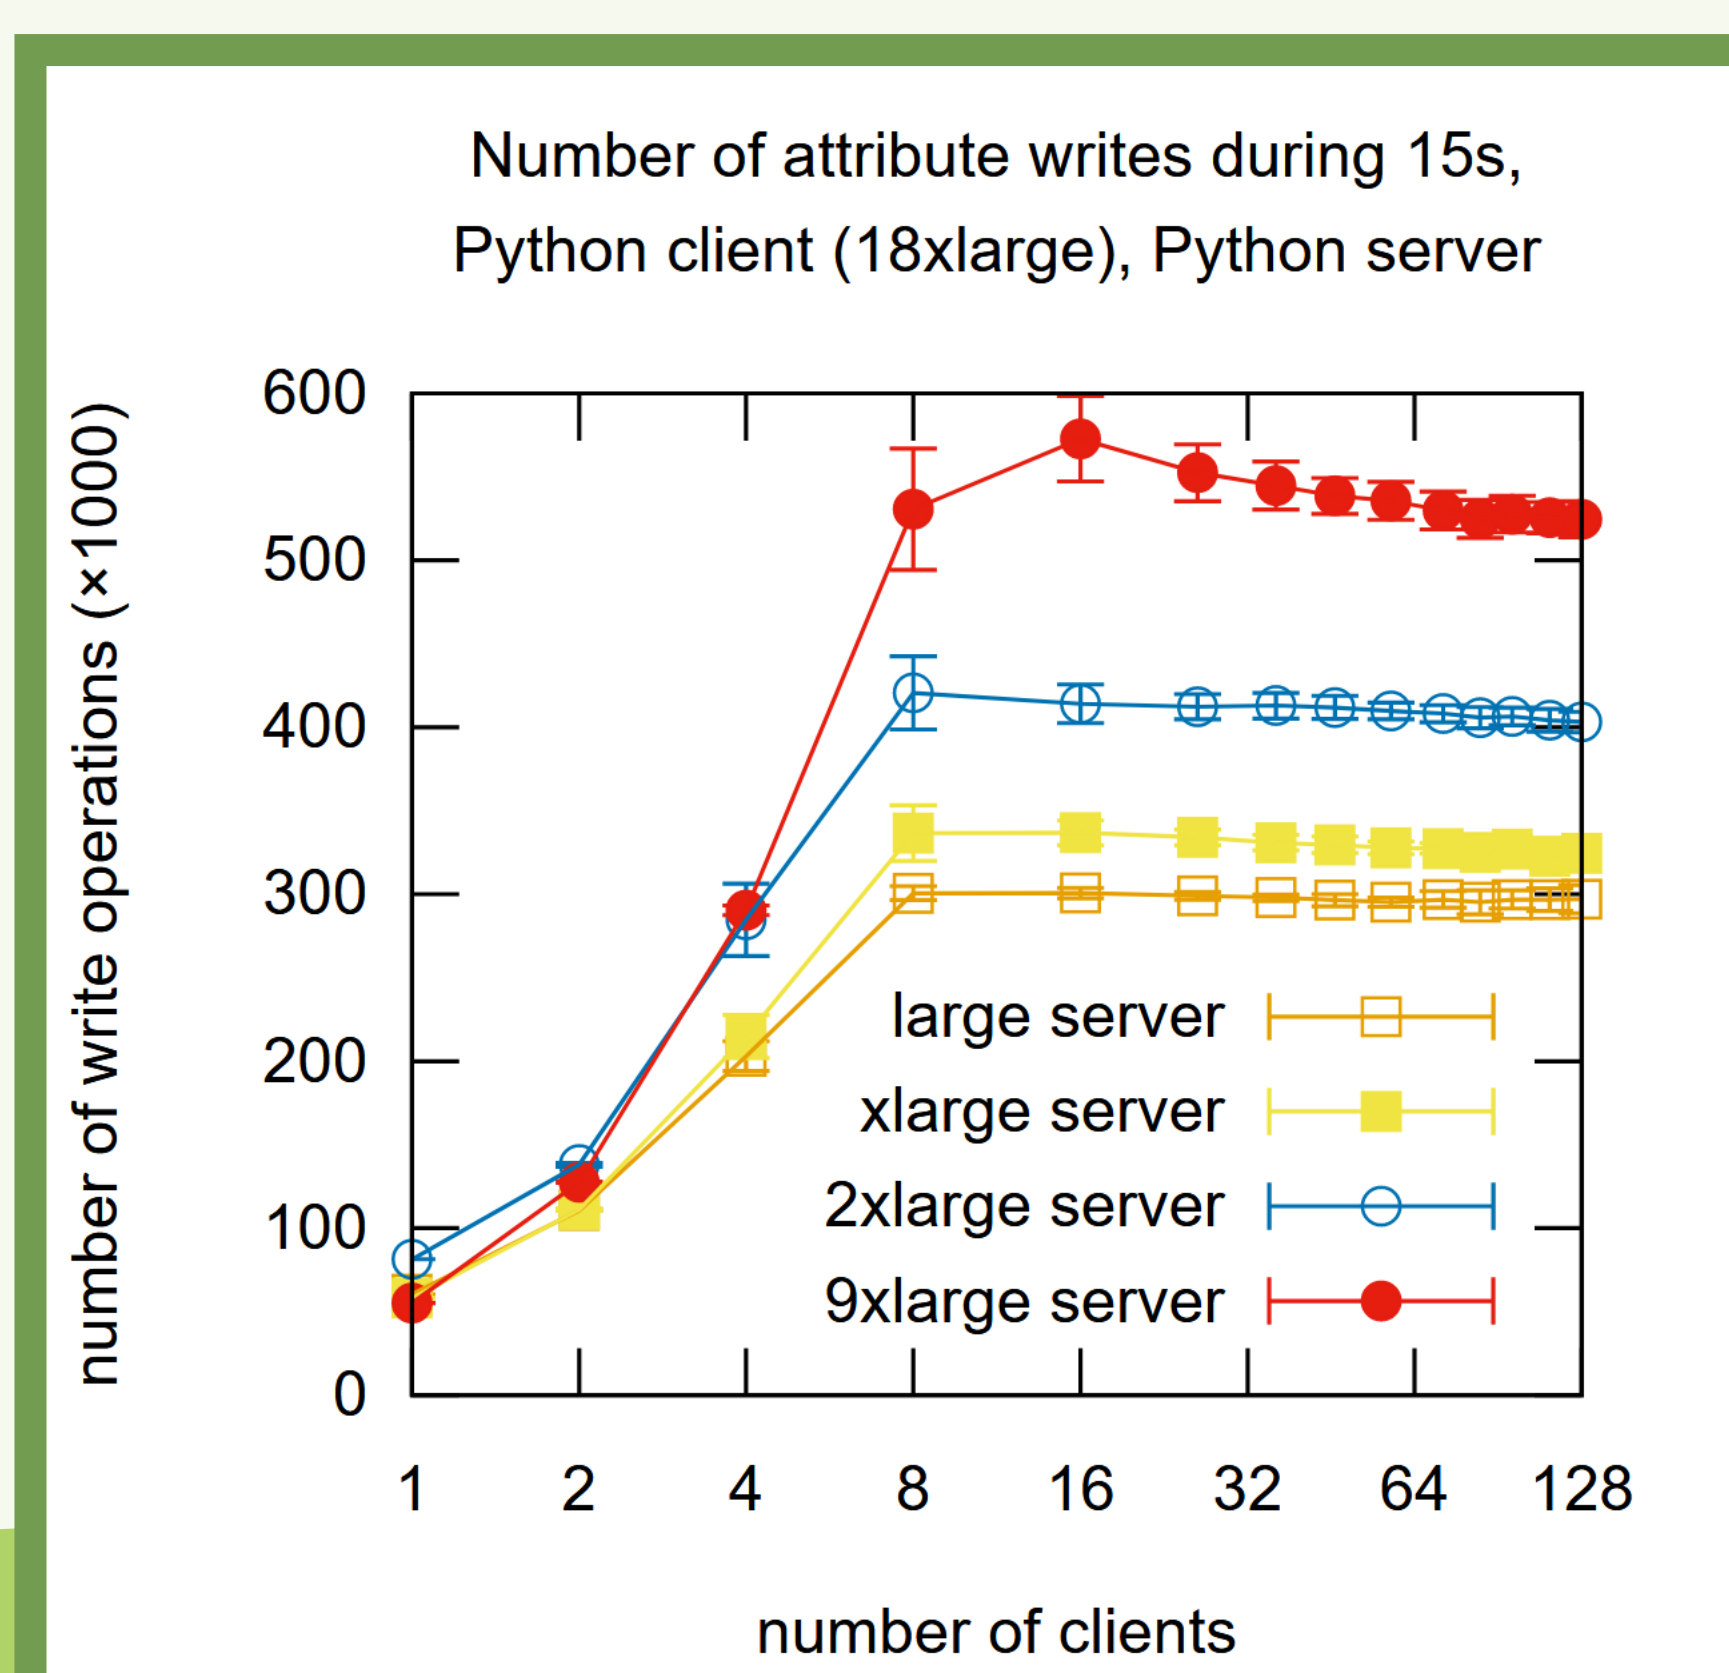
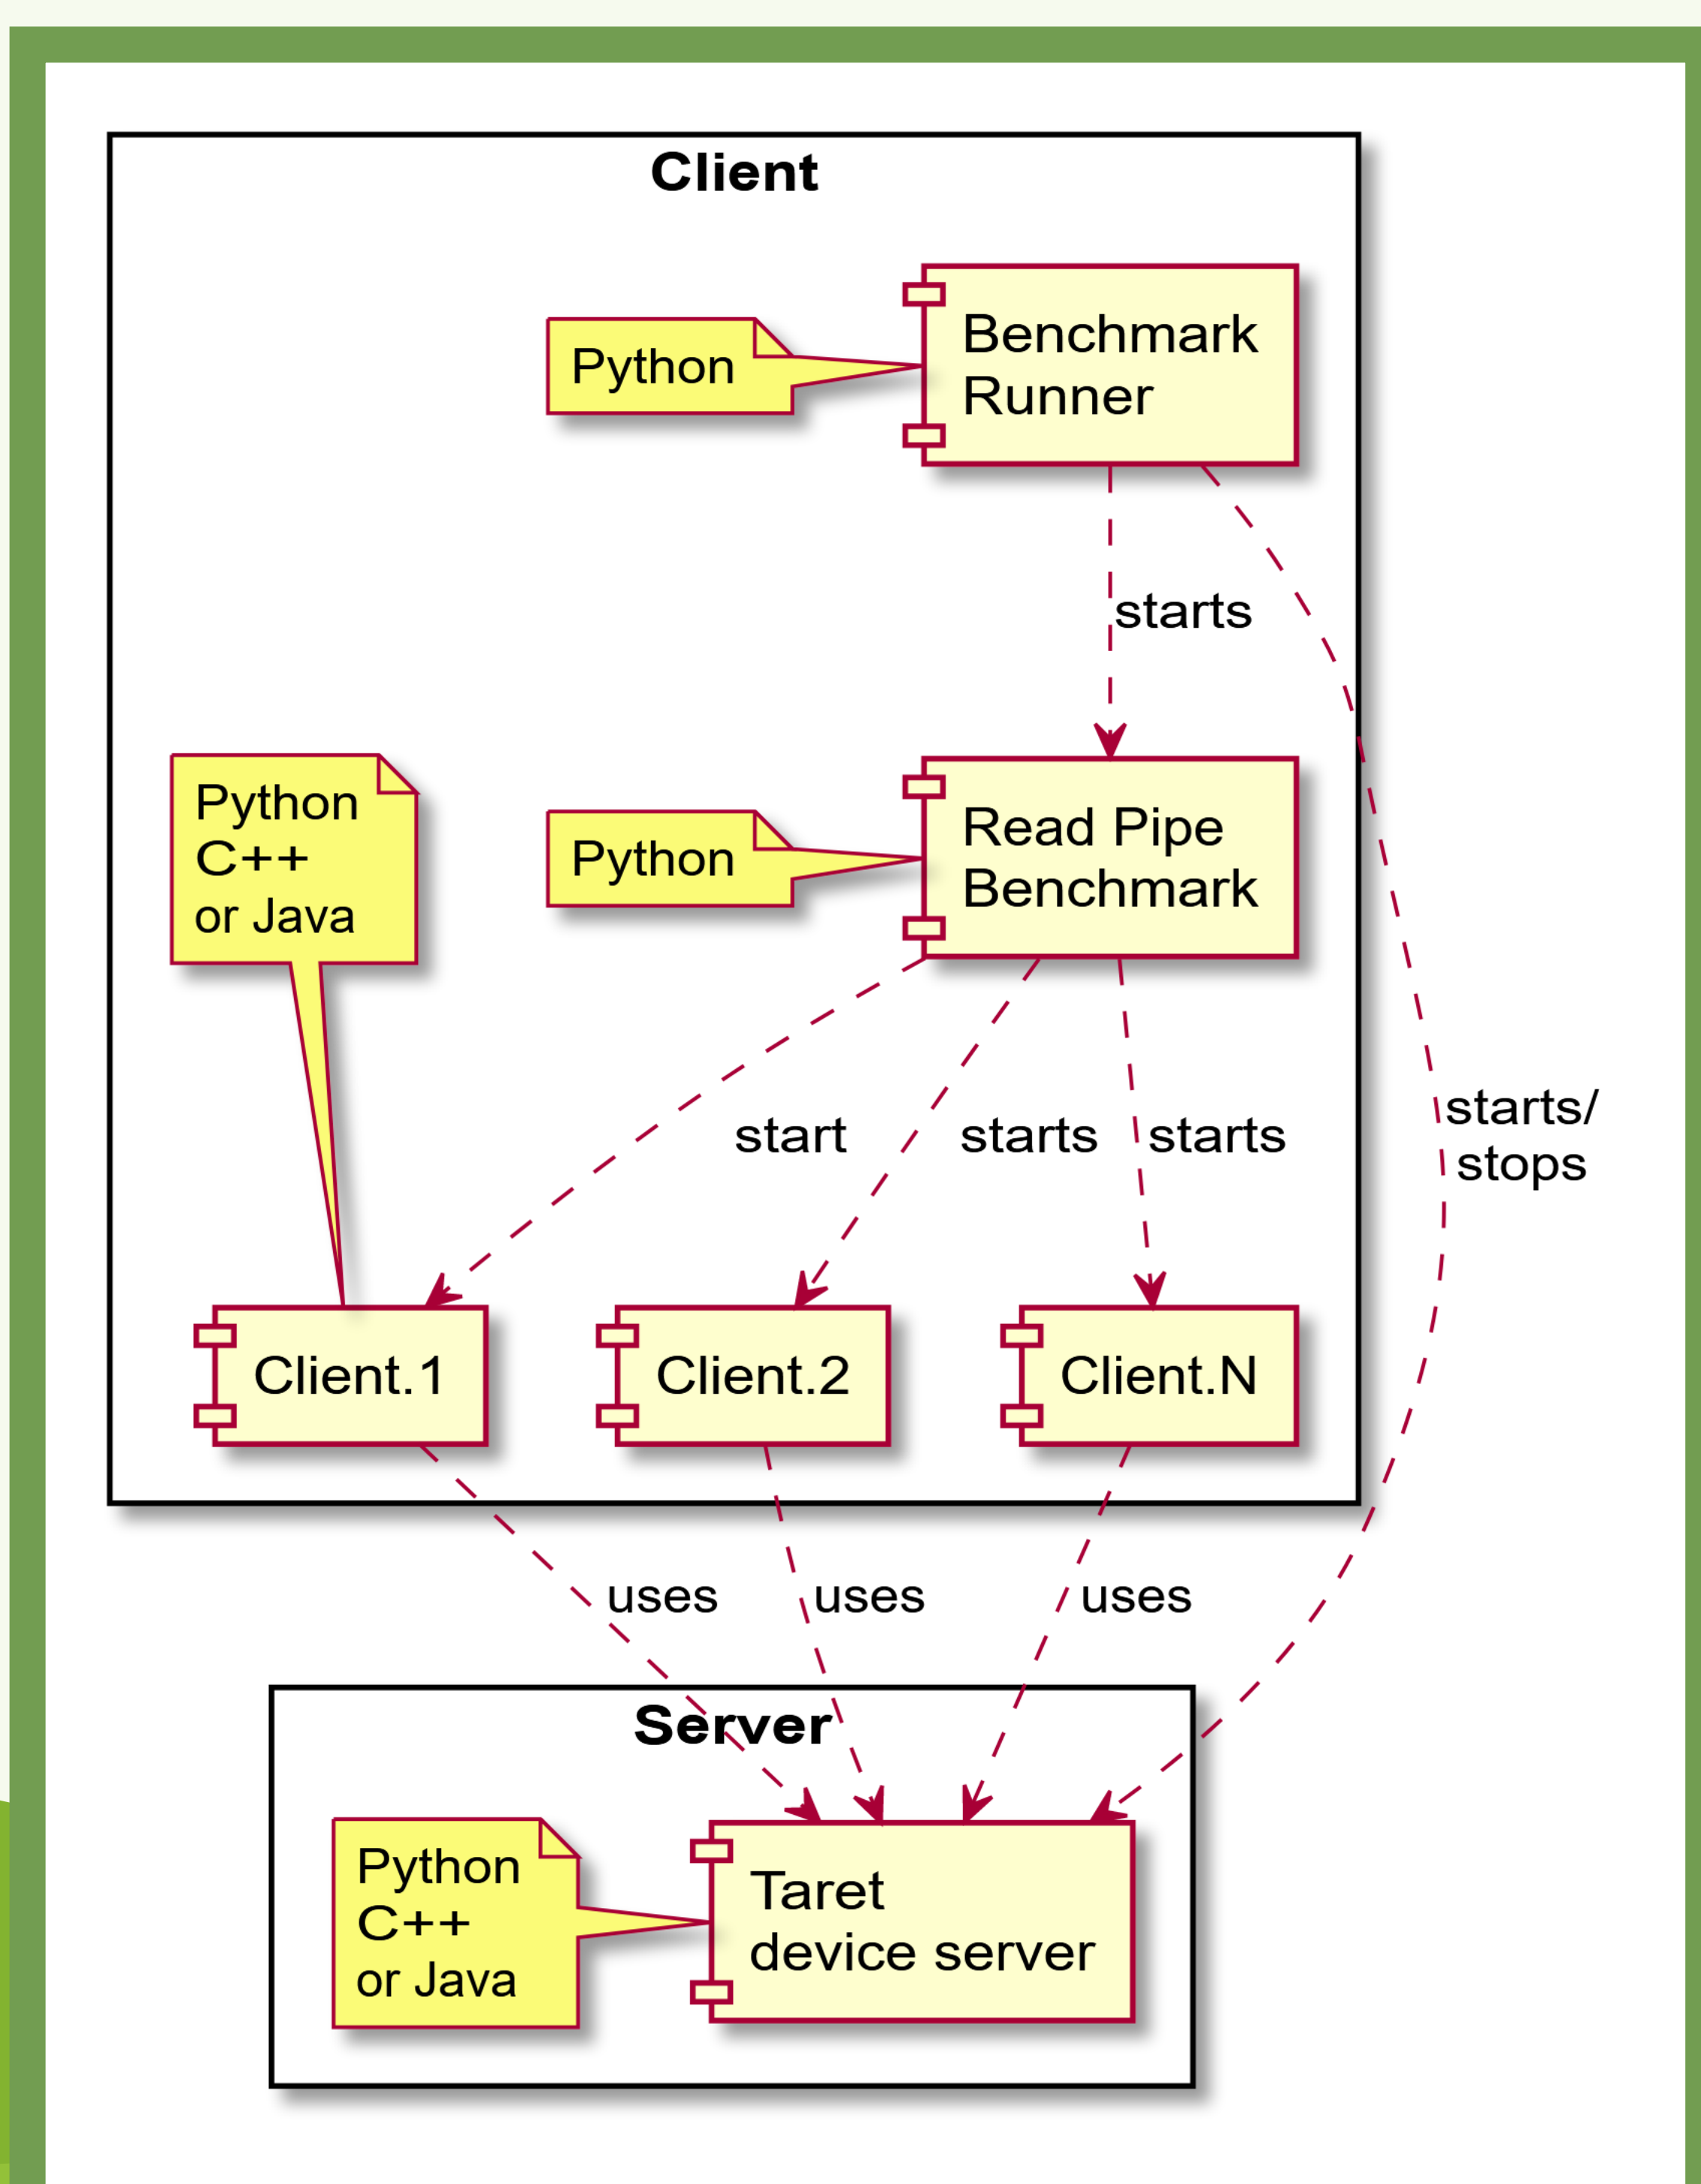
- ▶ A test setup for performance measurements was assembled on Amazon AWS EC2 platform,
- ▶ Different client and server instance sizes were compared.

Instance	vCPUs	Mem. [GiB]
c5n.large	2	5.25
c5n.xlarge	4	10.50
c5n.2xlarge	8	21.00
c5n.4xlarge	16	42.00
c5n.9xlarge	36	96.00



Attribute access performance

- ▶ Server performance increases with the number of available CPUs,
- ▶ Client performance is often limited by the server performance,
- ▶ C++ server implementation is much faster than Python or Java.



Acknowledgements: The project is funded by the Tango Controls Collaboration. This work not be possible without the contribution from Jan Kotanski (DESY) who worked on the first version of the Benchmarking Suite and developed most of the benchmarks. The authors also thank Andy Götz (ESRF) who suggested and made the tests on AWS possible and for his comments.

