

# PEDSEWAN: Platform for the Evaluation of Distributed Systems in Emulated Wide-Area Networks

Johannes Köstler, Hans P. Reiser  
{jk,hr}@sec.uni-passau.de

December 18, 2015

Our current research project tries to optimize reliable and distributed middleware. In order to track down the components' interoperation and to verify possible optimizations a comprehensive analysis of the system's runtime behavior under varying environmental conditions is required. This could be done by simulating relevant system aspects or by executing the system in testbeds or real deployment environments. Simulation works well for particular system aspects, but covering all implications from the components' interactions is a much more difficult task. The execution in the real deployment environment on the other hand has its own drawbacks. Besides the fact that real hardware may cause high costs the main disadvantage is that the changing environmental conditions make it impossible to obtain deterministic results.

Emulation forms a compromise combining the advantages of both approaches. The unmodified system can be executed in an environment that emulates any desired conditions – resulting in reproducible results. As our research area is primarily concerned with reliable group communication systems, our focus lies on the condition of the interconnecting network and on the health state of the participating execution hosts. But emulation also comes with setup and bootstrapping efforts which in most cases exceed those from simulation or real deployments as it requires both a real deployment and the modelling of certain aspects.

Basically all emulation platforms use tools like KauNet<sup>1</sup> or stress-ng<sup>2</sup> to emulate network conditions and workloads. Modelnet [2] requires a set of dedicated machines and the definition of par-

ticipating nodes as well as the network topology. Some nodes mimic this topology, whereas the unmodified test application is executed on the other nodes. Thus Modelnet provides basic emulation functionality, but it comes with high bootstrapping efforts and is no longer maintained. DieCast [1] – a successor of Modelnet – introduces time dilation where the execution time is stretched to emulate more processing resources. This allows the emulation of more virtual nodes, but does not cover the bootstrapping issue. The list could be continued, but the bottom line is that we could not find any existing solution that fit our needs.

This talk gives an overview over existing concepts and points out identified drawbacks. We introduce our own approach that reliably emulates all required external conditions and fully automates all bootstrapping efforts. We are especially interested in feedback and further suggestions as we don't want to come up with another isolated application, but would like to see that others can also benefit from this solution.

## References

- [1] D. Gupta, K. V. Vishwanath, M. McNett, A. Vahdat, K. Yocum, A. Snoeren, and G. M. Voelker. Diecast: Testing distributed systems with an accurate scale model. *ACM Trans. Comput. Syst.*, 29(2):4:1–4:48, May 2011.
- [2] K. Vishwanath, D. Gupta, A. Vahdat, and K. Yocum. Modelnet: Towards a datacenter emulation environment. In *Peer-to-Peer Computing, 2009. P2P '09. IEEE Ninth International Conference on*, pages 81–82, Sept 2009.

<sup>1</sup><http://www.kau.se/en/kaunet>

<sup>2</sup><http://kernel.ubuntu.com/~cking/stress-ng/>