

This manuscript proposes a new approach for live virtual machine (VM) migration based on checkpointing and logging techniques and it demonstrates that, compared to traditional pre-copy based live VM migration, this approach can significantly reduce migration overhead for certain workloads. Compared to this manuscript's previous conference version, this version includes new techniques for VM disk migration and multi-processor VM migration as well as new experimental evaluation of emulated wide-area VM migration.

The authors have made substantial improvement upon their last version of this manuscript. However, as pointed out in the first-round reviews, the authors need to spend enough efforts to improve the discussions on VM disk migration and multi-processor VM migration, which are the main contributions made by this manuscript compared to its original conference version.

In the discussion on VM disk migration (Section 5.2), this manuscript proposes an approach based on pre-copy and post-copy techniques, but it is lacking a detailed comparison study between this approach and the related work. Although the existing solutions for VM disk migration are summarized in Section 2, the authors should also at least qualitatively compare them to this manuscript's approach. (In fact, this seems to be a general problem for the entire related work section. When readers read such a discussion, they expect to see not only what the related work is about but also the connection between the existing work and this new work, e.g., what is the similarity or difference and what can be leveraged or improved.)

There are a couple of issues with the evaluation of VM disk migration in the manuscript too. In the evaluation of network throughput of wide-area VM migration (Section 6.4, Fig. 14), the assumption of the COW image being small because of the short period of migration time seems to be problematic. If the base VM disk image is pre-deployed on both source and destination hosts, then the COW image would include all the changes made by the VM since it starts running on the source. If the base disk image is pre-copied on demand before every VM migration, then the COW image would include all the changes made by the VM while the base image is copied to the destination, which could take a significant amount of time over WAN. In either way, a VM hosting a typical I/O active application (such as kernel compilation) could generate a considerable amount of dirty data that are far greater than 2MB. So the authors need to clarify what is the real assumption underneath of this experiment. This assumption may also exist for the other WAN-based experiments in the manuscript and if that is the case the authors should clearly state that in the experimental setup.

In addition, the performance overhead due to the proposed VM disk migration is not evaluated in the manuscript. In Section 6.5, only LAN-based migration is studied which does not require the migration of VM disks. However, in WAN-based migration, there could be considerably more overhead due to the use of COW and read logging and redirection, and the readers would be interested in seeing a quantitative evaluation of such overhead.

With respect to multi-processor VM migration, this manuscript proposes to use the CPU hot plug technique to address the issue caused by multi-processor logging. However, it still lacks a

discussion on this technique's implementation specifics and an evaluation of its overhead for VM migration. These details would be important to support the use of this technique as a new contribution of this manuscript. Otherwise, it should be entirely considered as part of the future work.

Finally, this manuscript has quite a few grammar errors and improper uses of English:

Page 4, line 13, right column: "lowest" is inappropriate ("little" sounds better);
Page 5, line 41, right column: "evolution" should be "evaluation";
Page 5, line 44, right column: "special" should be "specific";
Page 6, line 13, left column: "discriminate" should be "difference";
Page 6, line 57, left column: "are" should be removed;
Page 7, line 38, right column: "speed" should be "speed up";
Page 11, line 42, left column: "a little" is inappropriate (the difference is in fact quite significant);
Page 11, line 44, left column: "more close" should be "closer";
Page 13, caption of Fig. 17: "reversed" should be "reserved";
Page 13, line 2, right column: "makes" is inappropriate ("achieves" sounds better);
Page 13, line 3, right column: "unperceived" should be "negligible".

The authors have addressed many comments raised in the first-round reviews, but more work is still necessary to provide a better discussion on the two key issues, VM disk migration and multi-processor VM migration, which are the main differentiators between this manuscript and its previous conference version. Therefore, I recommend the authors to prepare another revision to address these remaining issues.