

仮想マシンの高速マイグレーション機構を利用した 動的なサーバ集約システムの提案

広瀬 崇宏 中田 秀基 伊藤 智 関口 智嗣
産業技術総合研究所 情報技術研究部門

Overview

In recent years, dynamic server consolidation with live VM (Virtual Machine) migration is one of the hot topics in the system software area [5, 1]. By exploiting live migration, it is possible to realize over-commit assignments of VMs for improving resource usage. If VMs are not consuming their assured amounts of computing resources, these VMs are migrated into fewer physical nodes to reduce necessary physical nodes.

To the best of our knowledge, however, these studies rely on precopy live migration mechanisms, which are already available in widely-used VM monitors. We consider such precopy live migration is not suitable for dynamic consolidation systems; because it takes long time to switch the execution host of a running VM, and it is hard to estimate when migration is completed. In precopy live migration, memory pages, being intensively updated by a guest OS at a source host, need to be iteratively copied to a new execution host, before the execution host is switched.

On the other hand, *postcopy* live migration postpones memory page copies after the execution host is switched; it is possible to change the execution host in several hundred milliseconds, and live migration is completed in a deterministic period. For dynamic consolidation systems, postcopy live migration will greatly contribute to relocating VMs on overloaded hosts with minimum latency. There are research papers focusing on postcopy migration mechanisms[2, 4], however, which do not discuss the effectiveness of postcopy for dynamic VM consolida-

tion.

In our ongoing project, we are now developing a new VM consolidation system that is based on our postcopy migration mechanism [3, 6, 7]. The results of preliminary experiments will be presented in the poster session, in which pre/postcopy-based consolidation systems are compared at the viewpoint of performance guarantees and power savings.

Demonstration

The prototype system in our laboratory will be demonstrated through the Internet. The system dynamically optimizes the locations of VMs in response to the number of client requests, so that minimizes power-on physical hosts and assures VM performance guarantees.

参考文献

- [1] F. Hermenier, X. Lorca, J.-M. Menaud, G. Muller, and J. L. Lawall. Entropy: a consolidation manager for clusters. In *Proceedings of the 5th International Conference on Virtual Execution Environments*, pages 41–50. ACM Press, 2009.
- [2] M. R. Hines and K. Gopalan. Post-copy based live virtual machine migration using adaptive pre-paging and dynamic self-ballooning. In *Proceedings of the 5th International Conference on Virtual Execution Environments*, pages 51–60. ACM Press, 2009.

- [3] T. Hirofuchi, H. Nakada, S. Itoh, and S. Sekiguchi. Enabling instantaneous relocation of virtual machines with a lightweight VMM extension. In *Proceedings of the 10th IEEE/ACM International Conference on Cluster, Cloud and Grid Computing*, pages 73–83. IEEE Computer Society, May 2010.
- [4] H. A. Lagar-Cavilla, J. A. Whitney, A. Scannell, P. Patchin, S. M. Rumble, E. de Lara, M. Brudno, and M. Satyanarayanan. SnowFlock: Rapid Virtual Machine Cloning for Cloud Computing. In *Proceedings of the Fourth ACM European Conference on Computer Systems*, pages 1–12. ACM Press, 2009.
- [5] T. Wood, P. J. Shenoy, A. Venkataramani, and M. S. Yousif. Black-box and gray-box strategies for virtual machine migration. In *Proceedings of the 4th Symposium on Networked Systems Design and Implementation*, pages 229–242. USENIX Association, 2007.
- [6] 広瀬崇宏, 中田秀基, 伊藤智, and 関口智嗣. 瞬間的な実行ホスト切り替えを可能とする仮想マシン的高速ライブマイグレーション機構. In *日本ソフトウェア科学会研究会資料シリーズ No.62 (インターネットコンファレンス 2009 論文集)*. 日本ソフトウェア科学会, Oct 2009.
- [7] 広. 崇宏, 中. 秀基, 伊. 智, and 関. 智嗣. 既存 VMM への適用が容易でゲスト透過なポストコピー型仮想マシン再配置機構. *情報処理学会論文誌: コンピューティングシステム*, ACS31, Aug 2010.