Strategies for Reliable, Cloud-based Distributed Real-time and Embedded (DRE) Systems

Kyoungho An (kyoungho.an@gmail.com)
ISIS, Dept. of EECS, Vanderbilt University

1. Background and Challenges

- **Advantages of Cloud computing**
  - Elastic
  - Scalable
  - Cost-effective
- **Challenges of Cloud computing**
  - Hard to guarantee the user's expectations of Quality of Service (e.g., reliability, timeliness)
  - Shared resources by a number of users
    - Compute-intensive processes
    - Network traffic
- **Research Challenges of Cloud Computing to Support DRE Systems**
  1. Need effective resource monitoring
  2. Real-time hypervisors and data-center networks
  3. Predictable dependability mechanisms

2. Literature Survey

- **Research criteria based on literature survey**
  1. Effective resource monitoring (common to all)
  2. Timeliness in data-center networks
  3. Real-time scheduling in hypervisors
  4. High availability via replication of virtual machines.
- **Timeliness in data-center networks**
  1. DCTCP
     - TCP modified protocol
     - Better throughput than TCP
  2. D^3
     - Deadline aware control protocol
  3. D^2TCP
     - Reducing deadline miss ratio
- **Real-time scheduling in hypervisors**
  1. RT-Xen
     - 4 Fixed priority real-time schedulers used in Xen
  2. Scheduler S
     - The modified Xen scheduler for soft-real-time tasks
- **High availability via replication of virtual machines**
  1. Remus
     - Asynchronous replication and speculative execution
     - Supported with Xen
  2. Kemari
     - Lock-stepping and continuous check-pointing approach
     - Supported with KVM
  3. HydraVM
     - Storage-based and memory-efficient approach

3. Addressing Challenge 1 – Need for Effective Resource Monitoring

- **SQRT-C (A scalable and QoS-enabled cloud monitoring system)**
  - Using OMG Data Distribution Service (DDS) real-time publish/subscribe (pub/sub) middleware
- **SQRT-C software artifacts**
  - DDS-based pub/sub communication
  - Monitoring Manager
- **DDS-based pub/sub communication**
  - Disseminating monitoring information for virtual resources from the source (i.e., publishers) to the sinks (i.e., the subscribers)
  - Supporting the QoS requirements
- **Monitoring Manager**
  - Serving as the orchestrator for the deployment of data-writers and data-readers of the DDS pub/sub mechanism
- **Performance Evaluation**
  - SQRT-C outperforms RESTful services in terms of response time (Figure 2) and jitter for real-time applications

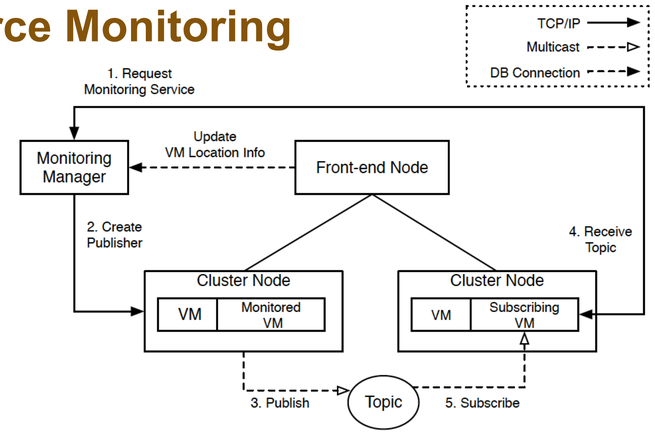

Figure 1. SQRT-C System Architecture

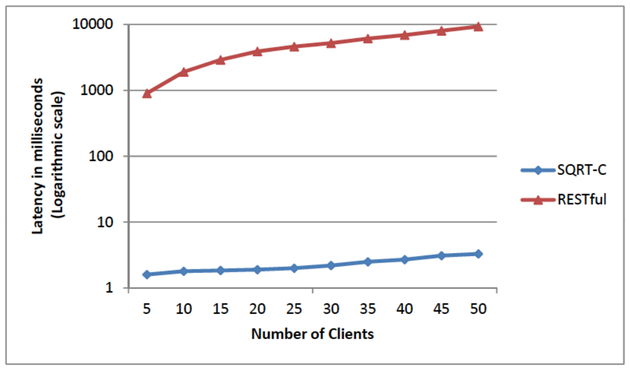

Figure 2. Average Message Latency Comparison of SQRT-C and RESTful by Number of VMs

4. Ongoing and Future Work

- **Unresolved challenges**
  - Trade-off between timeliness and high-availability with strong consistency
  - Tradeoffs between response time and consistency
    - BASE (Basically Available replicated Soft state with Eventual consistency)
    - ACID (atomicity, consistency, isolation, and durability) database models
- **Redundancy-based fault recovery mechanisms for DRE systems**
  - Replication using primary-backup
  - A proactive, resource-aware fail-over strategy
  - A resource-aware allocation based on backup resource overbooking
- **Our proposed research**
  - Implementation of a fault-tolerant cloud architecture applying redundancy-based fault recovery mechanisms
  - Performance analysis for trade-off between strict timeliness and strong consistency
  - Integration of real-time hypervisors and deadline-aware data-center networks
- **Work in progress**
  - A framework for automated placement of virtual machine replicas for DRE systems
  - Bin-packing heuristics developed

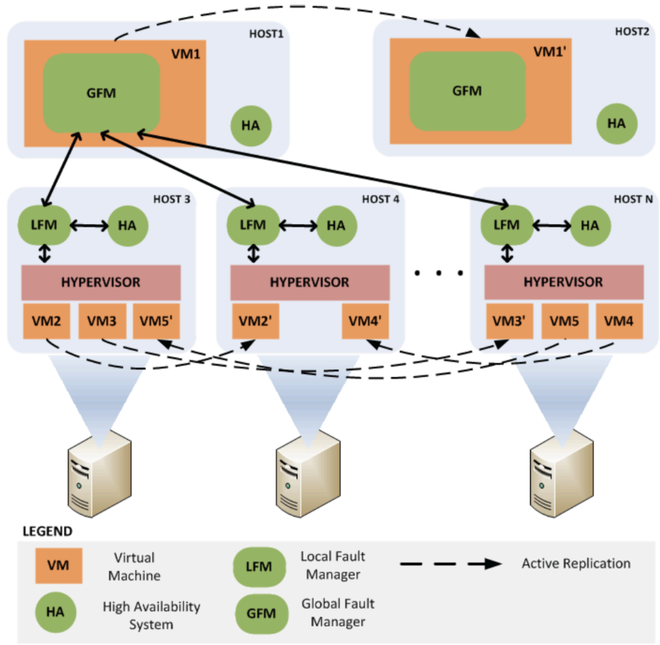

Figure 3. Fault-tolerant Cloud Architecture

5. References

[1] K. Keahey, I. Foster, T. Freeman, and X. Zhang, "Virtual workspaces: Achieving quality of service and quality of life in the grid," Scientific Programming, vol. 13, no. 4, pp. 265–275, 2005.
[2] C. Wilson, H. Ballani, T. Karagiannis, and A. Rowstron, "Better never than late: Meeting deadlines in datacenter networks," in Proceedings of the ACM SIGCOMM 2011 conference on SIGCOMM. ACM, 2011, pp. 50–61.
[3] S. Xi, J. Wilson, C. Lu, and C. Gill, "Rt-xen: towards real-time hypervisor scheduling in Xen," in Proceedings of the ninth ACM international conference on Embedded software. ACM, 2011, pp. 39–48.
[4] M. Lee, A. Krishnakumar, P. Krishnan, N. Singh, and S. Yajnik, "Supporting soft real-time tasks in the Xen hypervisor," in ACM SIGPLAN Notices, vol. 45, no. 7. ACM, 2010, pp. 97–108.
[5] B. Cully, G. Lefebvre, D. Meyer, M. Feeley, N. Hutchinson, and A. Warfield, "Remus: High availability via asynchronous virtual machine replication," in Proceedings of the 5th USENIX Symposium on Networked Systems Design and Implementation. USENIX Association, 2008, pp. 161–174.
[6] K. Birman, D. Freedman, Q. Huang, and P. Dowell, "Overcoming cap with consistent soft-state replication," Computer, no. 99, pp. 1–1, 2011.
[7] J. Balasubramanian, S. Tambe, C. Lu, A. Gokhale, C. Gill, and D. Schmidt, "Adaptive failover for real time middleware with passive replication," in Real-Time and Embedded Technology and Applications Symposium, 2009. RTAS 2009. 15th IEEE. IEEE, 2009, pp. 118–127.
[8] J. Balasubramanian, A. Gokhale, A. Dubey, F. Wolf, D. Schmidt, C. Lu, and C. Gill, "Middleware for resource-aware deployment and configuration of fault-tolerant real-time systems," in Real-Time and Embedded Technology and Applications Symposium (RTAS), 2010 16th IEEE

Research supported in part by NSF CNS 0915976