An Analysis of Virtual Machine Live Migrations in the Wild

Mukil Kesavan

Advisors: Dr. Karsten Schwan & Dr. Ada Gavrilovska

Live Migration Use Cases

- Fault Tolerance
- Repair/Maintenance/Upgrade
- Dynamic Resource Allocation (DRA)

➢ Context: We have been performing Large-Scale DRA in a 750 Server Datacenter for over a year now.

Measurement Environments

- Infrastructure 1: GT Techway Datacenter
  - 256 Servers, VMware vSphere 4 (DRS + CCM)
  - Workload: Datacenter Trace + Applications
  - Data Timeline: 1 Month
- Infrastructure 2: CMU vCloud Cluster
  - 15 Servers, VMware vCloud (DRS)
  - Workload: Academic Courses + Research Projects
  - Data Timeline: 3 Months
  - Newer hardware, VMs with Larger Memory

Migration Times

GT

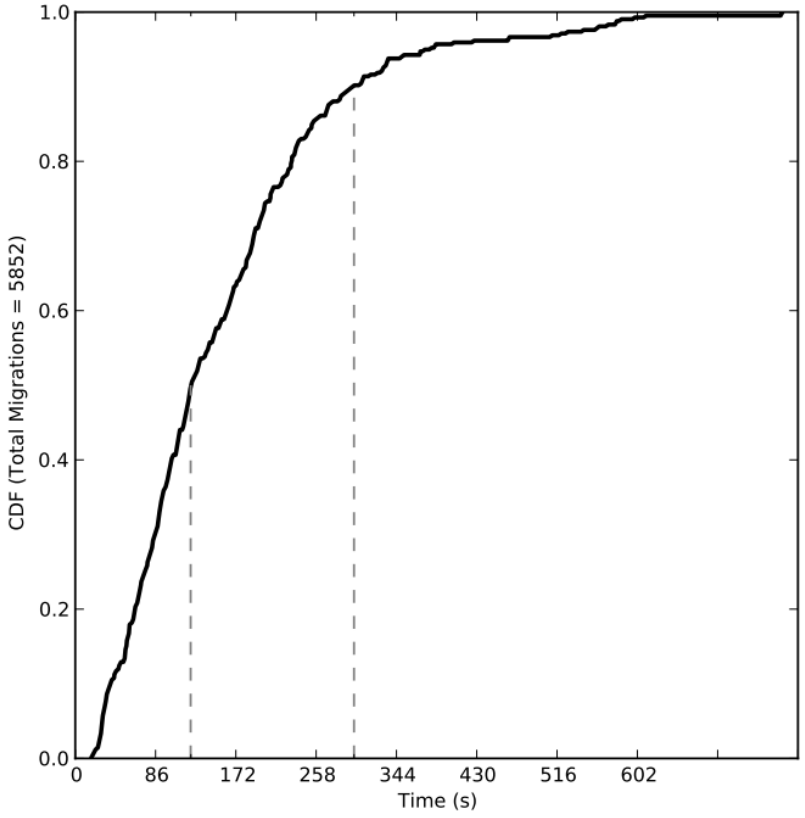

50%ile = 110s, 90%ile = 287s

CMU

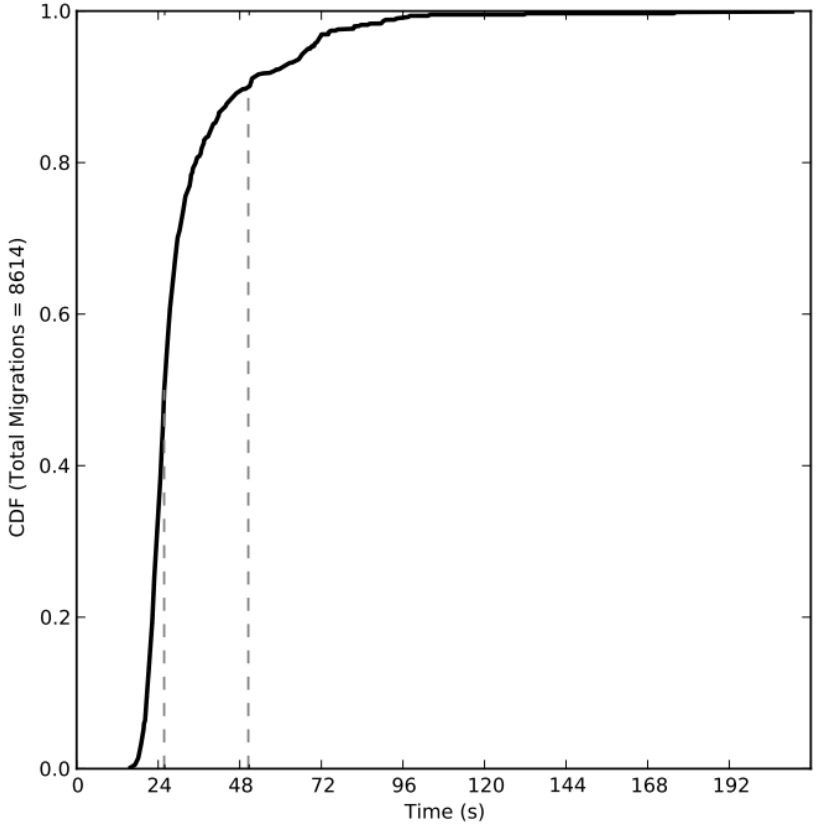

50%ile = 25s, 90%ile = 49s

* 56 migrations > 16 minutes!

Migration Failure Rates

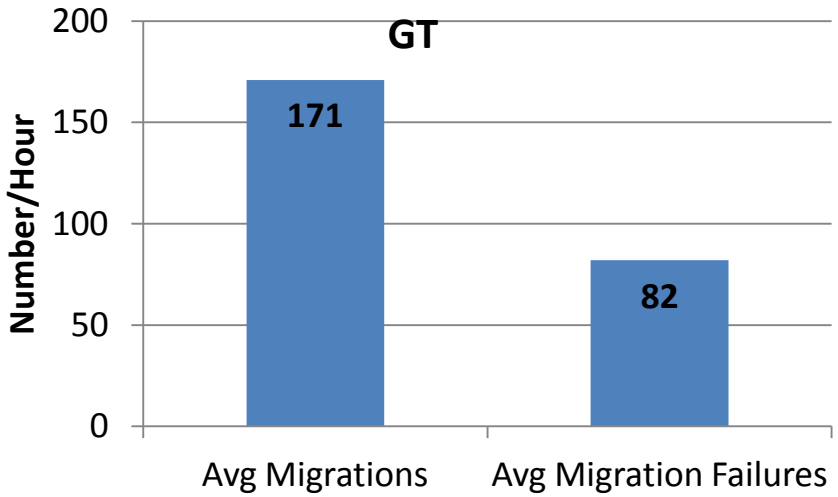

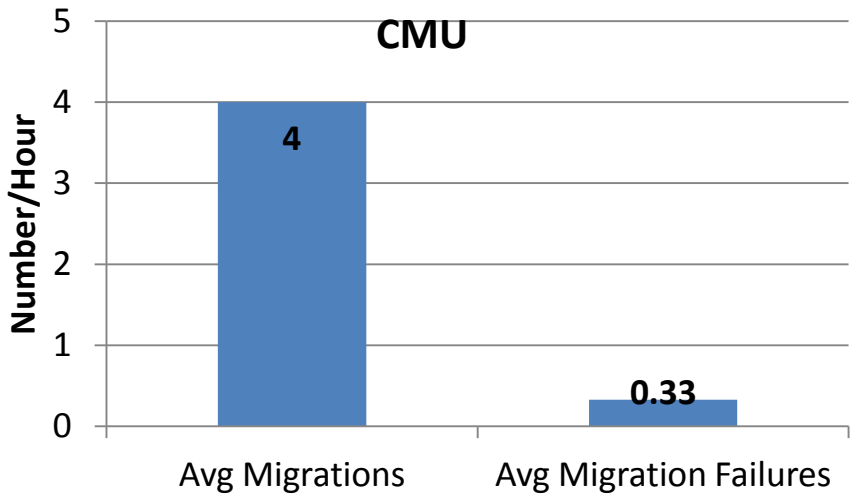

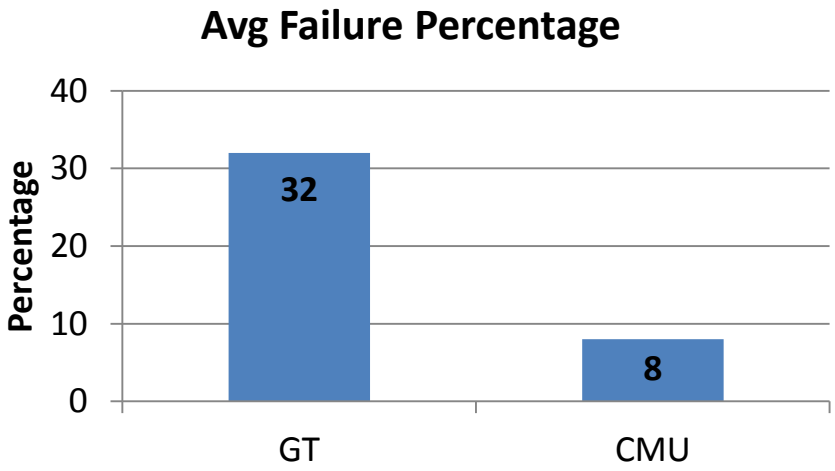

Why do these failures happen?

- Software Timeouts
- S/W Misconfig
- Software Bugs
- N/W Connectivity
- H/W Errors/Failures
- …

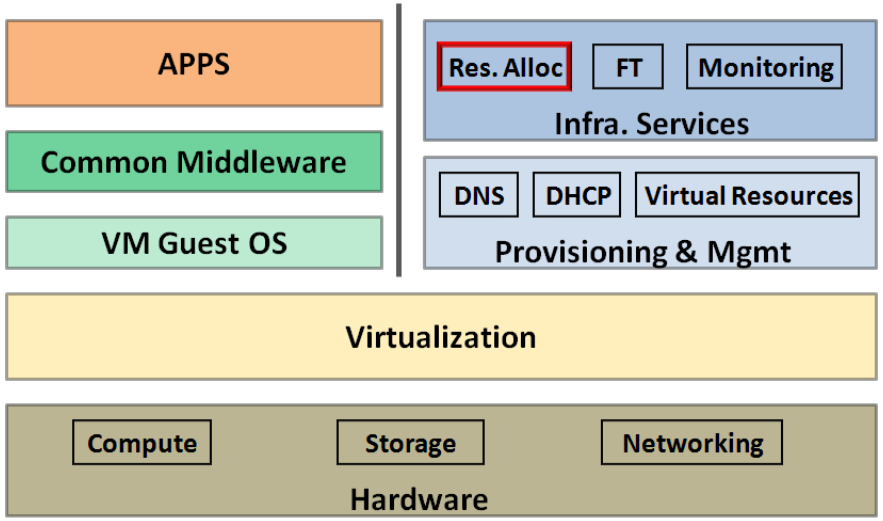

Datacenter Software Stack

Statistical Analysis of Failures

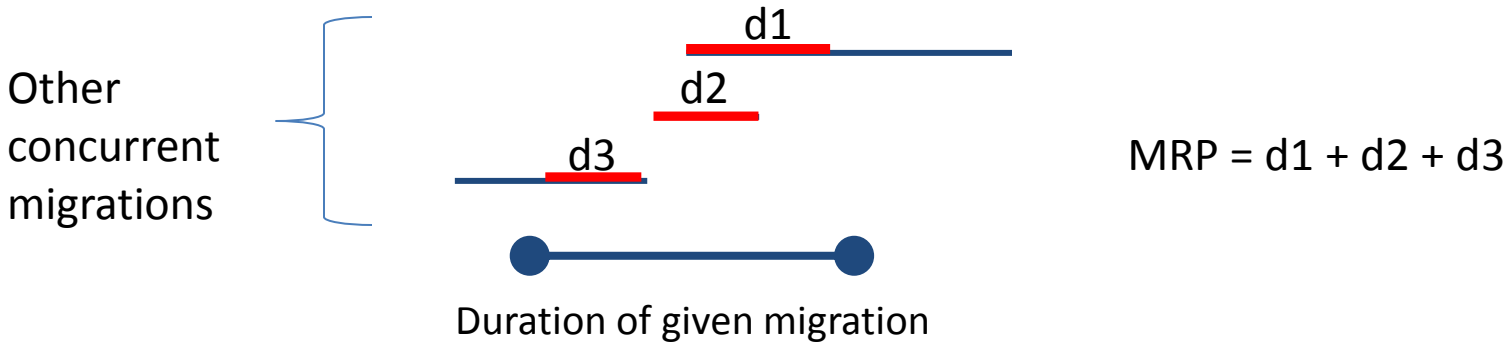

- **Management Resource Pressure (MRP):** Relationship between too many intensive concurrent migrations and their outcomes (success/failure)
- Migration duration α Dirty Memory (Res. Usage)

MRP vs. Migration Outcomes (GT)

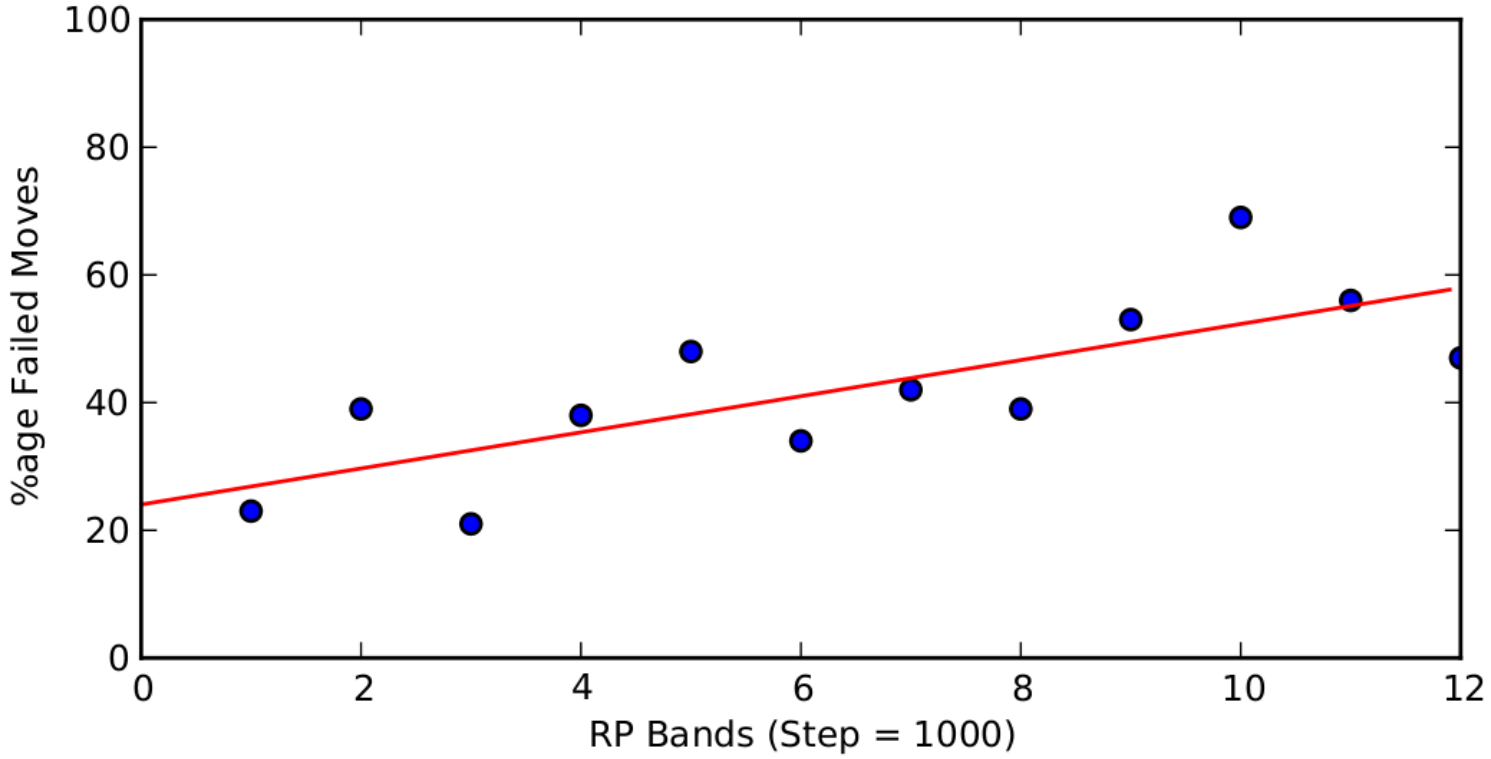

- Correlation Coefficient (R) = 0.8 – Strong Positive
- Goodness of Fit (R^2) = 0.6

Takeaways

- Large fraction of observed migration errors are related to resource insufficiency for mgmt workload.
- Mgmt workload cost varies widely.
- Need intelligent throttling of variable-cost mgmt workload to *avoid failures in addn to cost optimization*.