

Seagull: Intelligent Cloud Bursting For Enterprise Applications

Tian Guo, Upendra Sharma, Timothy Wood[‡],
Sambit Sahu[†], Prashant Shenoy

University of Massachusetts Amherst,
The George Washington University[‡], IBM Research[†]



Cloud Computing


■ Cloud Computing:

- Pay-as-you-go service
- Rent Resources
- Infrastructure as a Service
 - Virtualization technology, rent VMs
 - Popular for Apps with dynamic workload



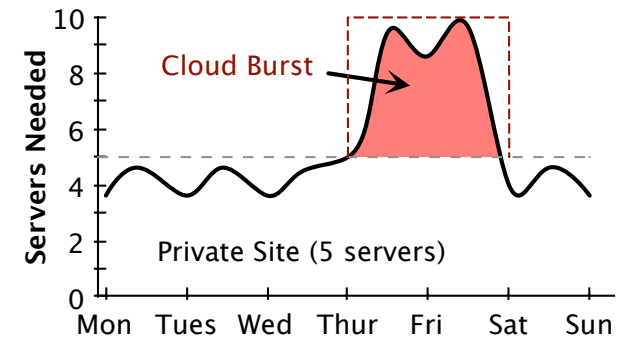
■ Benefits

- Flexible pricing model
- Agile to workload changes

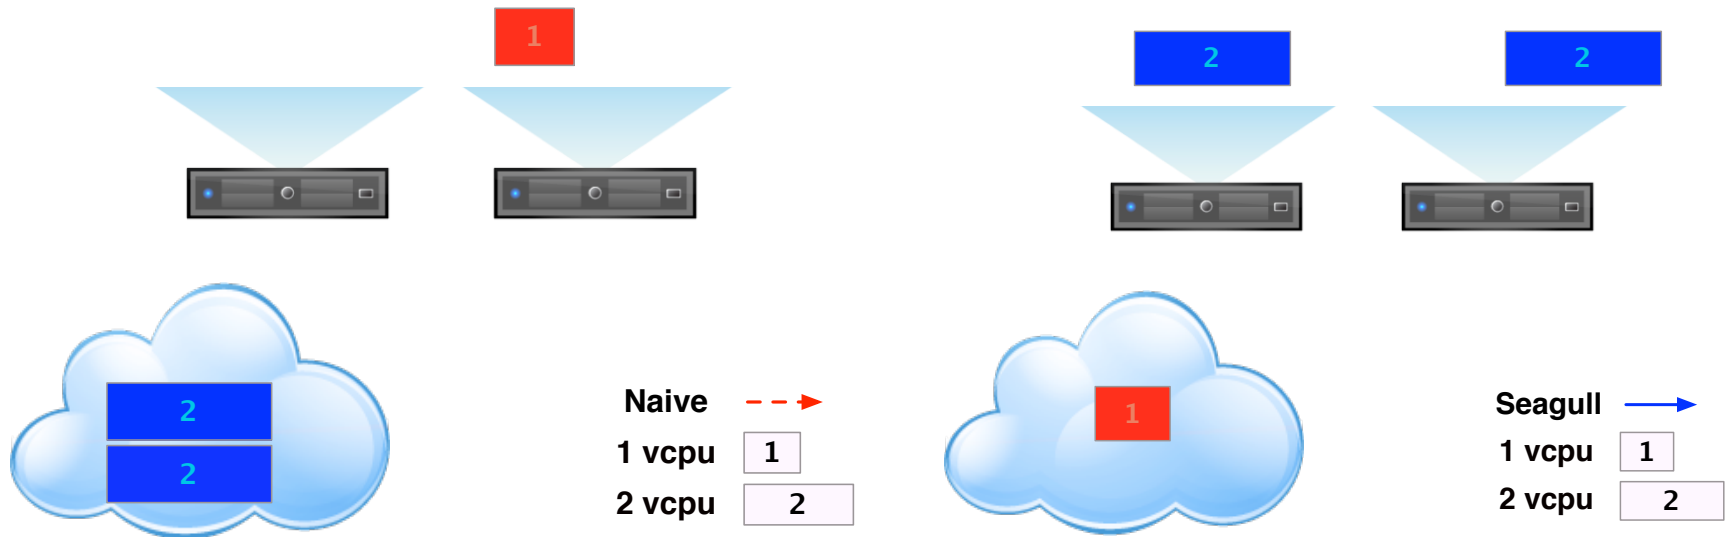


Cloud Bursting

- Enterprises own private data centers
 - Try to use the existing infrastructure (hybrid)
- Cloud Bursting
 - Enables Enterprise to use local data center
 - rents public resource upon workload changes
 - seamless and transparent resource sharing between local and public cloud
- Challenges
 - When to trigger cloud bursting?
 - Which Apps to cloud Burst?
 - How to balance cost and time trade-off?
- Seagull
 - Cloud Bursting Algorithm
 - Precopying Algorithm



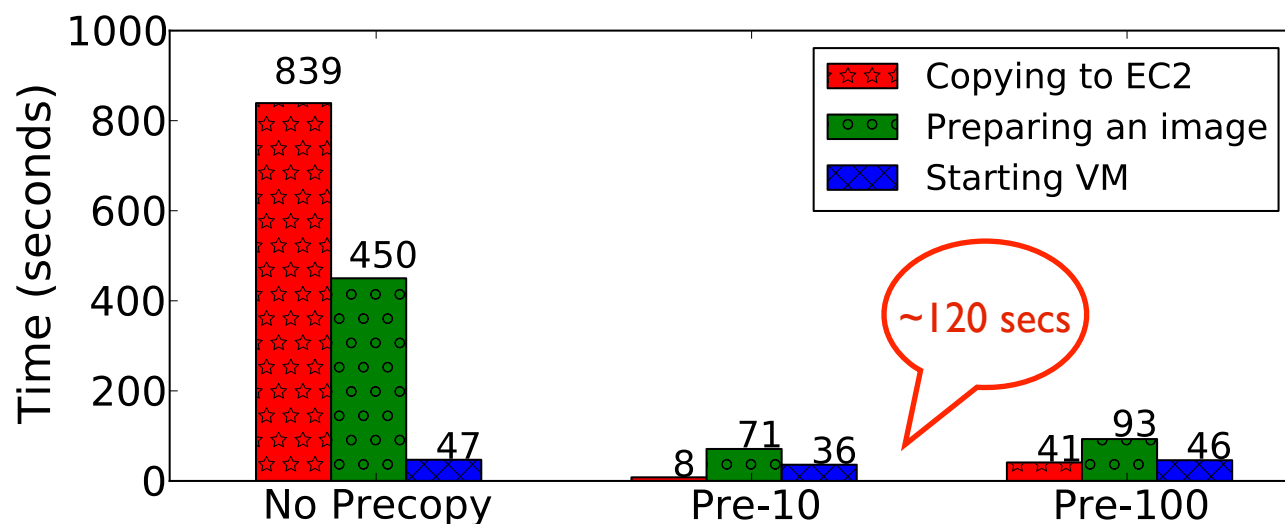
Seagull Cloud Bursting Algorithm



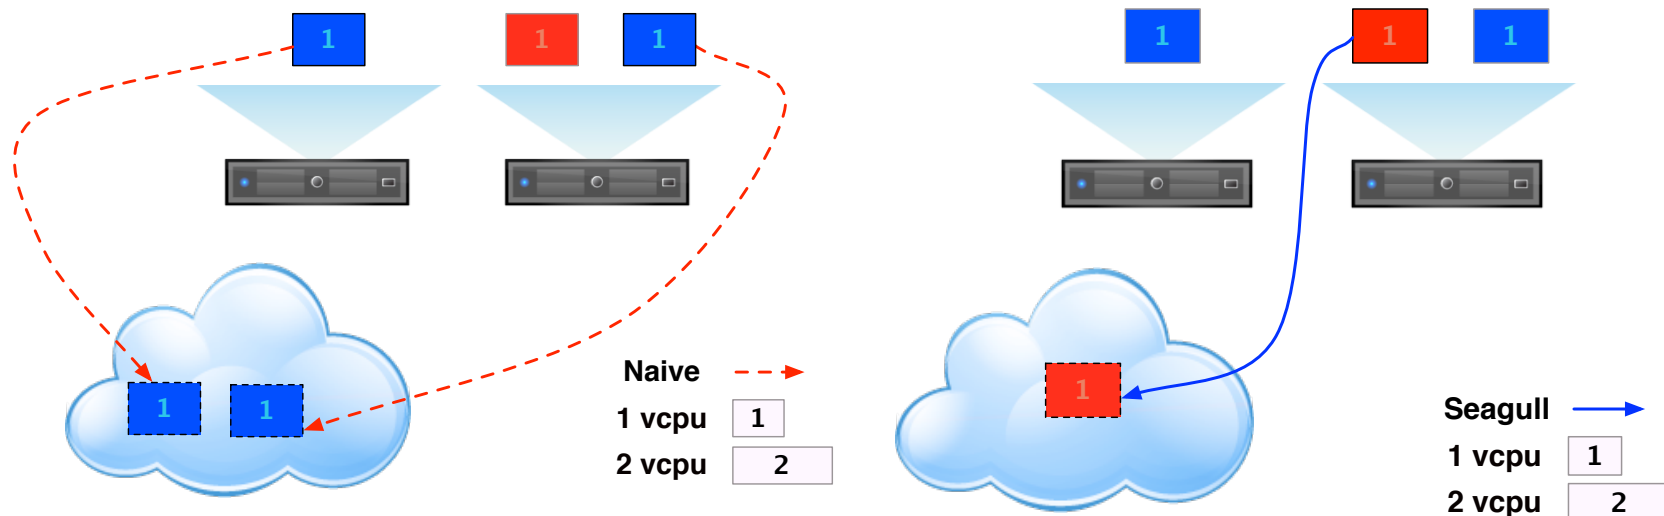
- Which applications to cloud burst?
 - Naive approach: move overloaded applications
 - Incurs high **cost and overhead**
 - Seagull approach: Pick the cheapest applications
 - Multi-resource bin packing problem
 - Greedy approach
 - Metric: $\text{App_Costs} / \text{VM_cores}$ to run in public cloud

How to Lower Migration Time ?

- Cloud bursting on demand
 - e.g 5 GB disk state, takes a long time (~22 mins)
- Opportunistic Precopying
 - Copys app vm state to the public cloud in the background
 - Benefit: Dramatically shortens the migration time
 - Some experiments:



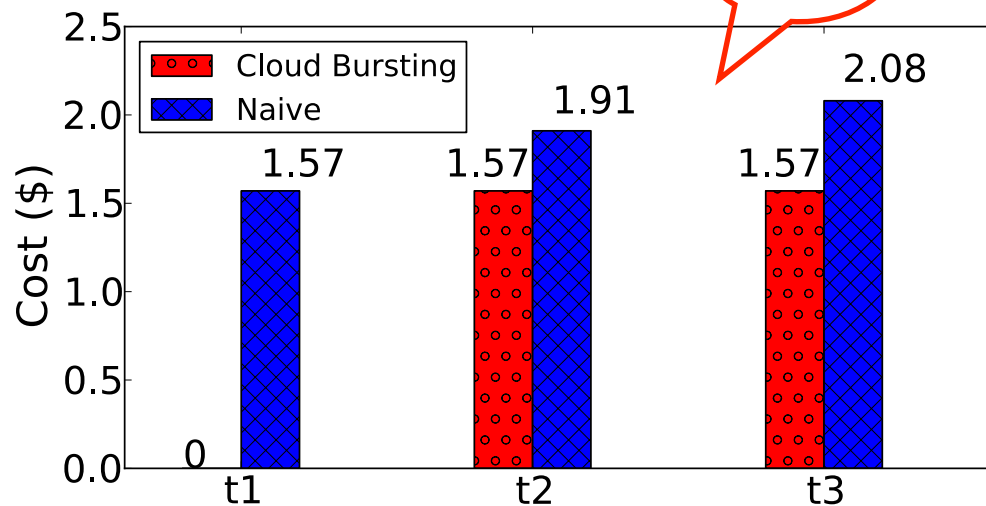

Seagull Precopying Algorithm



- How to balance cost and time trade-off?
 - Naive Precopying: Precopying overloaded applications
 - Not Necessary lower migration time
 - Intelligent Precopying
 - Intuition: Choose the apps that are most likely to be migrated

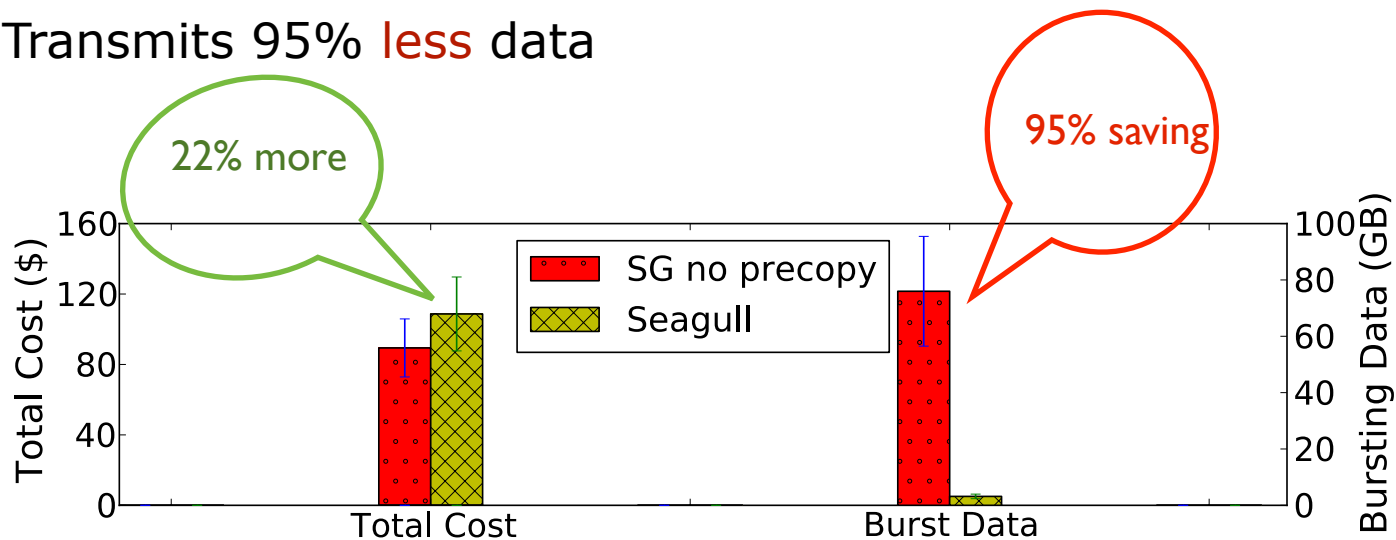
Cloud Bursting Algorithm Evaluation

- Experiment Setup
 - 3 hosts and 5 Apps
 - Varying workload of A for 4 hours
- Seagull is cost Efficient
 - Lowers cost by 25% over 4 hours



Precopying Algorithm Evaluation

- Experiment Setup
 - Emulation with 200 quad-core hosts
 - 40 applications, 30% were overloaded
 - Precopying frequency: 1 hr & total time: 24 hrs
- Seagull balances time and cost well
 - Spends 22% **more** money
 - Transmits 95% **less** data



Summary

- Cloud Bursting
 - Hybrid solution for dynamic workload
 - Good for Enterprises with private data centers
- Seagull: Intelligent and automated Cloud Bursting
 - Determines which Apps to Cloud Burst
 - Lowers Cost by 25%
 - Determines which Apps to Precopy
 - Saves 95% Data Transmission