# Towards Intelligent Agent-Based Energy-Aware Workload Consolidation in Clouds

Eugen Feller, Louis Rilling, Christine Morin

INRIA Myriads research team, Rennes, France
Green Days @ Paris, 31 May - 1 June, Paris, France

June 1, 2011

INRIA RENNES KERLABS

# Motivation

## Motivation

- Cloud computing paradigm is becoming ubiquitous
  - Increasing numbers of energy hungry data centers
  - Energy conservation $\rightarrow$ Decrease operation costs $\rightarrow$ Increase system reliability
- Some of the challenges
  - Dynamic energy and QoS-aware workload (e.g., VM) consolidation
- Answer requires
  - Models for performance, energy and migration costs
  - Consolidation algorithms which consider multiple objectives (i.e., Energy and QoS) and resources (e.g., CPU, RAM, I/O)

# Workload Consolidation

## Overview

- Workload (i.e., VMs) is typically load balanced across the cluster
  - Servers are rarely fully idle
  - Resource demand is usually of a bursty nature $\rightarrow$ Low average utilization of approximately 15-20%
- *Workload consolidation*
  - Aims at *unbalancing* the workload $\rightarrow$ Makes use of workload migration
  - Creates the necessary idle-times required to turn off the over-provisioned server (i.e., save energy)

# Workload Consolidation

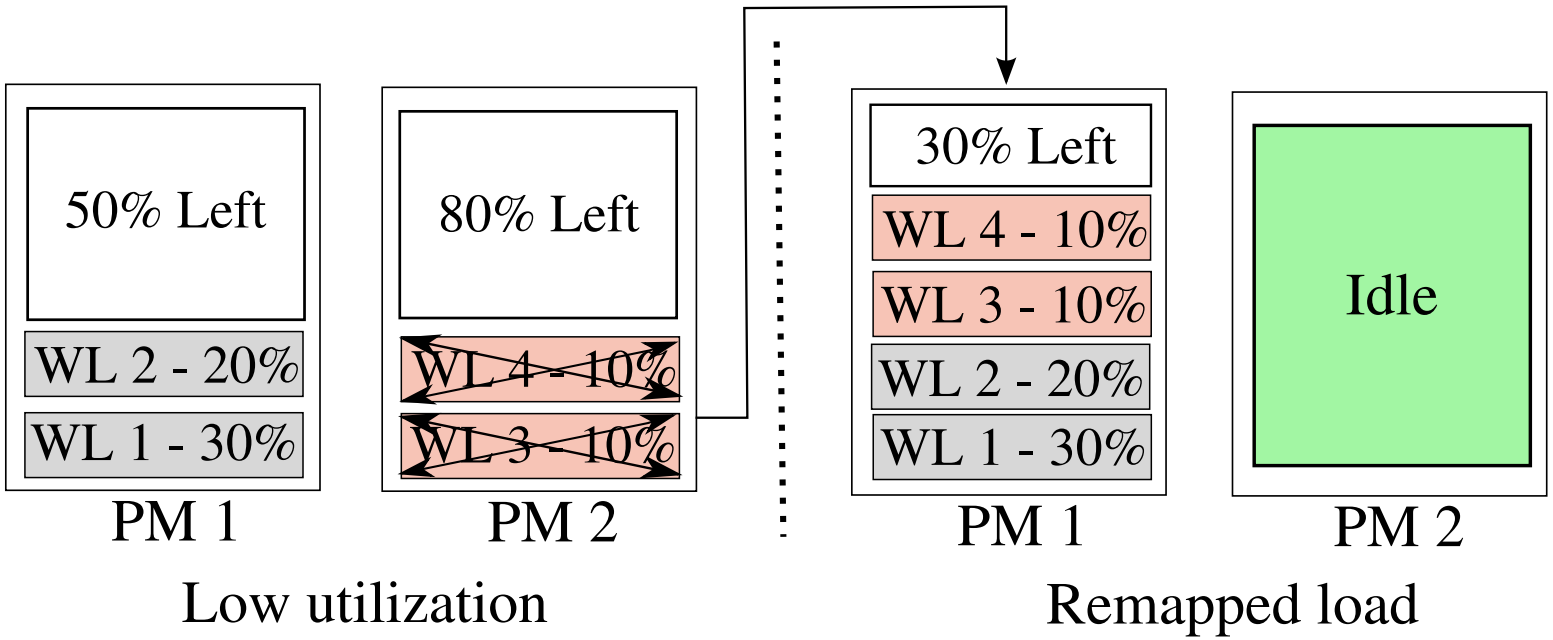


Figure: Use Case

# Research Challenges

## Research Challenges

- Traditional workload consolidation heuristics
  - Tend to waste alot of resources (e.g., First-Fit Decreasing)
  - Bad scalability and fault-tolerance properties
- *Our approach*
  - Definition of the workload consolidation problem $\rightarrow$ Instance of the multi-dimensional bin-packing problem (MDBPP)
  - *Nature-inspired* workload consolidation algorithm based on the *Ant Colony Optimization (ACO)*
- More information:
  - "**Energy-Aware Ant Colony Based Workload Placement in Clouds**", *RR-7622*, Rennes, France, May 2011