

# Green Clouds: Power Consumption as a First Order Criterion?



ICAC Panel – Jun 17<sup>th</sup> 2009



# Power consumption is expensive

- Society increasingly depends on IT infrastructure
- Growth in data center deployments [IDC 2006]
  - Total power consumption doubled over 2000-2005:
    - Almost all growth due to larger # of inexpensive servers
  - Cost of power+cooling:
    - 1996: 16% of new server spend; 2010: 70%
  - Top concerns, IT executives: security, power, performance
- Cloud infrastructures
  - Demand for larger, distributed data centers



# Environmental responsibility

- Every 2kW of power consumed by servers requires 1kW to cool emitted heat [IDC 2006]
  - 1.2% of total US electricity in 2005 [Kooimey 2007]
  - Carbon footprint
- Environmental impact beyond power consumption
  - Manufacturing, distribution, disposal



# Implications

- Performance no longer only criterion for optimization and management
- Power (servers and cooling) is a key resource:
  - Need to design components, platforms, data centers, cloud infrastructures for power-efficiency
  - At run-time, need to factor in power consumption when managing deployed resources
- More expensive to manage than purchase servers
- Management – including power – cuts across various layers in the software/hardware stack



# Meters and knobs across the stack

- Platforms: multi-core spare processing for monitoring/decision processing; modern management technologies
- Virtualization: utilization monitoring/aggregation, workload placement/migration, resource control
- Cloud computing: consolidation; economic model; geographical distribution; replication; new programming and execution models



# We like self-management

- Where in the stack is there most potential benefit from autonomic approaches? least benefit?
- What sensors and managed control points at different stacks are key to enable effective power management?
- What models/paradigms for autonomic power management are proven, or promising?
- What are the challenges and opportunities for adaptive power management across geographically distributed data center sites?



# Panelists

- Jeff Chase - Duke University
- Jeff Kephart - IBM Research
- Hartmut Schmeck - Universitat Karlsruhe
- Vanish Talwar - HP Labs
- Eno Thereska - Microsoft Research



# Panel organizers

- Karsten Schwan,  
Georgia Tech



- Renato Figueiredo,  
University of Florida





# Format

- Each panelist: short (10-minute) position statement
- Panel discussion round after all presentations
- Short/clarification questions after each presentation
  - Save your discussion questions to the end





Green Clouds Panel, Barcelona, Spain, June 17 2009

