Energy Efficiency in Virtualized Systems

Segoy Thiel

MIN Faculty
Department of Informatics
University of Hamburg

10. December 2014

Outline

Outline

Outline

Outline

Motivation

- Data Centers consume large amounts of energy
  - 56% rise from 2005-2010
  - 2010 about 1.3% of total energy use
  - estimated to 2% of CO_2 emissions
- Virtualization allows to create several Virtaul Machines (VMs) on a physical server to reduce amount of hardware in use
- VMs can be moved between physical hosts
- Resource usage can be adjusted to current requirements

Focus

- Live migration of VMs to dynamically reallocate VMs to current resource reqiurements
- Idle physical nodes can be switched off
- Provide reliable Quality of Service (QoS) defined in Service Level Agreements (SLA) characteristics:
  - maximal throughput
  - minimal response time
    or latency delivered by deployed system
- Thermal Optimization - temperature of physical nodes is considered in reallocation decisions
  - reducing workload of overheated nodes
  - decrease cooling system load

Focus

- Live migration of VMs to dynamically reallocate VMs to current resource reqiurements
- Idle physical nodes can be switched off
- Provide reliable Quality of Service (QoS) defined in Service Level Agreements (SLA) characteristics:
  - maximal throughput
  - minimal response time
    or latency delivered by deployed system
- Thermal Optimization - temperature of physical nodes is considered in reallocation decisions
  - reducing workload of overheated nodes
  - decrease cooling system load

Important requirements for a Cloud computing environment is providing reliable QoS. Defined in terms of Service Level Agreements (SLA) that describe characteristics such as minimal throughput, maximal response time or latency delivered by the deployed system. modern virtualization technologies can ensure performance isolation between VMs sharing the same physical node, due to aggressive consolidation and variability of the workload some VMs may not get the required amount of resource when requested. This leads to performance loss in terms of increased response time, time outs or failures. Cloud providers have to deal with energy-performance trade-off – minimization of energy consumption, while meeting QoS requirements.

Challenges

How to:

1. Solve trade-off: energy savings vs. delivered performance
2. Determine when which and where to migrate VMs in order to:
   - minimize energy consumption
   - minimize migration overhead
   - ensuring SLA
3. Develop efficient decentralized and scalable algorithms for resource allocation
4. Develop comprehensive solution by combining several allocation policies with different objectives

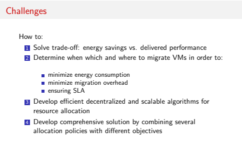

1. How to optimally solve the trade-off between energy savings and delivered performance?
2. How to determine when, which VMs, and where to migrate in order to minimize energy consumption by the system, while minimizing migration overhead and ensuring SLA?
3. How to develop efficient decentralized and scalable algorithms for resource allocation?
4. How to develop comprehensive solution by combining several allocation policies with different objectives?

Infrastructure for large-scale Cloud data center

- Comprises physical nodes charackterized by:
  - CPU: perfomance defined by MIPS
  - RAM
  - network bandwidth
- Users submit requests for provisioning VMs requiring resources from nodes
- SLA violation occurs when a VM cannot get requested amount of resources
- Tiered comprising: global and local managers

Infrastructure for large-scale Cloud data center

- Comprises physical nodes charackterized by:
  - CPU: perfomance defined by MIPS
  - RAM
  - network bandwidth
- Users submit requests for provisioning VMs requiring resources from nodes
- SLA violation occurs when a VM cannot get requested amount of resources
- Tiered comprising: global and local managers

Each node has a CPU, which can be multicore, with performance defined in Millions Instructions Per Second (MIPS). Besides that, a node is characterized by the amount of RAM and network bandwidth. Users submit requests for provisioning of m heterogeneous VMs with resource requirements defined in MIPS, amount of RAM and network bandwidth. SLA violation occurs when a VM cannot get the requested amount of resource, which may happen due to VM consolidation. The software system architecture is tiered comprising, global and local managers.

Local and Global Managers

Local managers:

- Reside on nodes as part of VMM
- Choose VMs to be migrated in cases:
  - utilization of resource is close to 100% - risk of SLA Violation
  - utilization is low - all VMs reallocated to another node
  - VMs have intensive network communication allocated to different hosts
  - temp exceeds limit - VMs are migrated to cool node down
- Send information and VMs to migrate to global manager
- Issue VM resizing, applicate DVFS, turn on/off idle nodes

Global manager:

- Attached to set of nodes and process data from local managers

Local and Global Managers

Local managers:

- Reside on nodes as part of VMM
- Choose VMs to be migrated in cases:
  - utilization of resource is close to 100% - risk of SLA Violation
  - utilization is low - all VMs reallocated to another node
  - VMs have intensive network communication allocated to different hosts
  - temp exceeds limit - VMs are migrated to cool node down
- Send information and VMs to migrate to global manager
- Issue VM resizing, applicate DVFS, turn on/off idle nodes

Global manager:

- Attached to set of nodes and process data from local managers

Local manager choose VMs to be migrated when:
-The utilization of some resource is close to 100-The utilization is low, all the VMs should be reallocated and the idle node should be turned off.
-A VM has intensive network communication with another VM allocated to a different physical host.
-The temperature exceeds a limit and VMs have to be migrated to reduce load on the cooling system and allow the node to cool down naturally.
Local managers send to the global managers the information about the utilization of resources and VMs chosen to migrate. They issue commands for VM resizing, application of DVFS and turning on/off idle nodes. The Global manager is attached to a set of nodes and processes data from the local managers. The global managers apply a heuristic for semi-online multidimensional bin-packing, where bins represent physical nodes and items are VMs that have to be allocated. The decentralization removes a Single Point of Failure and improves scalability.

The System Model

http://beloglazov.info/thesis.pdf Figure 4.1

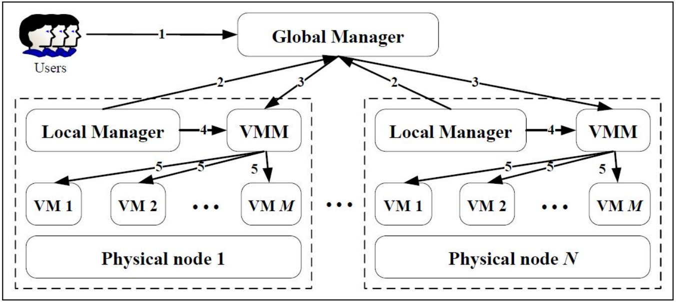

The System Model

http://beloglazov.info/thesis.pdf Figure 4.1

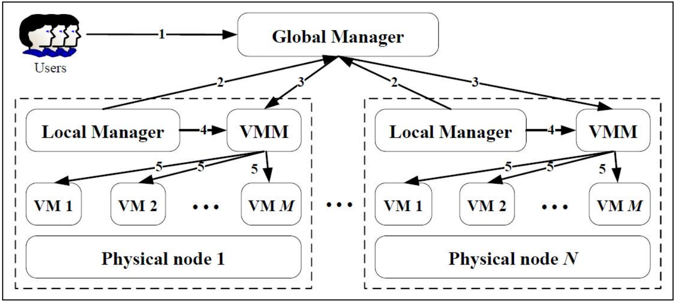

1. New requests for VM provisioning

The System Model

http://beloglazov.info/thesis.pdf Figure 4.1

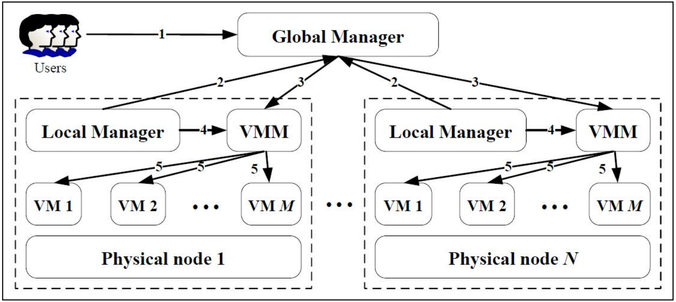

1. New requests for VM provisioning
2. Data about utilization of resources and VMs chosen to migrate

The System Model

http://beloglazov.info/thesis.pdf Figure 4.1

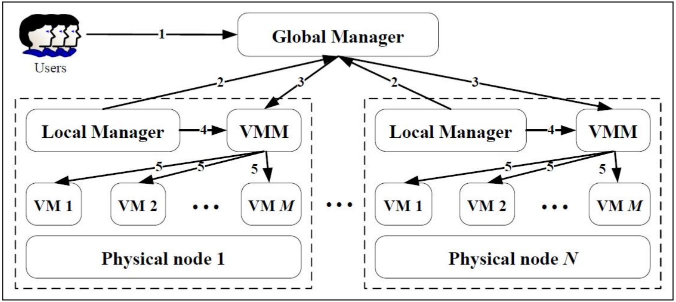

1. New requests for VM provisioning
2. Data about utilization of resources and VMs chosen to migrate
3. Migration commands

The System Model

http://beloglazov.info/thesis.pdf Figure 4.1

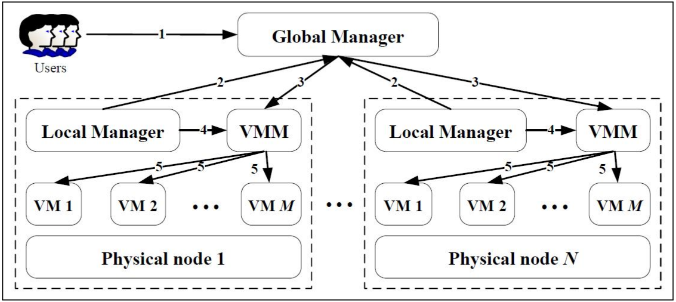

1. New requests for VM provisioning
2. Data about utilization of resources and VMs chosen to migrate
3. Migration commands
4. Commands for VM resizing and adjusting of power states

The System Model

http://beloglazov.info/thesis.pdf Figure 4.1

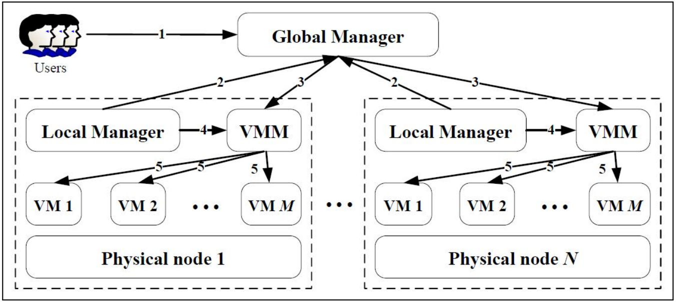

1. New requests for VM provisioning
2. Data about utilization of resources and VMs chosen to migrate
3. Migration commands
4. Commands for VM resizing and adjusting of power states
5. VM resizing, scheduling and migration actions

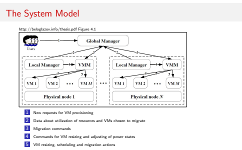

1. New requests for VM provisioning. Users submit requests for provisioning of VMs.
2. Data about utilization of resources and VMs chosen to migrate. The local managers propagate information about resource utilization and VMs chosen to migrate to the global manager.
3. Migration commands. The global manager issues VM migration commands in order to optimize current allocation.
4. Commands for VM resizing and adjusting of power states. The local managers monitor their host nodes and issue commands for VM resizing and changes in power states of nodes.
5. VM resizing, scheduling and migration actions. According to the received commands, VMM performs actual resizing and migration of VMs as well as resource scheduling.

Three stages of VM placement optimization

1. Reallocation according current utilization of multiple system resources
   - Main idea: set upper and lower utilization thresholds and keep utilization of CPU created by VMs sharing node inbetween
2. Optimization of virtual network topologies established between VMs
   - place communicating VMs in a way to minimize overhead of data transfer over network
3. VM reallocation considering thermal state of resources
   - monitor nodes thermal states, reallocate workload from overheated nodes to allow natural cooling

Three stages of VM placement optimization

1. Reallocation according current utilization of multiple system resources
   - Main idea: set upper and lower utilization thresholds and keep utilization of CPU created by VMs sharing node inbetween
2. Optimization of virtual network topologies established between VMs
   - place communicating VMs in a way to minimize overhead of data transfer over network
3. VM reallocation considering thermal state of resources
   - monitor nodes thermal states, reallocate workload from overheated nodes to allow natural cooling

Three stages of VM placement optimization: reallocation according to current utilization of multiple system resources:
-The main idea is to set upper and lower utilization thresholds and keep total utilization of CPU created by VMs sharing the same node between these thresholds. If the utilization exceeds the upper thresholds, some VMs have to be migrated to reduce the risk of SLA violation. If the utilization goes bellow the lower thresholds, all VMs have to be migrated and the node has to be switched off to save the energy.
Optimization of virtual network topologies established between VMs:
-It is crucial to consider network communication behavior of VMs in reallocation decisions. The aim is to place communicating VMs in a way minimizing the overhead of data transfer over network.
VM reallocation considering thermal state of the resources:
-A cooling system of a data center consumes a significant amount of energy, therefore this optimization stage is aimed at optimization of cooling system operation.

Allocation Policies

Algorithms have to meet following requirements:

- Decentralization and parallelism
- High performance
- Guaranteed QoS
- Independence of the workload type

VM reallocation problem can be devided in two:

1. Selection of VMs to migrate
   - has to be considered separately for each stage
2. determining new placement of these on physical hosts
   - solved by heuristic for multidimensional bin-packing

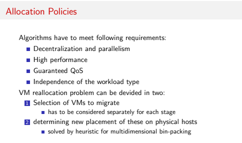

Decentralization and parallelism – to eliminate SPF and provide scalability.
-High performance – the system has to be able to quickly respond to changes in the workload.
-Guaranteed QoS – the algorithms have to provide reliable QoS by meeting SLA.
-Independence of the workload type – the algorithms have to be able to perform efficiently in mixedapplication environments.

The VM reallocation problem can be divided in two: selection of VMs to migrate and determining new placement of these VMs on physical hosts. The first part has to be considered separately for each optimization stage. The second part is solved by application of a heuristic for semionline multidimensional bin-packing problem.

Example Simulation of a Data Center

- Data Center consists of 100 heterogeneous physical nodes
- Each node have one CPU core, perfomance:
  1000, 2000 or 3000 MIPS, 8Gb RAM, 1TB storage
- Users submit requests for provisioning 290 heterogenous VMs that fill full capacity of Data Center
- Simulated policies:
  - Non Power Aware (NPA)
  - DVFS
  - Single Threshold(ST)
  - Two-threshold policy aimed at Minimazation of Migrations (MM)

Example Simulation Results

Policy	Energy	SLA	Migr.	Avg. SLA
NPA	9.15 KWh	-	-	-
DVFS	4.40 KWh	-	-	-
ST 50%	2.03 KWh	5.41%	35 226	81%
ST 60%	1.50 KWh	9.04%	34 231	89%
MM 30-70%	1.48 KWh	1.11%	3 359	56%
MM 40-80%	1.27 KWh	2.75%	3 241	65%
MM 50-90%	1.14 KWh	6.69%	3 120	76%

http://seelab.ucsd.edu/virtualefficiency/related_papers/26_4039a826.pdf

Example Simulation Results

Policy	Energy	SLA	Migr.	Avg. SLA
NPA	9.15 KWh	-	-	-
DVFS	4.40 KWh	-	-	-
ST 50%	2.03 KWh	5.41%	35 226	81%
ST 60%	1.50 KWh	9.04%	34 231	89%
MM 30-70%	1.48 KWh	1.11%	3 359	56%
MM 40-80%	1.27 KWh	2.75%	3 241	65%
MM 50-90%	1.14 KWh	6.69%	3 120	76%

http://seelab.ucsd.edu/virtualefficiency/related_papers/26_4039a826.pdf

The results show that dynamic reallocation of VMs according to current utilization of CPU can bring higher energy savings comparing to static allocation policies. MM policy allows to achieve the best energy savings: by 83%, 66% and 23% less energy consumption relatively to NPA, DVFS and ST policies with thresholds 30-70ensuring percentage of SLA violations of 1.1%; and by 87%, 74% and 43% with thresholds 50-90% and 6.7% of SLA violations. MM policy leads to more than 10 times fewer VM migrations than ST. The flexibility of the algorithm, as the thresholds can be adjusted according to SLA requirements. Strict SLA (1.11%) allow achievement of the energy consumption of 1.48 KWh. If SLA are relaxed (6.69%), the energy consumption is further reduced to 1.14 KWh.

Cloud Computing & Virtualiztion

Cloud Computing leverage virtualization

- Provides ability to provision resources on demand
- Organizations can outsource computation needs
- Eliminate necessity to maintain computing infrastructure

Cloud Computing leads to energy efficiency providing following charakteristics:

- Economy of scale and elimination of redundancies
- Increased utilization of computing resources
- Location independence
- Scaling up/down and in/out

Cloud Computing & Virtualiztion

Cloud Computing leverage virtualization

- Provides ability to provision resources on demand
- Organizations can outsource computation needs
- Eliminate necessity to maintain computing infrastructure

Cloud Computing leads to energy efficiency providing following charakteristics:

- Economy of scale and elimination of redundancies
- Increased utilization of computing resources
- Location independence
- Scaling up/down and in/out

Cloud Computing leverages virtualization and provides the ability to provision resources on-demand on a pay-as-you-go basis. Organizations can outsource their computation needs to the Cloud, thereby eliminating the necessity to maintain their own computing infrastructure.
Cloud computing naturally leads to energy efficiency by providing the following characteristics:
-Economy of scale and elimination of redundancies.
-Increased utilization of computing resources.
-Location independence – VMs can be moved to a place where energy is cheaper.
-Scaling up/down and in/out – the resource usage can be adjusted to suit the current requirements.

Emissions

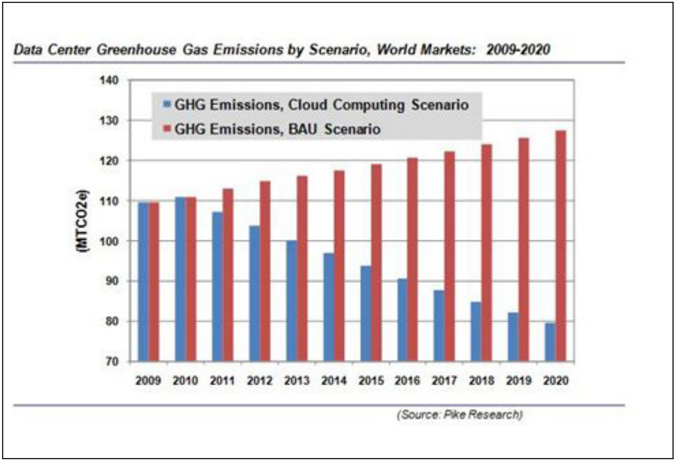

http://www.environmentalleader.com/2010/12/08/cloud-computing-could-cut-data-center-energy-use-38-by-2020/

Emissions

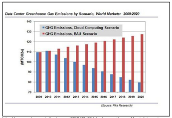

http://www.environmentalleader.com/2010/12/08/cloud-computing-could-cut-data-center-energy-use-38-by-2020/

The graph shows the Business as usual scenario and the Cloud Computing Scenario in camparison, in the business as usual scenario it is assumed that the increasing amount of data centers continue to handle their workloads as in time when there was no cloud computing, in the cloud computing scenario it is assumed that all large scale data centers will use the benefits of cloud computing in order to reduce their energy consumption. The total amount of greenhouse gases produced in 2020 differs by about 50MT thats 38% less Emissions by the Cloud Computing Scenario.

Examples for VMs

Xen Hypervisor

- P- and C-states
- Supports offline and live migration of VMs

Kernel-based Virtual Machine (KVM

- Open source as module for Linux kernel
- Linux as hypervisor

VM-Ware

- Similar to Xen, supports host-level power management via DVFS

Examples for VMs

Xen Hypervisor
- P- and C-states
- Supports offline and live migration of VMs

Kernel-based Virtual Machine (KVM
- Open source as module for Linux kernel
- Linux as hypervisor

VM-Ware
- Similar to Xen, supports host-level power management via DVFS

Three of the most popular virtualization technology solutions: Xen Hypervisor, KVM and VM-Ware. Xen Hypervisor is an open-source technology, licensed under GNU General Public License, it supports ACPI P-states, in addition it supports C-states(CPU sleepstates). Xen supports offline and live migration of VMs, offline migration moves a VM from one host to another by suspending, copying the memory and resuming on the destination host. Live migration transfers VMs without suspension, by starting a new instance on the destination and cyclically copy memory pages to it, so it is inconspicous from the user side. KVM is implemented as a module of the Linux Kernel supports S4(hibernate) and S3 (sleep/stand by) power states, S4 dumps the memory state to a hard disk and powers of the computer, S3 keeps memory powered, so content does not need to be saved to hard disk. VM-Ware is similar to Xen Hypervisor, also supports host-level power management and supports live migration.

Data Centers

- Data Center goals:
  - Save energy to reduce costs
  - Provide reliable Qos for heterogenous workloads
- Mostly achieved by DVFS
- Different Virtualization techniques achieve better results by:
  - swithcing off idle nodes
  - reallocate workload to provide better utilization of resources
- Provided Fault and perfomance isolation between VMs
- Problems handle arbitary workloads with reliable QoS

Data Centers

- Data Center goals:
  - Save energy to reduce costs
  - Provide reliable Qos for heterogenous workloads
- Mostly achieved by DVFS
- Different Virtualization techniques achieve better results by:
  - swithcing off idle nodes
  - reallocate workload to provide better utilization of resources
- Provided Fault and perfomance isolation between VMs
- Problems handle arbitary workloads with reliable QoS

Large scale Data Centers achieved goals are to save energy in order to reduce their expenses, while still providing the best possible Quality of service, this is mostly achieved by DVFS, but also different virtualization techniques can be very effective and achieve even better results. The key factors in virtualization techniques are switching off idle physical nodes as they are still cosume energy in this state, and reallocating workloads to provide better utilization of their resources in order to minimize the network communication between nodes, to reallocate VMs from nodes with low/high utilization,to idle/turn off nodes and to reduce workload of the cooling system. Thereby they provide loss of single points of failure, and perfomance between different VMs is isolated so that low and high traffic VMs can be handled differently. Problem for live migration is to handle arbitary workloads, with reliable quality of service, live migration of different workloads is hard to unify and to process.

Virtualization for Providers

- VM migration causes performance and energy overhead
- Provider is not aware of workloads
  - QoS must be defined workload independent
- Oversubscribe system resources (CPU)
  - most benefits from dynamic VM consolidation
  - risky from QoS perspective: perfomance degradation
- Eliminiation of single points of failure
  - if a compute or controller node fails system stays operable

Virtualization for Providers

- VM migration causes performance and energy overhead
- Provider is not aware of workloads
  - QoS must be defined workload independent
- Oversubscribe system resources (CPU)
  - most benefits from dynamic VM consolidation
  - risky from QoS perspective: perfomance degradation
- Eliminiation of single points of failure
  - if a compute or controller node fails system stays operable

Virtualization for providers causes different challenges and benefits, the migration of VMs produces overhead in perfomance and energy, the latter can be neglected by the advanteges it provides for the total energy cunsumption, the former can be solved by oversubscribing system resources. While the providers are not aware of the workloads on their nodes it is necessary to define a quality of service standard that is workload independent and can handle arbitary workloads. Oversubscribing the CPU benefits the dynamic VM Consolidation, it allows to allocate more VMs on a node than the CPU can handle if the are at maximum workloads; risky from the QoS perspective. Virtualization eliminates sinlge points of failure, if a node fails, the system stays operable it does not happen, that a node fails e.g. due to overheating that the entire system will have a shutdown, because the workloads are isolated.

Users

Average Users

- Create low amount of traffic
- Have little loss of perfomance when put on VM
- Can save much energy for providers when virtualized

Highperformance Users

- Create high amount of traffic
- Utilize big parts of machines
- Virtualization overhead can decrease perfomance significantly
- Decrease of perfomance when charing CPU

Users

Average Users

- Create low amount of traffic
- Have little loss of perfomance when put on VM
- Can save much energy for providers when virtualized

Highperformance Users

- Create high amount of traffic
- Utilize big parts of machines
- Virtualization overhead can decrease perfomance significantly
- Decrease of perfomance when charing CPU

The main majority of Users can be neglegted in decideing wether virtualization is practicable or not, as they create very low ammounts of traffic, have barely notable loss of perfomance when put on VMs, and the advantages for the providers are huge. The interesting users are which create high amounts of traffic and utilize big parts of machines for their computing efforts, the overhead created by virtualization decreses their perfomance significantly, though they rely on perfomance. Highperformance users charing their physical machine with other users creates the same result as the former.

Consequences for Users

- For average users:
  - disadvantages are overseeable
  - unsignificant decrease of perfomance
  - good virtualization techniques no difference will be noticed
  - note: security of data when migrated to multiple servers might be interesting to look at
- For high performance users:
  - dependend on how much perfomance is needed
  - migration should be avoided
  - sharing physical host can be critical
  - virtualization overhead can decrease perfomance

Consequences for Users

- For average users:
  - disadvantages are overseeable
  - unsignificant decrease of perfomance
  - good virtualization techniques no difference will be noticed
  - note: security of data when migrated to multiple servers might be interesting to look at
- For high performance users:
  - dependend on how much perfomance is needed
  - migration should be avoided
  - sharing physical host can be critical
  - virtualization overhead can decrease perfomance

The consequences for averae users are mainly positive or neutral, as they have little or no decrease in perfomance, and no other disadvantages. High perfomance users are to handle seperately, they can only be virtualized when the overhead is little and does not decreae their perfomance, they shall not be reallocated to a new node. Wether there is to place other VMs on the same note with high perfomance users must be considered independently.

Conclusion

- For saving energy on data centers virtualization usefull
- Benefits from cloud computing
- Through modern virtualization techniques
  - overhead and perfomance loss are reduced
  - becomes efficicient and applicable
- Problems with high performance users
- Optiminzation Algorithms long runtime not live applicable

Conclusion

- For saving energy on data centers virtualization usefull
- Benefits from cloud computing
- Through modern virtualization techniques
  - overhead and perfomance loss are reduced
  - becomes efficicient and applicable
- Problems with high performance users
- Optiminzation Algorithms long runtime not live applicable

Concluding is to say virtualization is very usefull, since cloud computing is developing t becomes an realistic alternative, modern virtualization techniques reduce the overhead and make it more efficient, therefor the perfomance loss is reduced, virtualization becomes applicable. There are unsolved problems with high perfomance users, that have to be handled. And the optimization algorithms for virtualization aspects as migration have long runtime and are not live applicable. In the future data centers will use these techniques evermore, because the advantages are increase, while the disadvanteges decrease.

Anton Beloglazov
Energy-Efficient Management of Virtual Machines in Data Centers for Cloud Computing
The University of Melbourne, 2013
http://beloglazov.info/thesis.pdf

Anton Beloglazov
Energy Efficient Resource Management in Virtualized Cloud Data Centers
The University of Melbourne, 2010
http://seelab.ucsd.edu/virtualefficiency/related_papers/26_4039a826.

http://www.environmentalleader.com
http://www.environmentalleader.com/2010/12/08/cloud-computing-could-cut-data-center-energy-use-38-by-2020/
on December 9th 2014