

Cutting power consumption in cloud computing



阳振坤

Senior scientist, Baidu

yangzhenkun@baidu.com

2009.6.26

Agenda



- Background information
- Power consumption in IDC
- Cutting power consumption
- Summary


What is cloud computing

- An emerging computing technology that uses the internet and central remote servers (IDC) to maintain data and applications
- Enabling much more efficient computing by centralizing storage, memory, processing and bandwidth
 - Machine idle ratio: ~70% vs. 90% in traditional
- A kind of large-scale distributed (usually heterogeneous) system

IDC: some numbers

- The data center in Dallas, Oregon: ~50 MW
 - $50\text{MW} \times 0.8 / 200\text{W} = 0.2\text{M}$
 - Average electricity consumption in USA:
~900kwh/month/family, or 1.25KW
- Power consumption is the major cost and constraint of IDC
- About 7000 IDCs in USA


IDC's infrastructure



- Usually it is critical to protect the cooling system by UPS

Power consumption in IDC

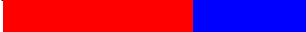
- Machines
 - 50,000 machines (300W each) consume 1.5 MW by themselves
- Network device
- Cooling system
- UPS array and their batteries
 - For both machines and cooling system
- Lighting
- Power Usage Effectiveness (PUE)



Cooling system

- Water-cooling vs. air-cooling
 - Heat transfer coefficient: 100:1
 - Air-cooling device is simpler and cheaper and more suitable for home and ordinary office
 - The vast majority of machines are air-cooling
- “Free cooling”

UPS array



- Contribute to about 11 percent of power consumption and also produce heat
- Obsolete UPS array?
 - Internal battery as standby power for machines
 - Activating standby power (usu. diesel generators) for cooling system in a few minutes after power failure

Lighting

- Contribute to about 3 percent of power consumption and also produce heat
- To cut power consumption by lighting
 - Energy-saving lamps
 - Voice-activated switch...

Machines

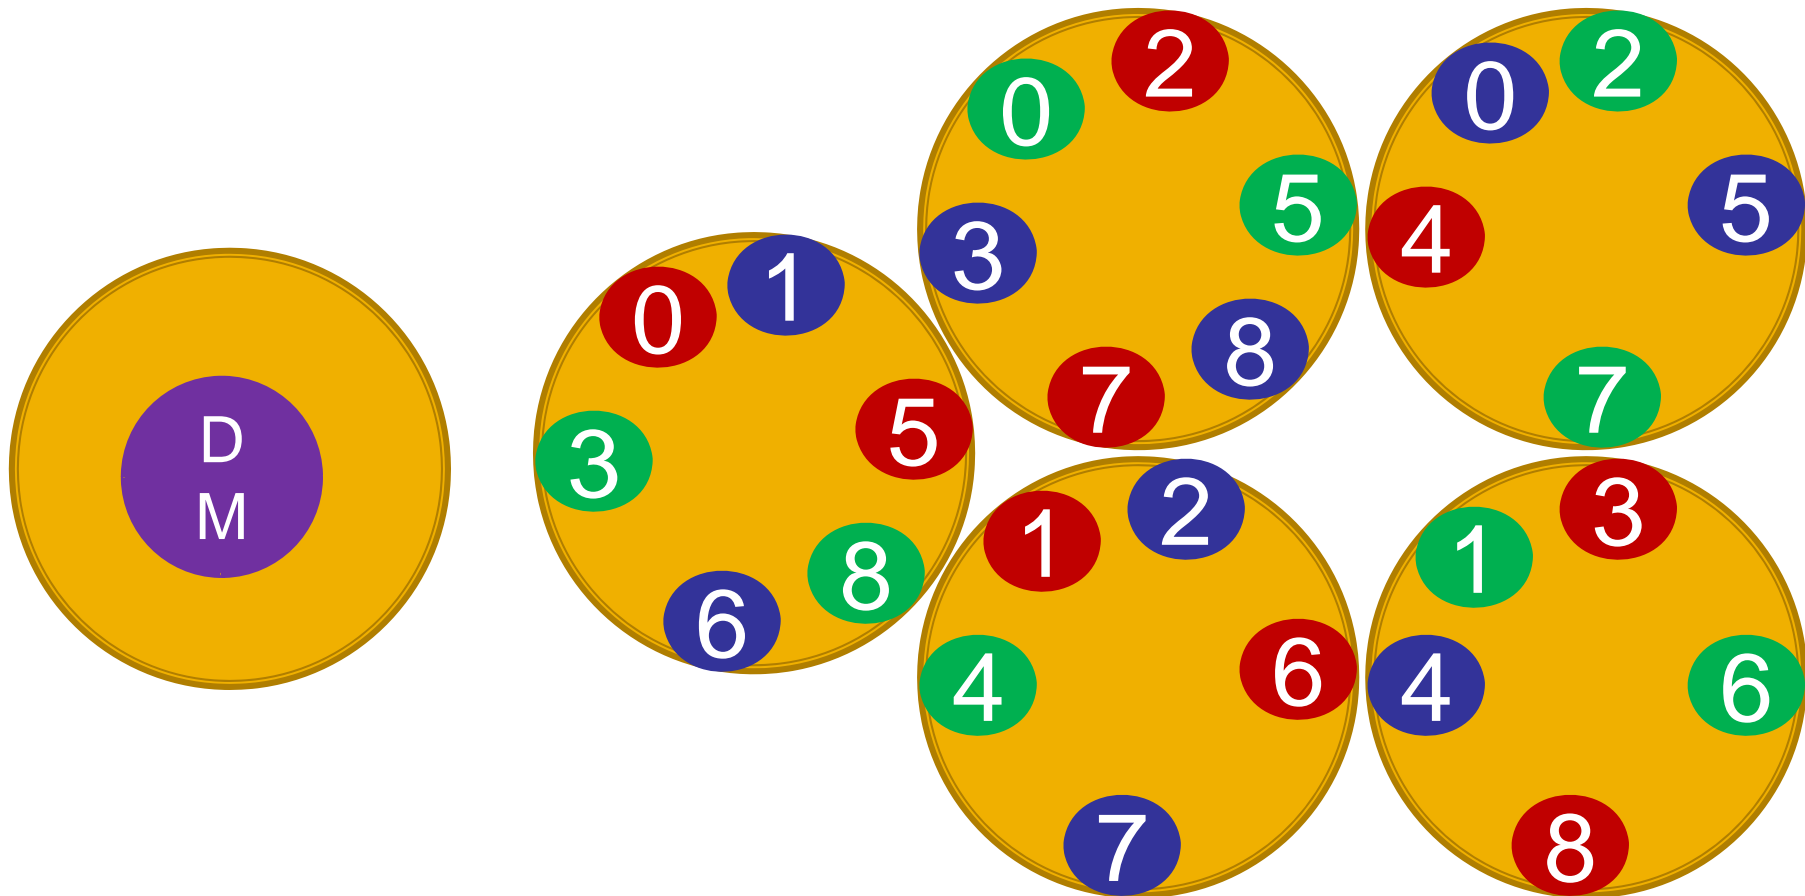
- Processor
- Memory
 - Larger capacity vs. smaller capacity
 - FB/ECC vs. ordinary RAM
- Power supply efficiency
 - ATX: 70%~75%
 - High (400) volt DC vs. AC
- Hard disk
 - HDD vs. SSD disk
 - 3.5" vs. 2.5" hard disk

Machines

- Redundant components
 - Graphics adapter, USB ports, DVD driver, sound chip...
- Cooling system
 - Water-cooling vs. air-cooling
- Machine idle ratio
 - Usually 70% or lower. Can it be cut idle ratio more?
- Standby or hibernation
 - Is it possible in cloud computing?

Data replication in cloud computing

- Data replication、heartbeat, ...
 - Some data become unavailable if many machines unavailable




Cluster: an example

- Requirements
 - 50TB of data, 2,000,000 queries/s and 90% cache hit rate
 - Single machine: 100~400 queries/s, 0.15~1.8TB disk capacity
- Cluster planning
 - $2,000,000 * (1 - 90\%) / 400 = 500$ machines
 - $50\text{TB} * 3 / 500 = 0.3\text{TB/machine}$
- Can we make half or 1/3 of machines standby or hibernated?

Cluster: an example (cont'd)

- Make machines standby/hibernated one by one
 - May lead to mass data shuffle, or
 - Requires data distribution



Summary

- Power is the major constraint and cost of IDC
- To cut power consumption
 - Computer machines
 - Cooling system
 - UPS
 - System infrastructure...
- Whole industry should be involved
- Q & A