Benchmarking With Your Head In The Cloud

Karl Huppler
huppler@us.ibm.com

A Cloud is not All Clouds

A Computing Services Cloud is not All Computing Services Clouds

- Seemingly obvious, yet we hear:

"Cloud is the latest thing – We need a cloud benchmark!"

  - About the same as "We need a multi-processing benchmark" or "We need a database benchmark"

- Before continuing, we need to decide:
  - What key aspects differentiate "cloud" from other forms of computing?
  - What general business model is to be simulated?
  - How can these aspects fit with the requirements of creating a good benchmark?
- Specific to the TPC, we also need to ask
  - How do these conditions fit with the overall benchmark strategy of the TPC?

Cloud Benchmarks Abound!

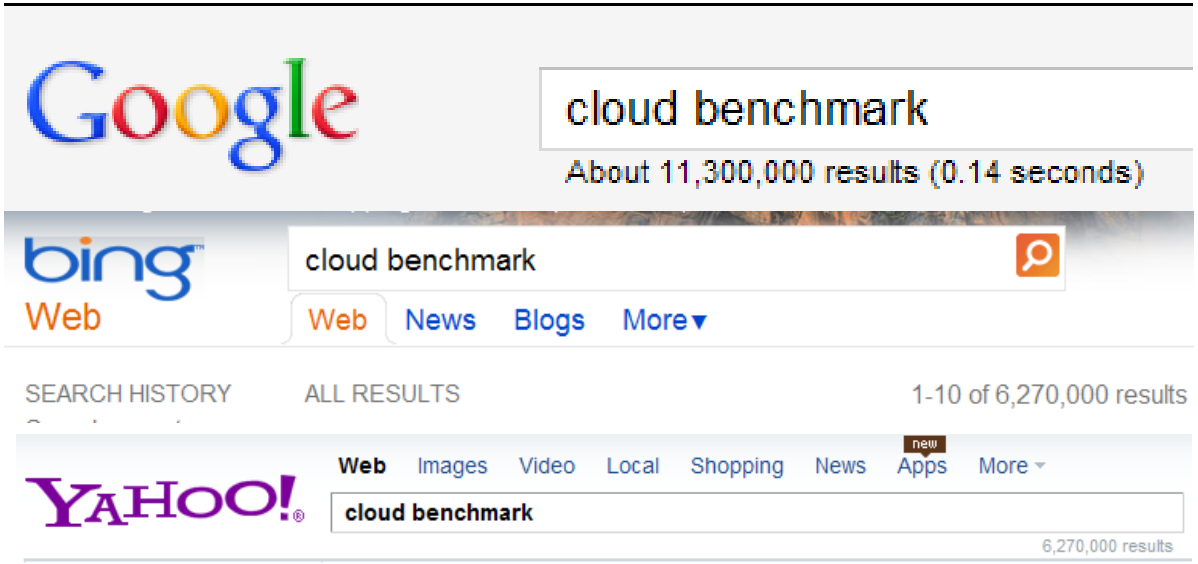

(Use of product-specific references in this presentation is not intended as a statement for or against the use of the product)

Many are focused on single-user applications.

Many are focused on key "cloud" requirements, like bandwidth, ability to expand resources and single-transaction response times

As the industry moves toward cloud computing as a mainstay of commercial use, the need for "run-your-business" database benchmarks grows.

Some options in DB Cloud Benchmarking already exist.

CloudHarmony provides a for-fee service using relatively simple benchmarks to compare a wide variety of application environments.

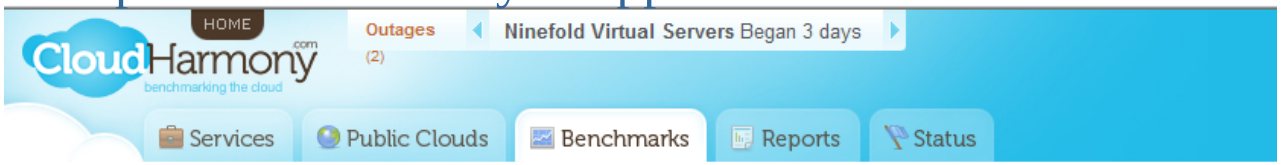

From Cloud Harmony's description:

The standard for this benchmark is published by the TPC (Transaction Processing Performance Council). The description they provide for this benchmark is provided below. The benchmark we perform uses 10 warehouses, 6 operators, a 60 second warm-up, and a 180 second benchmark. The benchmark value is the # of sustained transactions per second. This benchmark is primarily I/O bound, and thus the results can be interpreted as an indication of disk I/O performance. […]

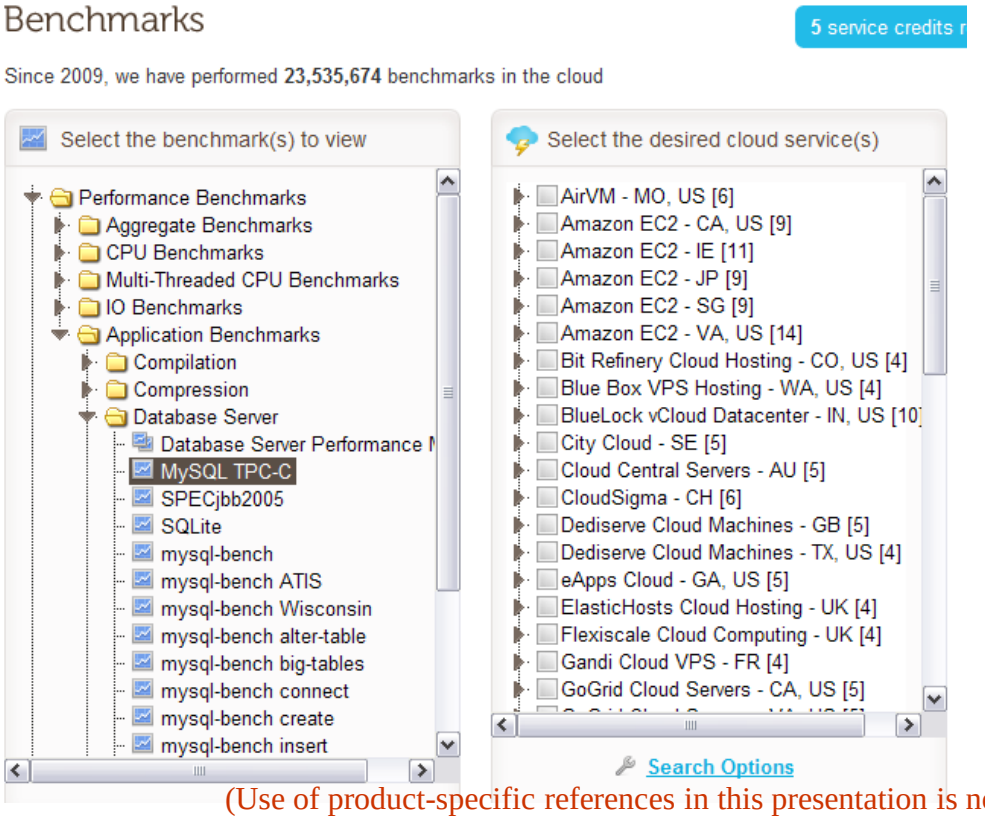

(Use of product-specific references in this presentation is not intended as a statement for or against the use of the product)

Some options in DB Cloud Benchmarking already exist.

Yahoo Research has developed a framework called Yahoo Cloud Serving Benchmark that can accommodate several benchmarks that relate to data processing.

Benchmarking Cloud Serving Systems with YCSB

Brian F. Cooper, Adam Silberstein, Erwin Tam, Raghu Ramakrishnan, Russell Sears

Yahoo! Research
Santa Clara, CA, USA
{cooperb,silberst,etam,ramakris,sears}@yahoo-inc.com

And, if one needs any additional incentive for an industry standard, there are examples of marketing materials developed with cloud database benchmarks

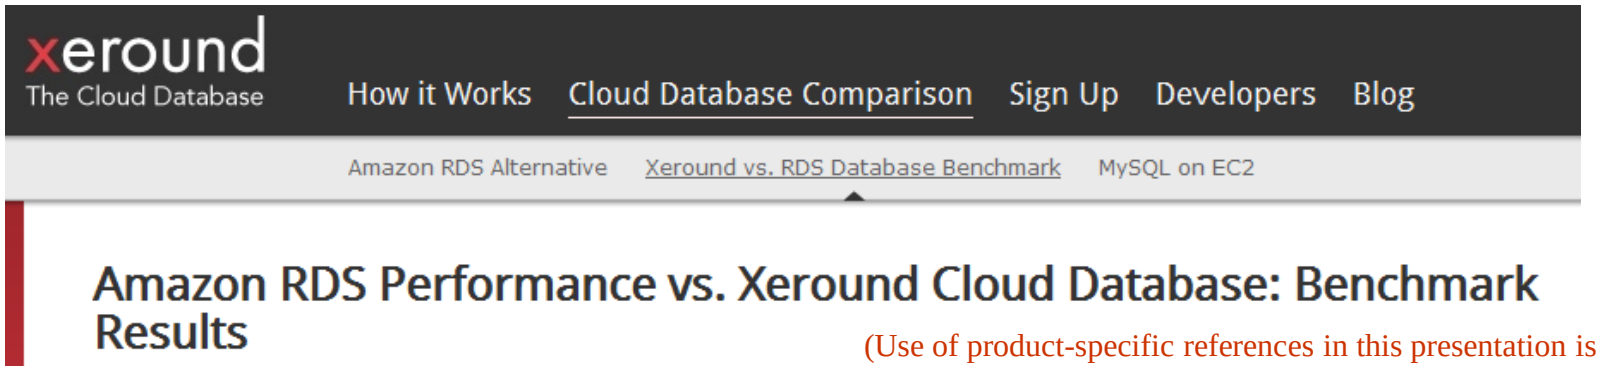

(Use of product-specific references in this presentation is not intended as a statement for or against the use of the product)

What makes a Cloud a Cloud?

Sharing resources is not enough

- Allocation/accounting of compute resources
- Using compute resources controlled by a third party
- Sharing 3rd-party compute resources with others, while seeming to be isolated
- Compute-usage accounting with bill-back for resources consumed/used
- Putting "stuff" out on the world wide web

Additional "cloud" features

- Location and allocation of contracted resources under control of supplier
- Resources (or use of) can be physically migrated, as needed
- Quantity of resource is guaranteed to be available, but resources do not sit idle
- Service Level Agreement established for guaranties during contract
  - QOS
  - Uptime
  - Back-up and Recovery
  - Growth and Surge options
  - Workload volume
  - Charges

Cloud Computing Models

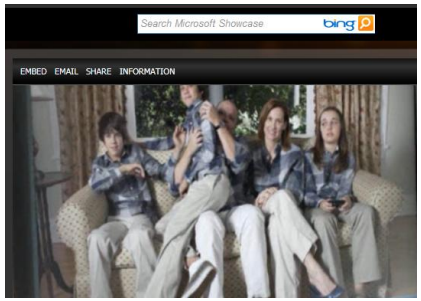

Software as a Service: Use application, storage and computing resources to manipulate personal photographs: "To the cloud!"

Platform as a Service: Use middleware services, storage and computing resources to manage a specific application, or part of an application. Web hosting is a typical example; Adding support from PayPal or a similar service is a business use of the above model.

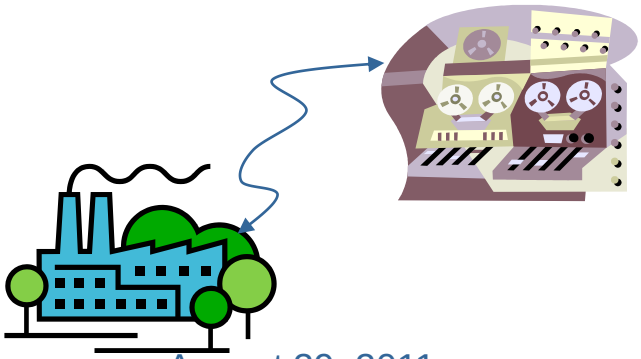

Infrastructure as a Service: Move your business software and data to storage and computing resources that are managed by another group.

(Use of product-specific references in this presentation is not intended as a statement for or against the use of the product)

Relating Cloud Delivery Models to Benchmark Environments


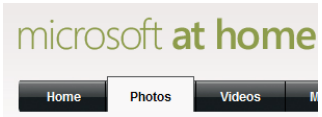

craigslist
ebay
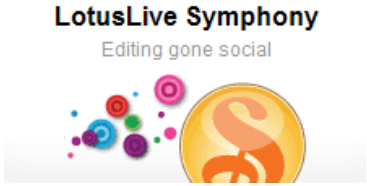

Cisco webex


- **Software as a Service**
  - Today: Most focused on individual user tasks; Response time important
  - Growing: Corporate use for collaborative and mail functions; Response time and bandwidth important
  - Few "mission critical enterprise" environments.
  - Many benchmarks or benchmark-like information exists
    - Not recommended for TPC-like enterprise DB benchmark

(Use of product-specific references in this presentation is not intended as a statement for or against the use of the product)

Relating Cloud Delivery Models to Benchmark Environments

- **Platform as a Service**
  - Web hosting, Application Development, Application Support software shifts from in-house environment to shared-use environments, along with physical resources where applications will execute
  - Application design, application ownership, data management retained within the enterprise
  - Opportunity for "run your business" enterprise applications grows
  - Candidate for public benchmark

(Use of product-specific references in this presentation is not intended as a statement for or against the use of the product)

Relating Cloud Delivery Models to Benchmark Environments

- **Infrastructure as a Service**
  - "Cloud Outsourcing" – Physical processor, storage, memory, usually OS managed by service provider.
  - Differs from traditional outsourcing in that only the service capacity is provided, without guaranteeing specific hardware
  - This is the environment that will see the most early enterprise usage, and is therefore of interest as a format for a public benchmark

(Use of product-specific references in this presentation is not intended as a statement for or against the use of the product)

Public and Private Clouds

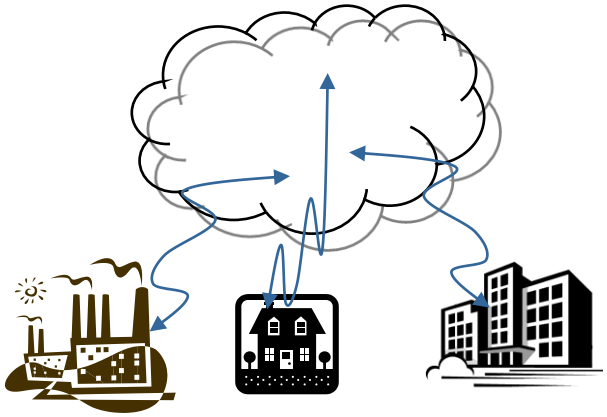

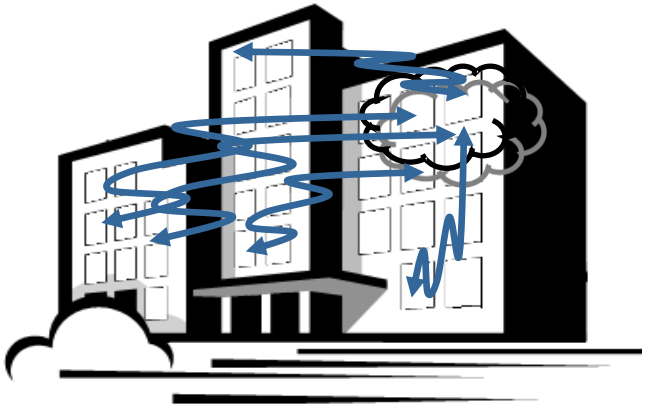

Public:

- "true" cloud
- Requires complex Service Level Agreements (SLAs) to compensate for physical unknown for the enterprise
- Shifts costs from capital to incremental expense

Private

- Combination of "mainframe" accounting with cloud deployment
- Improved control
- But enterprise pays for everything

Each offer unique challenges in creating a public performance benchmark.

TPC Benchmarks, in general

- Database
- Enterprise, "run your business"
  - Data integrity, transaction integrity, system resiliency
- Total System
  - Most benchmarks measure at 100% CPU
  - TPC benchmarks also focus on other resources needed to achieve 100% (storage, memory, networking, etc.)
- Purchase plus Maintenance Pricing
  - Upfront purchase of entire system is assumed
- Auditable, Verifiable, Repeatable

Mapping Cloud Characteristics to TPC Benchmarks - - Price for Public Clouds

- Price models for public clouds do not match TPC price model
  - "rental" – Contract for xx amount of compute power; guaranteed to get it when needed; charged always, used not as much
  - "utility" – Contract to pay-as-you use
  - "pre-paid phone" – combination with "rental" payment for some base plus "utility" payment for intermittent additional requirements.
- TPC's Pay-all-up-front pricing is far removed from this
  - Need some kind of total expense for 3 years, rather than capital investment plus minor maintenance
- Public Cloud pricing can help solve some TPC pricing challenges:
  - Difficulty matching massive benchmark configurations with customer reality
  - Prices of components that are not yet available
  - Verification of discounts
  - Discrepancy between maintenance required for TPC and that required by general customers
- Creating rules for a public cloud benchmark price would
  - Provide a public quote that should be verifiable
  - Use configuration parameters that are reasonable for consumers
  - Only use components that are currently available
  - Include maintenance specified in the SLA

Mapping Cloud Characteristics to TPC Benchmarks - - Price for Private Clouds

- Price models for private clouds do not match TPC price model
  - Lease servers and storage with planned rotating upgrades
    - Outright purchase with depreciation also employed
    - Lease often makes sense based on expense assessments
  - Assess expense costs against internal groups contracting for resources
- Cloud applications do not consume total resources of system
  - Price challenge: Devise a way to price the expense cost for the resources being consumed by the application
- Adjusting TPC price model to accommodate private cloud could help with other TPC price scenarios
  - Creation of new language for lease-based pricing yields an opportunity to close some perceived gaps in TPC Pricing requirements
  - Lease contract requirements should include full service typically required by consumers
  - Brings Hardware and Software rules in line (SW already able to use licenses that expire over time)

Mapping Cloud Characteristics to TPC Benchmarks - - ACID

- TPC's atomicity, consistency, isolation and durability requirements (ACID) are a hallmark of DB benchmarks
- AC&I are likely worth keeping in the cloud environment as they are in more traditional configurations
- Durability needs adjustments in several areas
  - Single points of failure
  - Individual tests
  - Additional requirements

- Public Cloud:
  - Concept of a single point of failure requires understanding of the overall physical configuration – also impractical to test
  - Soft failures can and should be introduced
  - Overall SLA requirements need to establish reliability and failure-recovery thresholds
- Private Cloud:
  - Current durability tests could be employed
  - SLA requirements may not need to be a stringent
  - Similar reliability and failure-recovery thresholds would be an excellent addition

Mapping Cloud Characteristics to TPC Benchmarks - - Availability

- Availability Date is one of the TPC Primary Metrics
  - Current benchmarks allow 185 days between when the result is published and when the configuration becomes available
  - Since the performance of public clouds may be dependent on outside influences, the availability date for such results likely must be "now"
  - Private cloud configurations can likely continue to follow the existing requirements.

Mapping Cloud Characteristics to TPC Benchmarks - - Minimum Functional Delivery Threshold

- It is not enough to just share resources
- Cloud computing has many other characteristics that must be guaranteed by an SLA
  - Particularly for public clouds
  - Also good for private clouds
- Public clouds have to be able to satisfy requirements via an SLA contract, since the physical resources are hidden
- Private clouds could provide proof of requirements, but an element of ease of benchmarking must be employed

- To call a result a "cloud result", TPC must establish minimum delivery rules for:
  - Response Time/Quality of Service
  - Guaranteed availability
  - Time for failure recovery
  - Migration time for server or application migration
  - Overall minimum/average/maximum bandwidth
  - Overall minimum/average/maximum compute power
  - Allocated storage
  - Capacity to scale to 2X or 3X or xX the capacity under contract
  - (Perhaps more)

Audit and Verification Challenges

- Clouds Change
  - How do you prove repeatability?
    - Public cloud – periodic retest?
    - Private cloud – perhaps suite of variable outside loads that consume portions of total system resource
  - How do you verify the claimed resources were used?
    - Vendor-specific utilization claims difficult to audit
    - May have to attest to what is reported, without attesting to its complete validity
  - How do you verify that only the claimed resources were used?
    - Public cloud - - Auditors need to be confident that environment was not customized for the benchmark, even though they cannot "see" the actual resources
- Challenges are not insurmountable – but also not trivial

Conclusions

- Cloud computing is clearly here, and growing
  - Mission Critical, DB applications are being considered for deployment within either private or public clouds
- The time for an enterprise DB cloud benchmark is upon us
  - The most likely delivery models for an enterprise cloud benchmark start with Infrastructure as a Service and move toward Platform as a Service
- Many paradigms of benchmarking must be altered for TPC to create a benchmark that is truly "cloud"
  - Price and Durability testing chief among them
  - Additional SLA-like functional requirements must be added
- Although there are significant challenges, there are also substantial possibilities to move benchmarking technology forward.