Introduction to the Transaction Processing Performance Council (TPC)

Robert M. Lane

Performance, Availability & Architecture Engineering

Sun Microsystems, Inc.

Agenda

- Purpose of the TPC
- Membership
- Organization Chart
- History
- Benchmarks
- Audits and Full Disclosure Reports
- Challenges
- Benchmark Development Process
- Questions and Answers

Purpose of the TPC

The TPC is a non-profit corporation founded to define vendor-neutral transaction processing benchmarks and to disseminate objective, verifiable performance data to the industry.

- Refinement of existing benchmarks
- Development of new benchmarks
- Publication of benchmark results
- Promotion of the TPC model and results
- Resolution of disputes and challenges

Path to Accomplish TPC's Purpose

- Refinement of existing benchmarks
- Development of new benchmarks
- Publication of benchmark results
- Promotion of the TPC model and results
- Resolution of disputes and challenges

Membership

Full Members

Associate Members

Full Members Influence Direction

- Meetings
  - Face-to-face: five or six annually
  - Committee and subcommittee conference calls
- Vote
  - General Council
  - Subcommittees
- Budget
- Benchmarks
  - Revisions
  - New
- Election to standing committees

Organization Chart

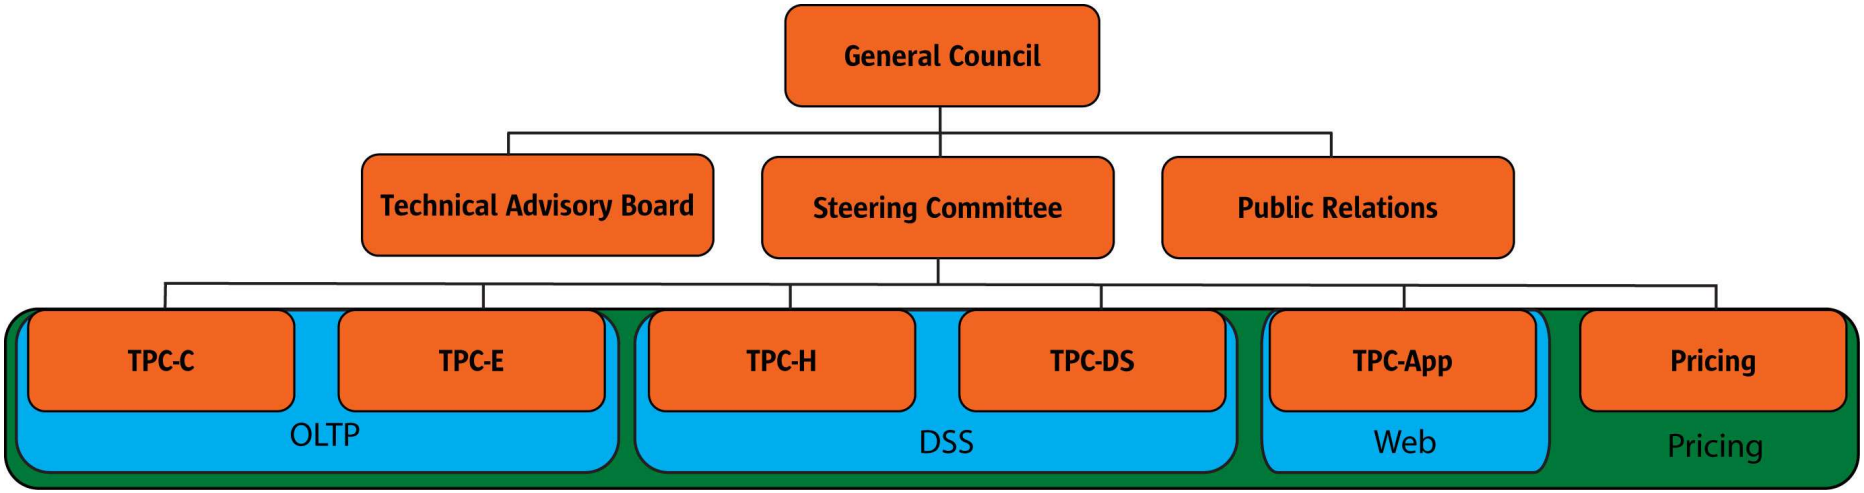

History

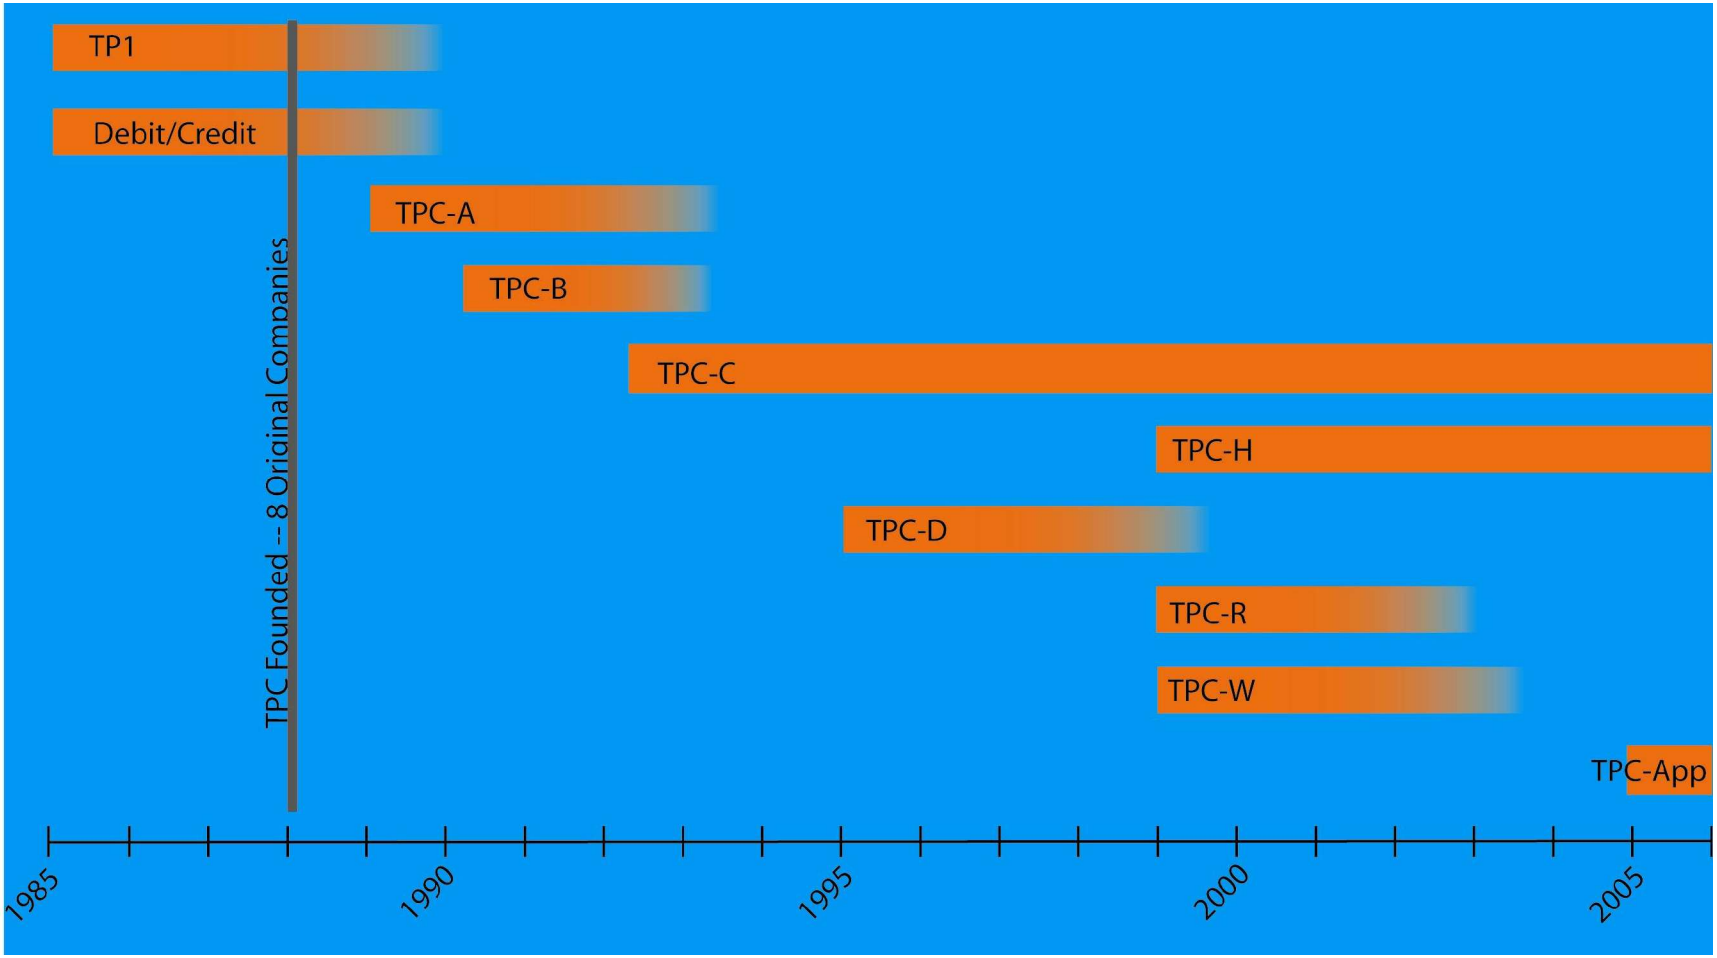
TP1
Debit/Credit
TPC-A
TPC-B
TPC-C
TPC-H
TPC-D
TPC-R
TPC-W
TPC-App
TPC Founded -- 8 Original Companies
1985
1990
1995
2000
2005

Benchmarks

- 3 active benchmarks
  - TPC-C: Online transaction processing (OLTP)
  - TPC-H: Decision support for ad hoc queries
  - TPC-App (application server benchmark)
- 2 benchmarks in development
  - TPC-E (OLTP)
  - TPC-DS (decision support)
- Specification rather than kernel based
- Costly and time consuming (months)

Audits and Full Disclosure Reports

- All benchmarks are audited by TPC certified auditors for compliance and correctness.
- Disclosures
  - Executive Summary (3-5 pages)
  - Full Disclosure Reports (100's of pages)
- Two types of publication
  - Leading primary metric (throughput)
  - Leading price/performance

Typical TPC-C Configuration

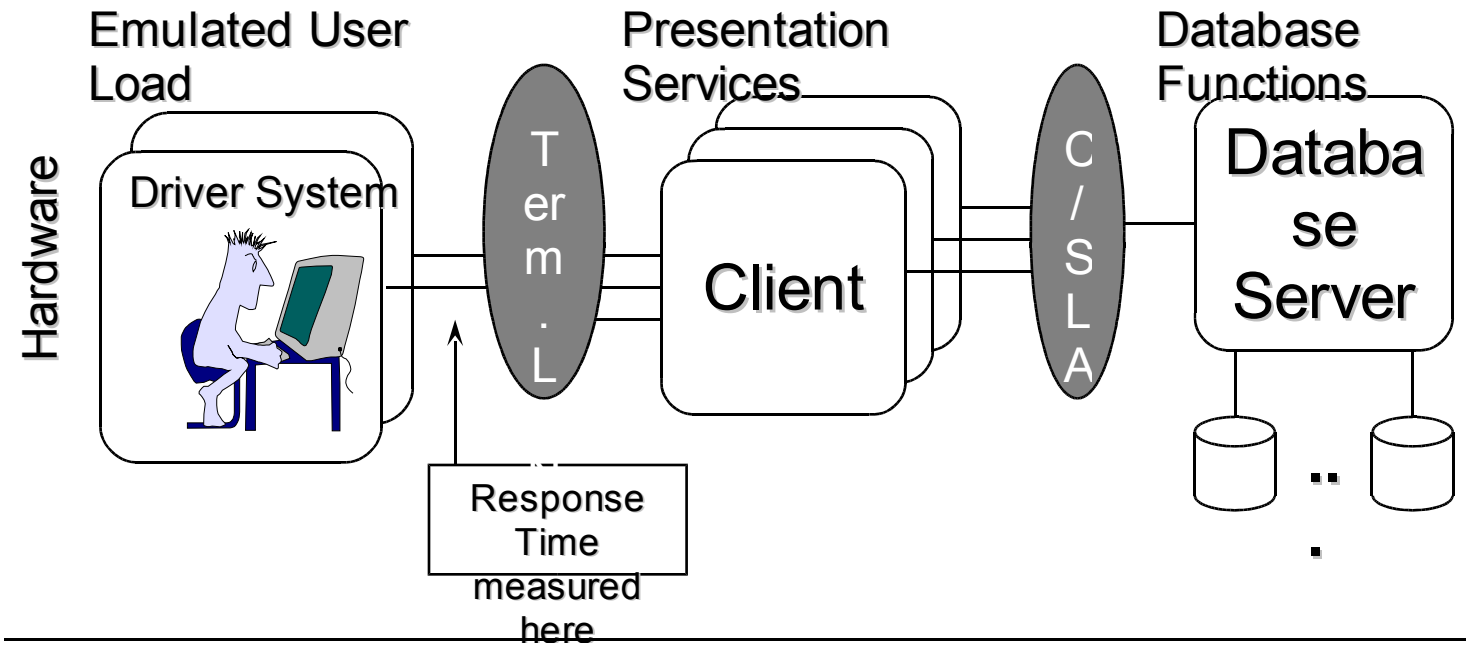

Software	Remote Terminal Emulator (RTE)	TPC-C application + Txn Monitor and/or database RPC library e.g., COM+, Tuxedo, ODBC	TPC-C application (stored procedures) + Database engine e.g., SQL Server, Oracle, DB2

TPC-C Workflow

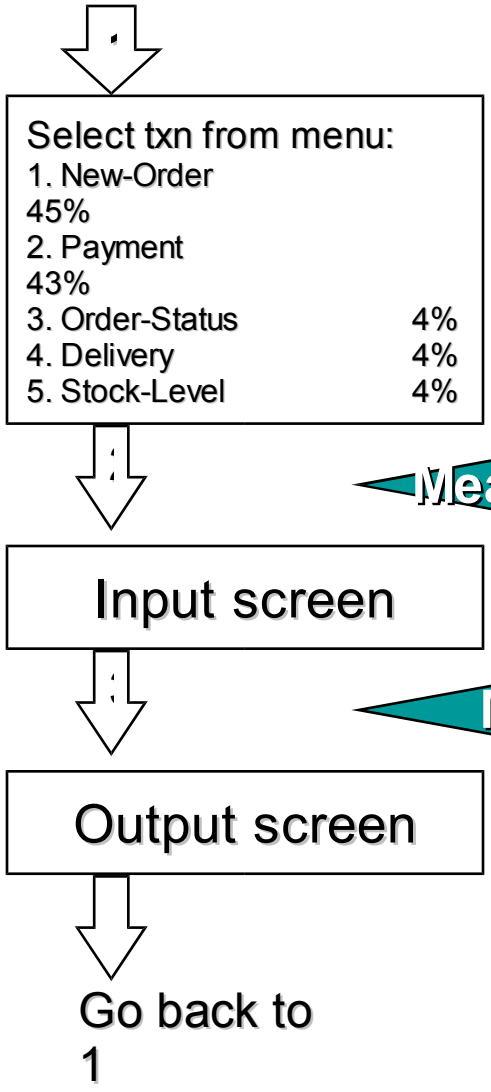

Measure menu Response Time

Keying time

Measure txn Response Time

Think time

Cycle Time Decomposition
(typical values, in seconds, for weighted average txn)

Menu = 0.11

Keying = 5.5

Txn RT = 0.43

Think = 7.5

Avg. cycle time = 13.5

TPC-C Rules of Thumb

- 1.2 tpmC per User/terminal (maximum)
- 10 terminals per warehouse (fixed)
- 65-70 MB/tpmC priced disk capacity (minimum)
- ~ 0.5 physical IOs/sec/tpmC (typical)
- 250-700 KB main memory/tpmC (how much $ do you have? Price vs. Performance trade off. Power will factor in here.)

FDR: Example TPC-C Equipment

System Components	Each of the 128 Clients		Server	
	Quantity	Description	Quantity	Description
Processors	2	3.2GHz 1MB L3 Xeon Processor	64	1.9GHz POWER5™
Memory	4	512 MB	4	512 GB
Disk Controllers	1	SATA	2 90 45	Integrated dual Ultra3SCSI 2Gb FC Adapters IBM DS4500 Controllers
Disk Drives	1	80 GB	6400 140 8	36.4GB 15K RPM FC 73.4GB 15K RPM FC 36.4GB 15K RPM SCSI
Total Storage		10240 GB		243,236GB
Terminals	4	System Console	1	System Console

TPC-H Execution Rules

- High level overview (timed and un-timed portions of the benchmark)

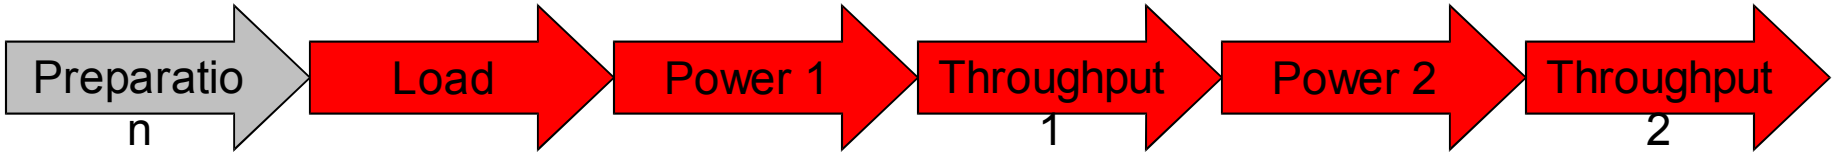

Load is timed → Load Metric
Slower of first and second
Power → Power Metric
Throughput → Throughput Metric

Power and Throughput Test Details

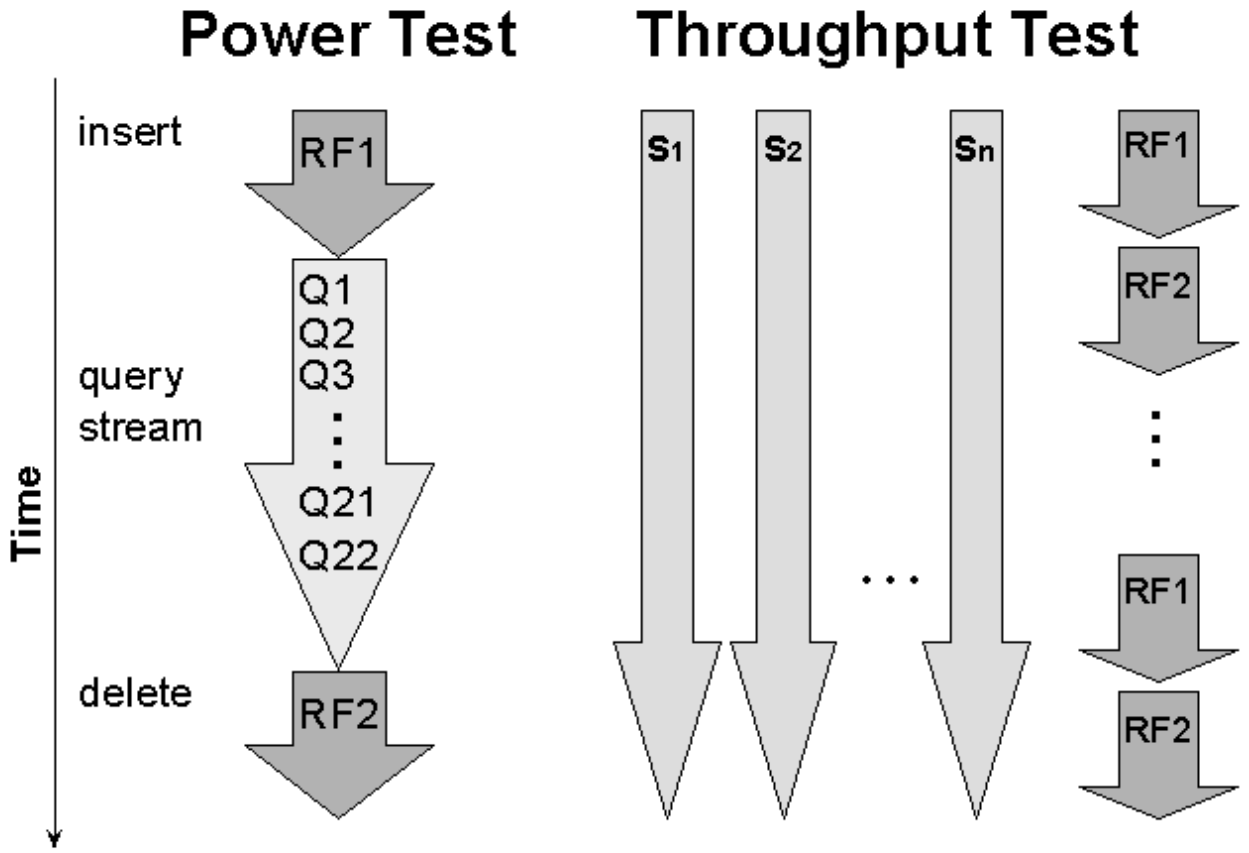

Legend: Q: Query, S: Stream, RF refresh function

FDR: Example TPC-H Equipment

- Processors: 72 UltraSPARC™ IV+ 1500 MHz processors on 72 chips, 144 cores, 144 threads
- Memory: 288GB memory
- Disks:
  - 108 StorEdge 3510 FC Arrays (12x73.4GB)
  - 1 SE6120 (14x73.4GB)
  - 4 S1 (3x73.4GB)
- Total Storage: 97,034.8 (in this calculation one GB is defined as 1024*1024*1024 bytes)

FDR: Lesson

- Large, complex configurations tend to be benchmarked.
- Benchmark equipment tends to be dominated by the disk subsystems.
- Memory and disks subsystems play important roles in TPC benchmarks. They also consume significant power.
- Discounts in the FDR's reflect pricing that any customer should be able to get.

  *The purpose of TPC benchmarks is to provide relevant, objective information to industry users. To achieve that purpose, publication of a TPC benchmark requires **pricing** that:*

  1. *Is no lower than what would be quoted to **Any customer** from the date of publication of the **FDR**.*
  2. *Is actively used by the vendor in the market segment that the **pricing** models or represents (e.g., small business customers, or large corporations, depending on the type of system being priced).*
  3. *A significant number of customers in the market segment that the **pricing** models or represents would plausibly receive in a purchase agreement.*
- Failure to obtain pricing shown in the FDR should be reported to pricing@tpc.org.

Challenges — Compliance

3.3.5.1 If the TAB finds that a Result failed to satisfy one or more specification requirements, the TAB will recommend to the council that either: (1) the Result has an insignificant deviation from the specification or (2) the Result is non-compliant.

3.3.5.2 Non-compliance is recommended to the council if and only if the TAB finds that at least one of the following conditions is applicable:

- Failure to satisfy one or more requirements of the specification that results in incorrect operation of the functions in the business environment the benchmark represents (e.g. Transparency, ACID) regardless of the impact on the primary metrics.
- Failure to meet any of the following items: Availability, Orderability, Clause 0.2, and requirements applied to any Numerical Quantities listed in the Executive Summary.
- The aggregate effect of one or more violations results in more than a 2% difference in price/performance or performance metrics.
- There is an excessive number of clauses violated even though the aggregate difference in price/performance or performance primary metrics is less than or equal to 2%.
- A violation against the same clause language has been voted twice before for the same Test Sponsor within the two year period prior to the result's submission date.

Compliance challenges are heard by the TAB and then affirmed or rejected by the General Council. Interpretation of some rules may seem counter-intuitive or be inconsistently applied.

Slippery slope phenomenon: Can my company gain competitive advantage or is the behavior so egregious the door should be shut? Is the genie already out of the bottle?

Challenges—Fair Use

- When Results are used in publicity, the use is expected to adhere to basic standards of fidelity, candor, and due diligence, the qualities that together add up to, and define, Fair Use of Results.
  - Fidelity: Adherence to facts; accuracy
  - Candor: Above-boardness; needful completeness
  - Due Diligence: Care for integrity of results
  - Legibility: Readability and clarity
- Generally required metrics (Primary Metrics)
  - Throughput
  - Price/Performance
  - Availability Date
- A set of esoteric rules, some with novel interpretations, govern fair use.
- Fair use challenges are heard by the Steering Committee and then affirmed or rejected by the General Council.

Benchmark Development Process

- Lengthy, political process.
  - TPC-DS has been in development for 6 years.
  - TPC-E has been in development for 3 years.
  - Balancing act: software vs. hardware, manufacturers vs. consumers, vendor vs. vendor, etc.
- Beta Benchmark proposal
  - Effort to shortcut the benchmark development process
- Proposals with the greatest chance of success should be well-formed, based on the template of an existing benchmark, and have industry momentum.
- Membership in the TPC gives voting rights and the ability to influence direction.

Who Can Benchmark?

- Anyone who adheres to the requirements of the specification.
- Anyone who hires a certified TPC auditor.
- Anyone who publishes an Executive Summary and Full Disclosure Report to the TPC web site (http://www.tpc.org).
- Anyone who survives the review period and challenge process.
- Generally, the system manufacturer and database software vendor sponsor the benchmark. They are usually, but not always, members of the TPC.

http://www.tpc.org

- Benchmark Results
- Benchmark Specifications
- Benchmark Status Reports
- Bylaws and Policies
- Historical Documents and Articles
- Membership Application
- If you want the TPC to incorporate a power metric into its benchmarks, you should become active in the TPC.

Issues to Consider

- Primary vs. Secondary Metrics
- Role of memory in performance (non-linear) and power consumption
- Role of disk subsystems in performance (queueing) and power consumption
- A power metric is required that works on uniprocessor, behemoth SMP, and large clustered systems so that fair accurate comparisons can be made.
- A metric should push technology forward and not freeze it at a point in time. It should encourage rather than inhibit innovation.
- A metric should be robust and not require tinkering.
- Vendors seek to exploit loopholes in specifications to their advantage as demonstrated by the TPC challenge process. The practice is for a group to slip to the lowest accepted practice.
- How might vendors offload benchmark work to avoid metric penalty?
- What unnatural undesirable benchmark configurations might result from a given metric?

Questions and Answers