8. Datenbank-Benchmarks

- Benchmark-Anforderungen
- TPC-Benchmarks
- OLTP-Benchmarks
  - TPC-C
  - TPC-E
- Decision Support Benchmark: TPC-H

Anforderungen an geeignete Benchmarks*

- Domain-spezifische Benchmarks
  - kein geeignetes Leistungsmaß für alle Anwendungsklassen möglich
  - spezielle Benchmarks für techn./wissenschaftliche Anwendungen, DB-Anwendungen, etc.
- Relevanz
  - Berücksichtigung “typischer” Operationen des jeweiligen Anwendungsbereichs
  - Messen der maximalen Leistung
  - Berücksichtigung der Systemkosten (Kosteneffektivität)
- Portierbarkeit
  - Übertragbarkeit auf verschiedene Systemplattformen
  - Vergleichbarkeit
- Skalierbarkeit
  - Anwendbarkeit auf kleine und große Computersysteme
  - Übertragbarkeit auf verteilte/parallele Systeme
- Einfachheit / Verständlichkeit

* J. Gray (ed.): The Benchmark Handbook for Database and Transaction Processing Systems. 2nd ed., Morgan Kaufmann, 1993. *http://research.microsoft.com/en-us/um/people/gray/BenchmarkHandbook/TOC.htm*

TPC-Benchmarks

- Herstellergremium zur Standardisierung von DB-Benchmarks (www.tpc.org)
  - Gründung 1988
  - erste Benchmarks für Kontenbuchung (“Debit-Credit”): TPC-A, TPC-B (1989/90)
- besondere Merkmale
  - Leistung eines Gesamt-Systems wird bewertet
  - Bewertung der Kosteneffektivität (Kosten / Leistung)
  - skalierbare Konfigurationen
  - verbindliche Richtlinien zur Durchführung und Dokumentation (Auditing; Full Disclosure Reports)
  - Ausschluß von “Benchmark Specials” innerhalb von DBMS etc. .
- aktuelle Benchmarks für
  - OLTP (TPC-C)
  - Web-OLTP (TPC-E)
  - Decision Support (TPC-H)

TPC-C

- 1992 eingeführter OLTP-Benchmark
- 9 Tabellen (Bestellverwaltung von Warenhäusern)
- 5 Transaktionstypen
  - New-order (45%)
  - Payment (43%), Delivery (4%) , Order-status (4%), Stock-level(4%)
- DB skaliert proportional zum Durchsatz
- Hohe Lokalität
- Metriken
  - Durchsatz von NewOrder-Transaktionen pro Minute (**tpmC**)
  - Price/performance **($/tpmC**)

TPC-C DB-Schema

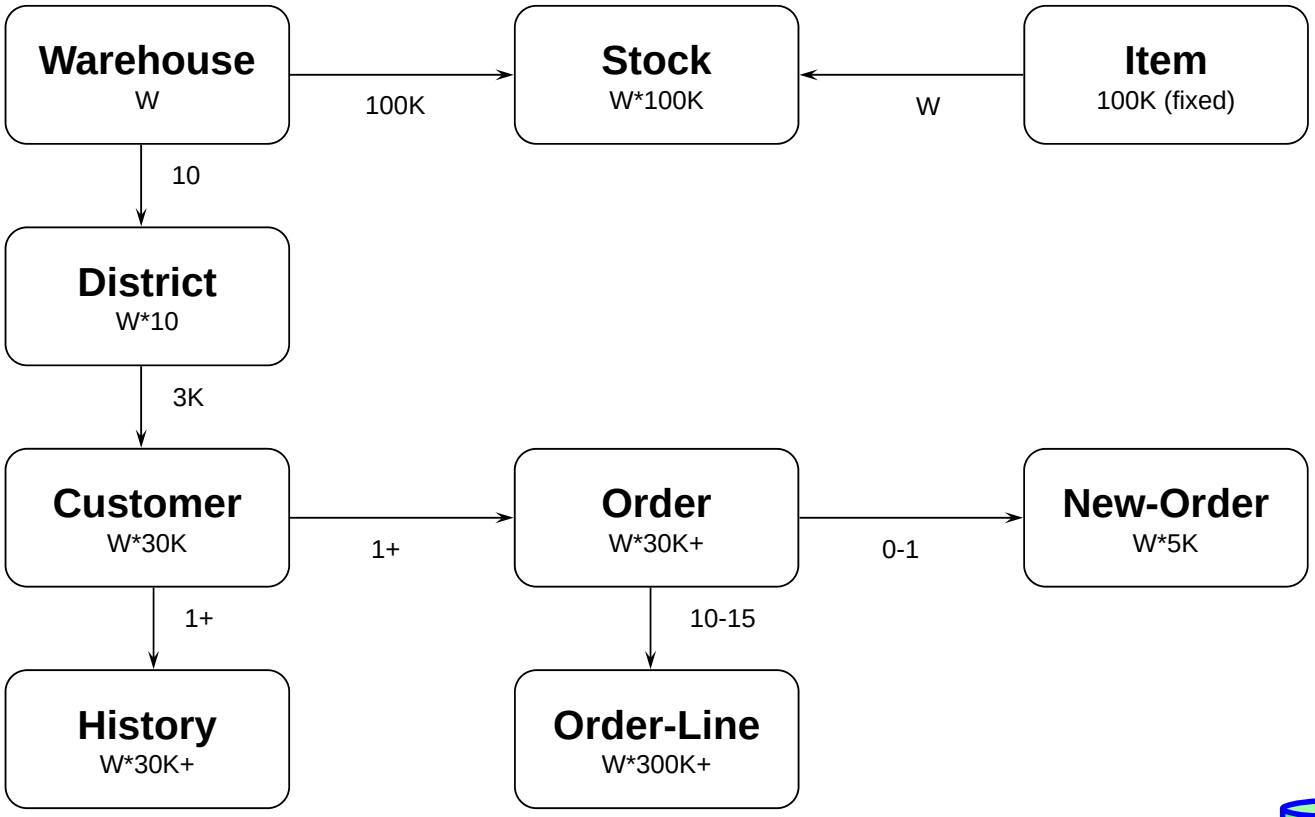

Transaktionstyp New Order (TPC-C)

- im Mittel 48 SQL-Anweisungen (BOT, 23 SELECT, 11 UPDATE, 12 INSERT, EOT)
- 1% der Transaktionen sollen zurückgesetzt werden

```
BEGIN WORK   { Beginn der Transaktion }
    SELECT ... FROM CUSTOMER
        WHERE c_w_id = :w_no AND c_d_id = :d_no AND c_id = :cust_no
    SELECT ... FROM WAREHOUSE WHERE w_id = :w_no
    SELECT ... FROM DISTRICT  (* -> next_o_id *)
        WHERE d_w_id = :w_no AND d_id = :d_no
    UPDATE DISTRICT SET d_next_o_id := :next_o_id + 1
        WHERE d_w_id = :w_no AND d_id = :d_no
    INSERT INTO NEW_ORDER ...
    INSERT INTO ORDERS ...
    pro Artikel (im Mittel 10) werden folgende Anweisungen ausgeführt:
        SELECT ... FROM ITEM WHERE ...
        SELECT ... FROM STOCK WHERE ...
        UPDATE STOCK ...
        INSERT INTO ORDER-LINE ...
COMMIT WORK   { Ende der Transaktion }
```

TPC-C: Entwicklung 1994-2004

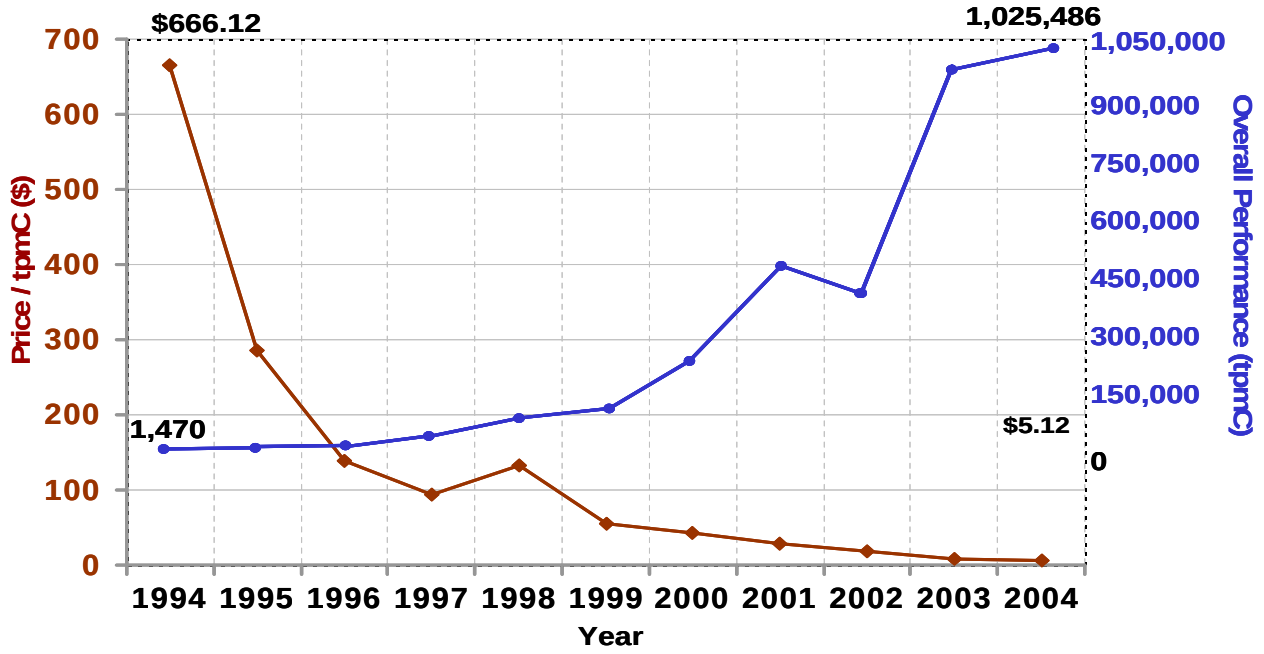

Aktuelle TPC-C-Ergebnisse: Top-Durchsatz

Rank	Company	System	Performance (tpmC)	Price/tpmC	Watts/KtpmC	System Availability	Database	Operating System	TP Monitor	Date Submitted	Cluster
1	ORACLE	SPARC SuperCluster with T3-4 Servers	30,249,688	1.01 USD	NR	06/01/11	Oracle Database 11g Release 2 Ent. Ed. w/Real Application Clusters w/P	Oracle Solaris 10 09/10	Tuxedo CFS-R	12/02/10	Y
2	IBM	IBM Power 780 Server Model 9179-MHB	10,366,254	1.38 USD	NR	10/13/10	DB2 9.7	AIX Version 6.1	Microsoft COM+	08/17/10	Y
3	ORACLE	Sun SPARC Enterprise T5440 Server Cluster	7,646,486	2.36 USD	NR	03/19/10	Oracle Database 11g Ent. Ed. w/Real Application Clusters w/Partitionin	Sun Solaris 10 10/09	Tuxedo CFS-R	11/03/09	Y
4	IBM	IBM Power 595 Server Model 9119-FHA	6,085,166	2.81 USD	NR	12/10/08	IBM DB2 9.5	IBM AIX 5L V5.3	Microsoft COM+	06/10/08	N
***	BULL	Bull Escala PL6460R	6,085,166	2.81 USD	NR	12/15/08	IBM DB2 9.5	IBM AIX 5L V5.3	Microsoft COM+	06/15/08	N
5	hp invent	HP Integrity Superdome-Itanium2/1.6GHz/24MB iL3	4,092,799	2.93 USD	NR	08/06/07	Oracle Database 10g R2 Enterprise Edt w/Partitioning	HP-UX 11i v3	BEA Tuxedo 8.0	02/27/07	N
6	IBM	IBM System p5 595	4,033,378	2.97 USD	NR	01/22/07	IBM DB2 9	IBM AIX 5L V5.3	Microsoft COM+	01/22/07	N

TPC-C: Top-Preis/Performance

Rank	Company	System	Performance (tpmC)	Price/tpmC	Watts/KtpmC	System Availability	Database	Operating System	TP Monitor	Date Submitted	Cluster
1	hp invent	HP ProLiant ML350 G6	290,040	.39 USD	4.22	08/16/10	Oracle Database 11g Release 2 Standard Ed One	Oracle Enterprise Linux	Microsoft COM+	08/16/10	N
2	hp invent	HP ProLiant DL580 G7	1,807,347	.49 USD	2.46	10/15/10	Microsoft SQL Server 2005 Enterprise x64 Edition SP3	Microsoft Windows Server 2008 R2 Enterprise Edition	Microsoft COM+	08/27/10	N
3	DELL	Dell PowerEdge T710	239,392	.50 USD	NR	11/18/09	Oracle Database 11g Standard Edition One	Microsoft Windows Server 2003 Enterprise x64 Edition	Microsoft COM+	11/18/09	N
4	hp invent	HP ProLiant ML350 G6	232,002	.54 USD	NR	05/21/09	Oracle Database 11g Standard Edition One	Oracle Enterprise Linux	Microsoft COM+	05/21/09	N
5	hp invent	HP ProLiant DL385G7	705,652	.60 USD	NR	09/01/10	Microsoft SQL Server 2005 Enterprise x64 Edition SP3	Microsoft Windows Server 2008 R2 Enterprise Edition	Microsoft COM+	04/08/10	N
6	DELL	Dell PowerEdge 2900	104,492	.60 USD	NR	02/20/09	Oracle Database 11g Standard Edition One	Microsoft Windows Server 2003 Standard Ed. x64	Microsoft COM+	02/20/09	N

Business Model – Financial Market

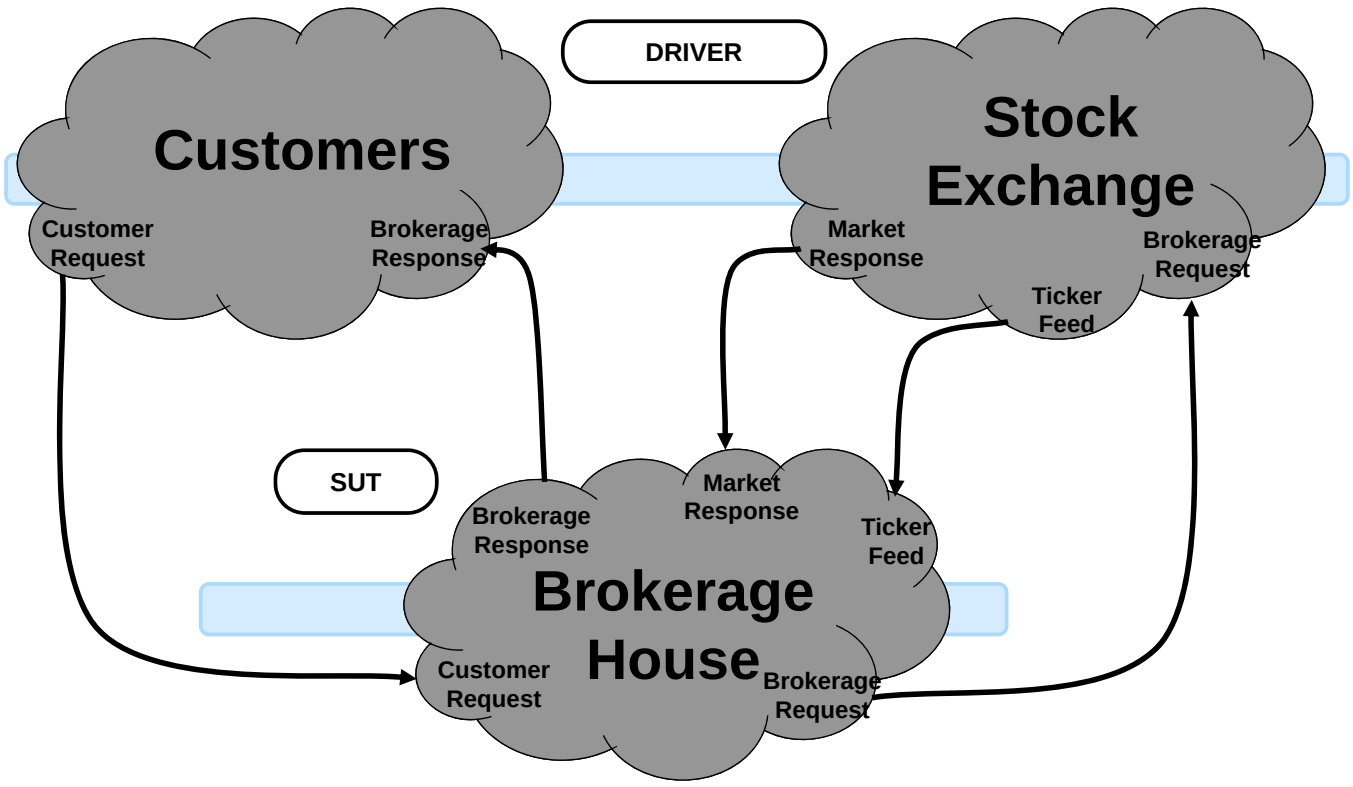

TPC-E

- 2007 verabschiedet
- wesentlich komplexer als TPC-C
  - 33 statt 9 Tabellen
  - 188 statt 92 Spalten
  - Referentielle Integrität
  - Obligatorische RAID-Datenspeicherung
  - Pseudo-reale Daten (z.B. für Kundennamen)
- 10 Transaktionstypen
  - Consumer-to-Business: Trade-Order, Trade-Result, Trade-Update, …
  - Business-to-Business: Broker Volume, Market-Watch, ..
- Metriken
  - Durchsatz von TradeResult-Transaktionen pro Sekunde (**tpsE**)
  - Price/performance **($/tpsE**)

TPC-E-Ergebnisse: Top-Durchsatz

Rank	Company	System	Performance (tpsE)	Price/tpsE	Watts/tpsE	System Availability	Database	Operating System	Processors / Cores / Threads	Date Submitted
1	FUJITSU	PRIMERGY RX900 S1	3,800.00	245.82 USD	NR	10/01/10	Microsoft SQL Server 2008 R2 Datacenter Edition	Microsoft Windows Server 2008 R2 Datacenter Edition	8 / 64 / 128	09/24/10
2	FUJITSU	PRIMEQUEST 1800E	3,800.00	283.03 USD	NR	10/01/10	Microsoft SQL Server 2008 R2 Datacenter Edition	Microsoft Windows Server 2008 R2 Datacenter Edition	8 / 64 / 128	10/26/10
3	NEC	NEC Express5800/A1080a-E	3,141.76	768.92 USD	NR	07/30/10	Microsoft SQL Server 2008 R2 Datacenter Edition	Microsoft Windows Server 2008 R2 Datacenter Edition	8 / 64 / 128	03/30/10
4	FUJITSU	PRIMERGY RX600 S5	2,046.96	193.68 USD	NR	09/01/10	Microsoft SQL Server 2008 R2 Enterprise Edition	Microsoft Windows Server 2008 Enterprise Ed x64 R2	4 / 32 / 64	06/18/10
5	IBM	IBM System x3850 X5	2,022.64	355.02 USD	NR	07/30/10	Microsoft SQL Server 2008 R2 Enterprise Edition	Microsoft Windows Server 2008 R2 Enterprise Edition	4 / 32 / 64	03/30/10

TPC-E: Top-Preis/Performance

Rank	Company	System	Performance (tpsE)	Price/tpsE	Watts/tpsE	System Availability	Database	Operating System	Processors / Cores / Threads	Date Submitted
1	FUJITSU	PRIMERGY RX300 S6	1,246.13	191.48 USD	NR	11/01/10	Microsoft SQL Server 2008 R2 Enterprise Edition	Microsoft Windows Server 2008 R2 Enterprise Edition	2 / 12 / 24	11/29/10
2	FUJITSU	PRIMERGY RX600 S5	2,046.96	193.68 USD	NR	09/01/10	Microsoft SQL Server 2008 R2 Enterprise Edition	Microsoft Windows Server 2008 Enterprise Ed x64 R2	4 / 32 / 64	06/18/10
3	FUJITSU	PRIMERGY RX900 S1	3,800.00	245.82 USD	NR	10/01/10	Microsoft SQL Server 2008 R2 Datacenter Edition	Microsoft Windows Server 2008 R2 Datacenter Edition	8 / 64 / 128	09/24/10
4	DELL	Dell PowerEdge T710	1,074.14	264.32 USD	NR	06/21/10	Microsoft SQL Server 2008 R2 Enterprise x64 Edition	Microsoft Windows Server 2008 R2 Enterprise x64 Edition	2 / 12 / 24	06/21/10
5	DELL	Dell PowerEdge T610	766.47	273.65 USD	NR	03/30/09	Microsoft SQL Server 2008 Enterprise x64 Edition	Microsoft Windows Server 2008 Enterprise x64 Edition	2 / 8 / 16	03/30/09

TPC-H

- Benchmark für Decision Support (Ad-Hoc-Queries)
  - 1999 eingeführt
- Unterschiedliche DB-Größen (100 GB – 30 TB) über Skalierungsfaktor
- 22 Query-Typen, 2 Update-Operationen
- Zeitbedarf für Einbenutzer- und Mehrbenutzerausführung (single/multiple streams)
- Metriken
  - Query-per-Hour (QphH@Size)
  - $/QphH@Size

TPC-H Schema

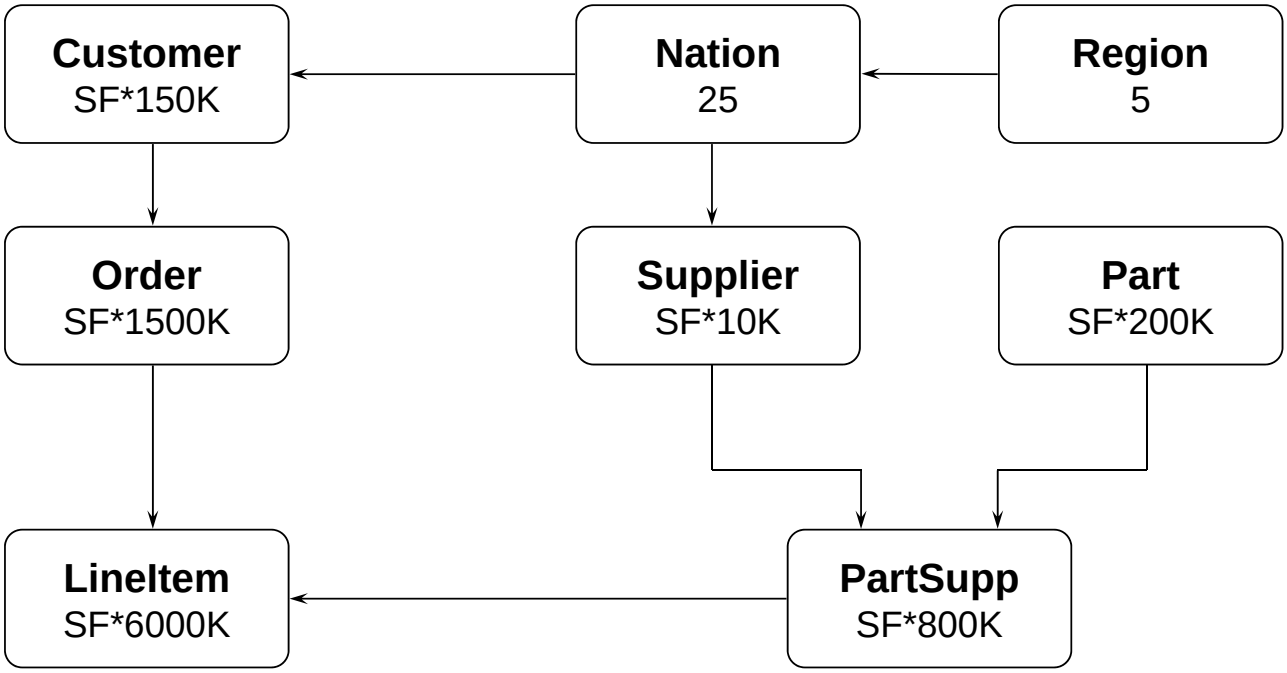

Beispiel-Query Q7 (Volume Shipping)

```
/* TPC_H Query 7 - Volume Shipping */
SELECT SUPP_NATION, CUST_NATION, L_YEAR, SUM(VOLUME) AS REVENUE
FROM ( SELECT N1.N_NAME AS SUPP_NATION
 , N2.N_NAME AS CUST_NATION, datepart(yy, L_SHIPDATE) AS L_YEAR
 , L_EXTENDEDPRICE*(1-L_DISCOUNT) AS VOLUME
FROM SUPPLIER, LINEITEM, ORDERS, CUSTOMER, NATION N1, NATION N2
WHERE S_SUPPKEY = L_SUPPKEY AND O_ORDERKEY = L_ORDERKEY
 AND C_CUSTKEY = O_CUSTKEY AND S_NATIONKEY = N1.N_NATIONKEY
 AND C_NATIONKEY = N2.N_NATIONKEY
 AND ((N1.N_NAME = 'FRANCE' AND N2.N_NAME = 'GERMANY')
  OR (N1.N_NAME = 'GERMANY' AND N2.N_NAME = 'FRANCE'))
 AND L_SHIPDATE BETWEEN '1995-01-01' AND '1996-12-31' ) AS SHIPPING
GROUP BY SUPP_NATION, CUST_NATION, L_YEAR
ORDER BY SUPP_NATION, CUST_NATION, L_YEAR
```

Zur Unterstützung bei der Verhandlung über neue Lieferverträge soll der Wert der zwischen Frankreich und Deutschland transportierten Güter festgestellt werden. Dazu werden jeweils die rabattierten Einnahmen in den Jahren 1995 und 1996 berechnet, die aus Auftragspositionen resultieren, bei denen der Lieferant aus dem einen, und der Kunde aus dem anderen Land stammt (also vier Ergebnistupel).

TPC-H: Performance

100 GB Results

Rank	Company	System	QphH	Price/QphH	Watts/KQphH	System Availability	Database	Operating System	Date Submitted	Cluster
1	CPI	CPI Phoenix IQ-201	209,298	1.25 USD	NR	01/14/08	EXASOL EXASolution 2.0	EXASOL EXACluster OS 1.3	01/14/08	Y
2	ORACLE	SunFire X4100	98,857	2.65 USD	NR	10/29/07	ParAccel Analytic Database	Red Hat Linux 4.4	10/29/07	Y
3	hp invent	HP ProLiant DL380 G7	73,974	.58 USD	5.93	07/02/10	Microsoft SQL Server 2008 R2 Enterprise Edition	Microsoft Windows Server 2008 R2 Enterprise Edition	07/02/10	N

3,000 GB Results

Rank	Company	System	QphH	Price/QphH	Watts/KQphH	System Availability	Database	Operating System	Date Submitted	Cluster
1	FUJITSU	PRIMERGY RX300 S4	1,608,920	1.36 USD	NR	08/01/08	EXASOL EXASolution 2.1	EXASOL EXACluster OS 2.1	06/02/08	Y
2	ORACLE	Sun SPARC Enterprise M9000 Server	198,907	16.58 USD	NR	12/09/10	Oracle Database 11g Release 2 Enterprise Edt.	Oracle Solaris 10	10/05/10	N
3	hp invent	HP ProLiant DL980 G7	162,601	2.68 USD	NR	10/13/10	Microsoft SQL Server 2008 R2 Enterprise Edition	Microsoft Windows Server 2008 R2 Enterprise Edition	06/21/10	N

TPC-H: Preis/Performance

100 GB Results

Rank	Company	System	QphH	Price/QphH	Watts/KQphH	System Availability	Database	Operating System	Date Submitted	Cluster
1	hp invent	HP ProLiant DL385 G7	71,438	.51 USD	6.48	07/14/10	Microsoft SQL Server 2008 R2 Enterprise Edition	Microsoft Windows Server 2008 R2 Enterprise Edition	07/14/10	N
2	hp invent	HP ProLiant DL380 G7	73,974	.58 USD	5.93	07/02/10	Microsoft SQL Server 2008 R2 Enterprise Edition	Microsoft Windows Server 2008 R2 Enterprise Edition	07/02/10	N
3	KICKFIRE	Kickfire Appliance 2300	49,228	.70 USD	NR	10/14/08	MySQL 5.1 w/ Kickfire database	Kickfire CentOS 5.0	04/14/08	N

3,000 GB Results

Rank	Company	System	QphH	Price/QphH	Watts/KQphH	System Availability	Database	Operating System	Date Submitted	Cluster
1	FUJITSU	PRIMERGY RX300 S4	1,608,920	1.36 USD	NR	08/01/08	EXASOL EXASolution 2.1	EXASOL EXACluster OS 2.1	06/02/08	Y
2	hp invent	HP ProLiant DL980 G7	162,601	2.68 USD	NR	10/13/10	Microsoft SQL Server 2008 R2 Enterprise Edition	Microsoft Windows Server 2008 R2 Enterprise Edition	06/21/10	N
3	ORACLE	SunFire X4540	11,435	5.26 USD	NR	04/09/09	Sybase IQ 15 SAS	Sun Solaris 10	03/04/09	N