

NCR Teradata Warehouse

Projektpraktikum im betrieblichen Umfeld 4.0 PR (WI/PPR) WS 2003/2004

Technische Universität Wien

Vincent Hager
Senior Consultant CRM Solutions
Teradata Austria
vincent.hager@teradata-ncr.com


Dec 2003

You've never seen your business like this before.

Agenda

- **Teradata Warehouse Architecture**
 - > Parallel Foundation
 - > HW + SW Architecture
 - > Data Management + Placement
 - > Optimizer Functions
 - > Scalability, Availability, Performance
- **Teradata Tools & Utilities**
- **Teradata Warehouse 7.0 Demo System (Demonstration)**


Teradata Warehouse Architecture



Client-Tier

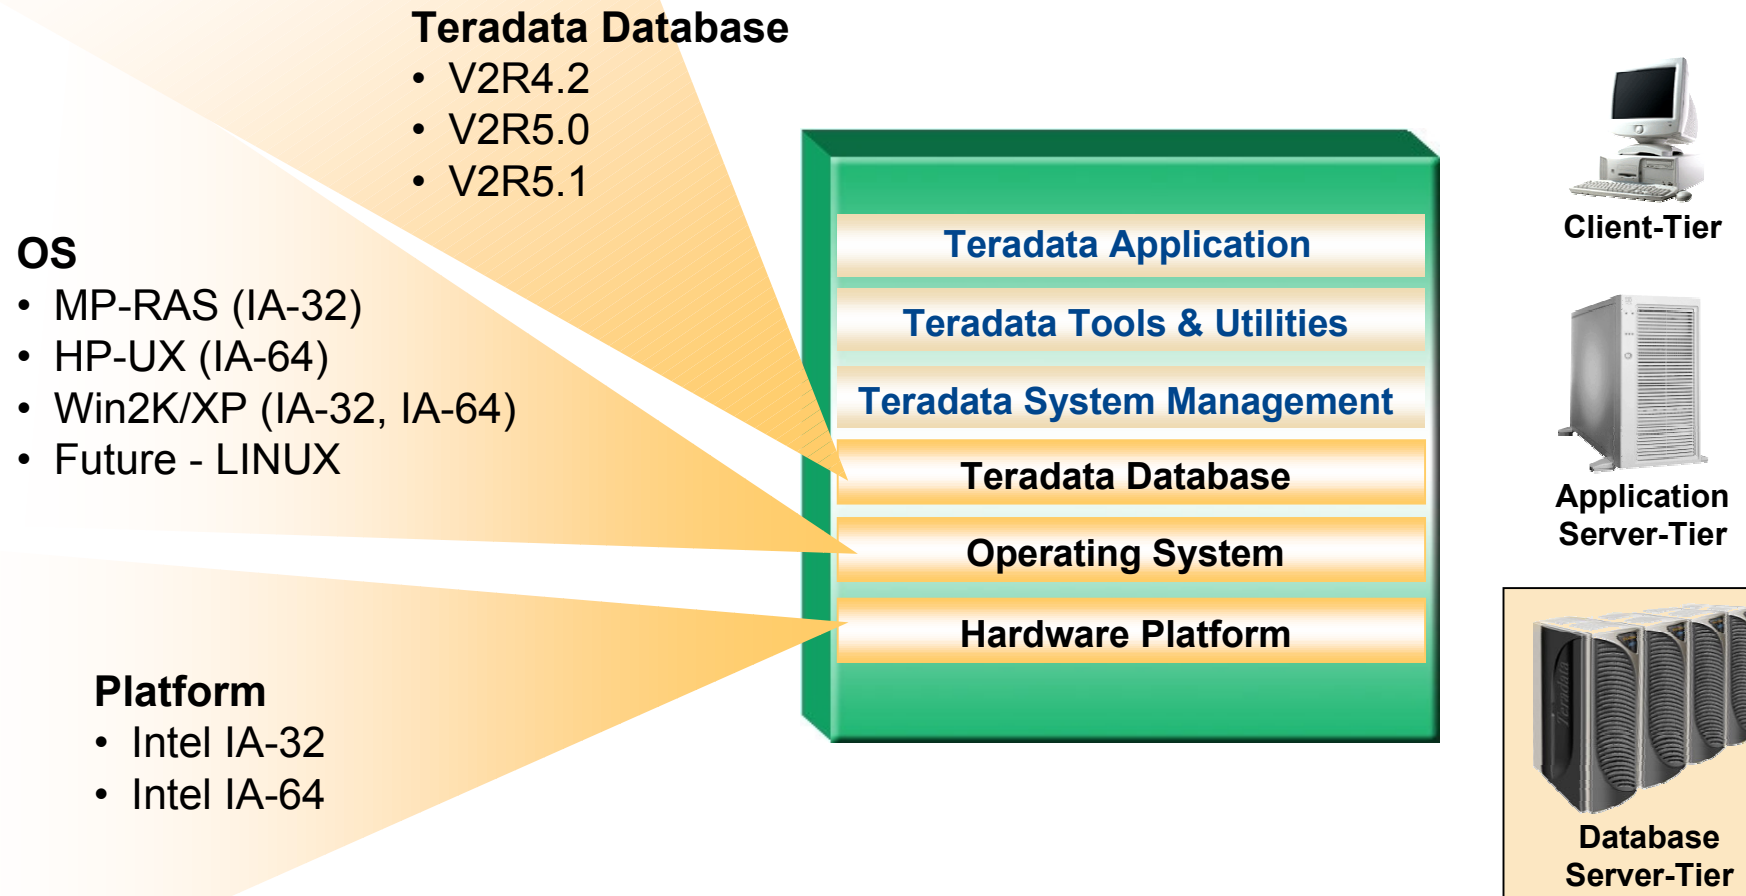


**Application
Server-Tier**



**Database
Server-Tier**

Teradata Warehouse Architecture




Teradata Warehouse Architecture

System Management

- Teradata Manager
- Teradata Dynamic Query Manager
- Teradata System Emulation Tool
- Teradata Visual Explain
- Teradata Index Wizard
- Teradata Statistics Wizard

Metadata

- Teradata Meta Data Services



Teradata Warehouse Architecture

Load/Unload

- Teradata Warehouse Builder
- FastLoad, MultiLoad & FastExport
- Teradata TPump
- Access Modules

Teradata Utility Pak



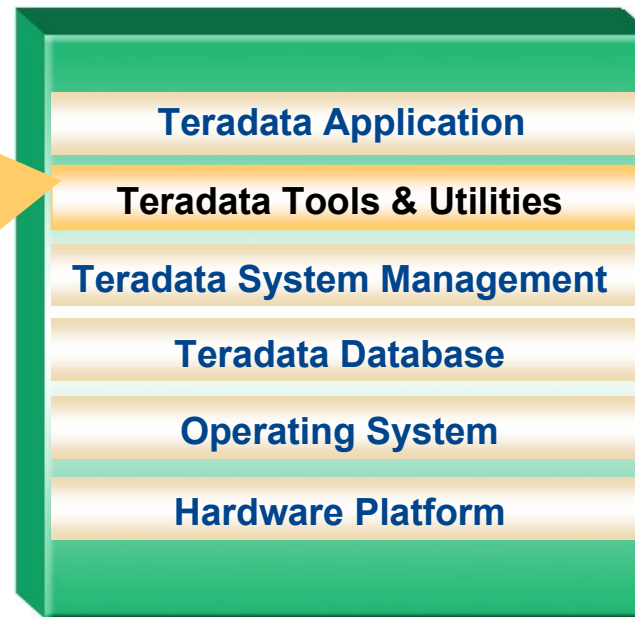
- Teradata Administrator, BTEQ
- Teradata SQL Assistant
- Teradata SQL Assistant/Web Edition
- ODBC, JDBC, CLI, OLE DB Provider

Mainframe Connectivity

- Mainframe Channel Connect
- TS/API, CICS, HUTCNS & IMS/DC

Backup and Restore


- NetVault
- ASF2 Reader



Teradata Warehouse Architecture

Applications

- CRM
- Supply Chain Management
- Sagetree
- Value Analyzer
- LifeTime Value
- Teradata Warehouse Miner



Client-Tier

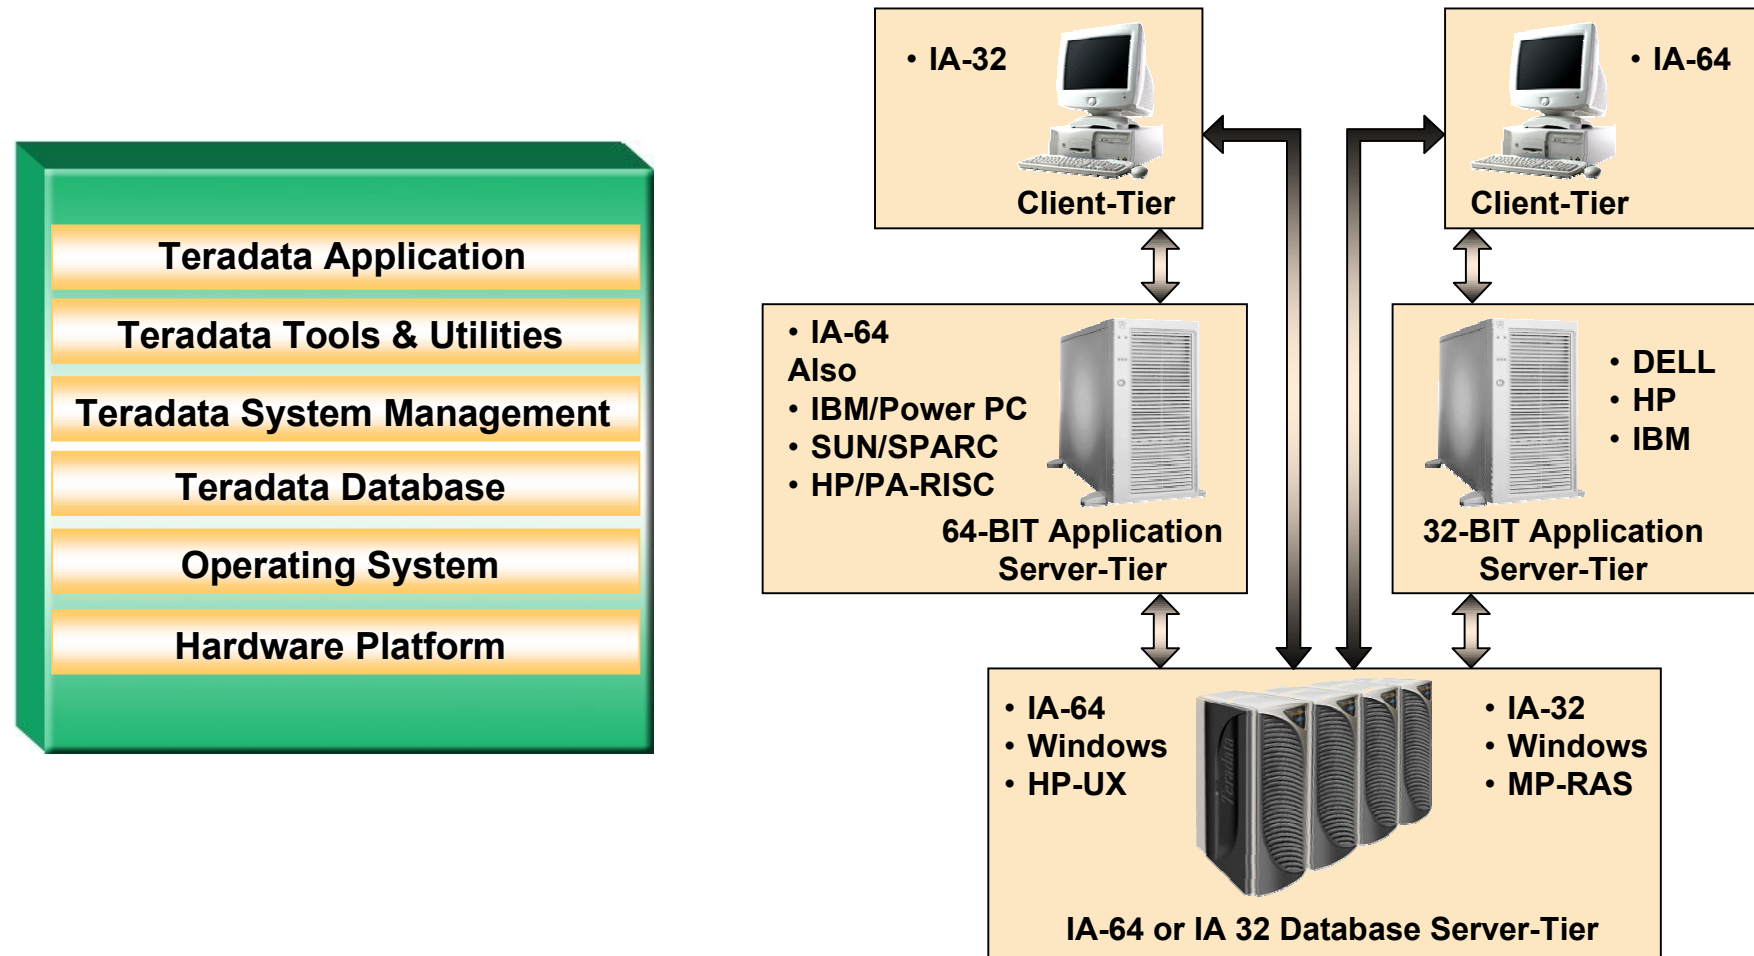


**Application
Server-Tier**



**Database
Server-Tier**


Teradata Warehouse Architecture



Teradata Database supports both 32-bit and 64-bit applications & clients concurrently.

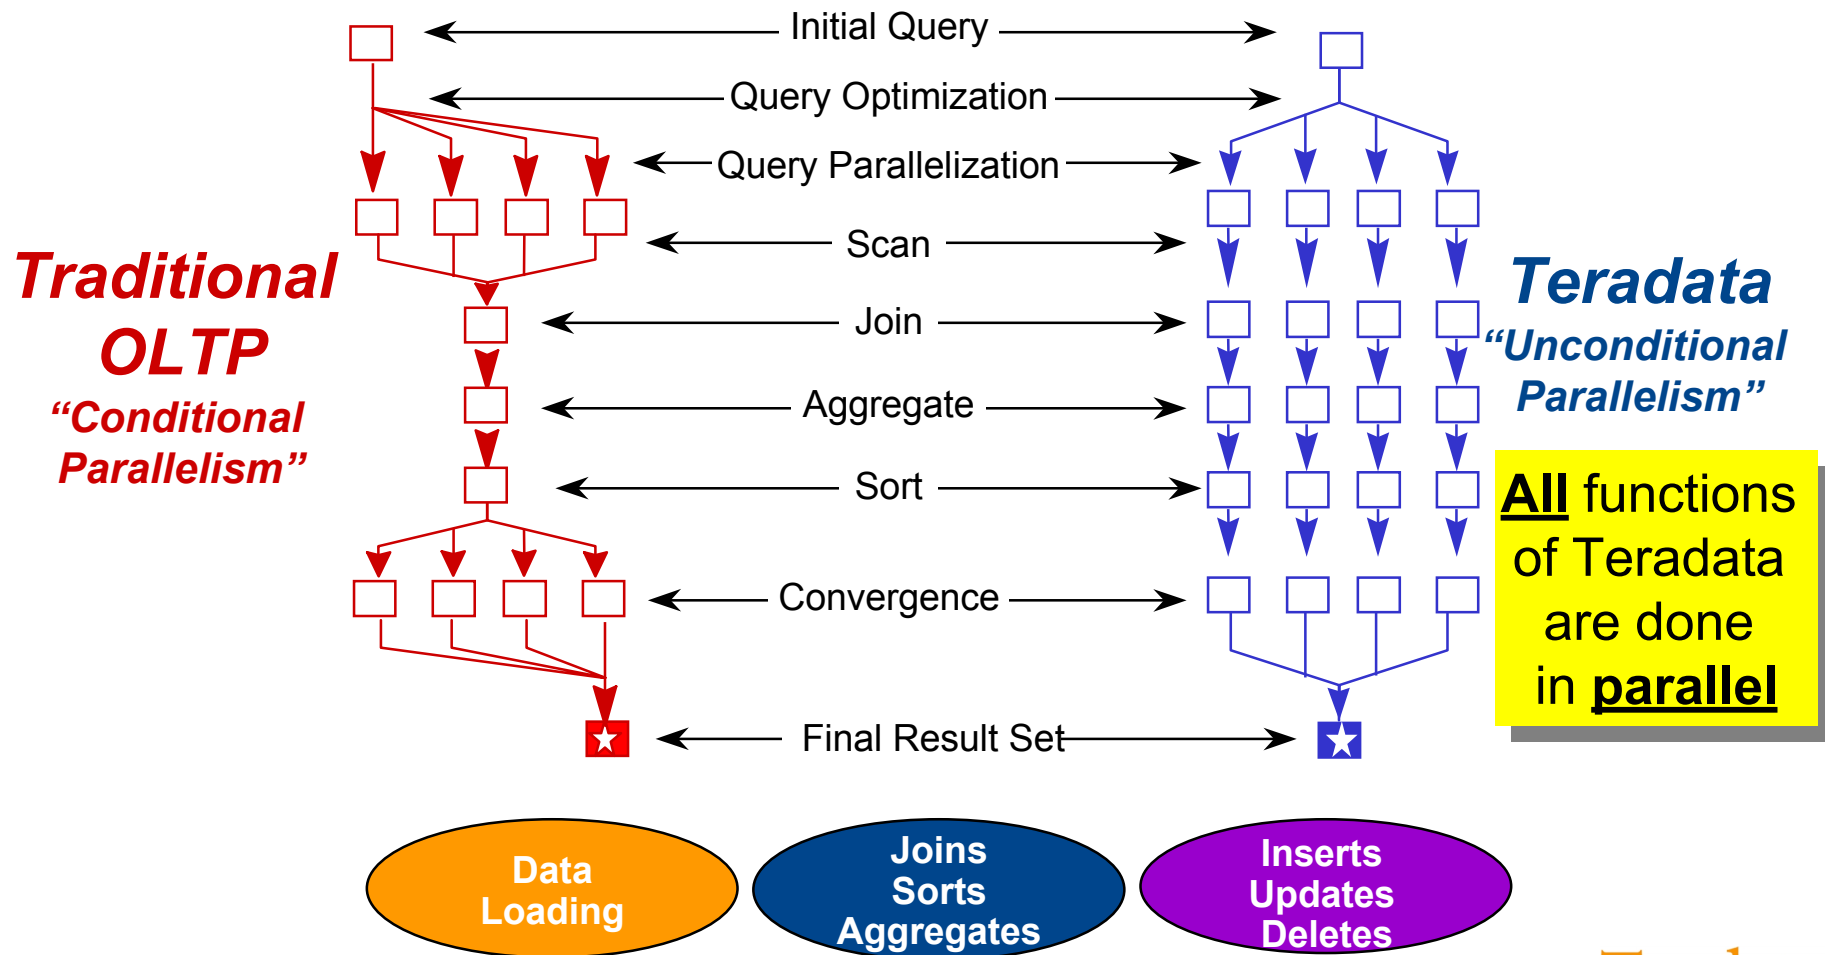
Data Warehousing Fundamentals: Parallel versus Single Workload

Counting Aces




Technical Issues of Parallelism

Each UoP (Unit of Parallelism) does an equal amount of work



Teradata Platform Architecture

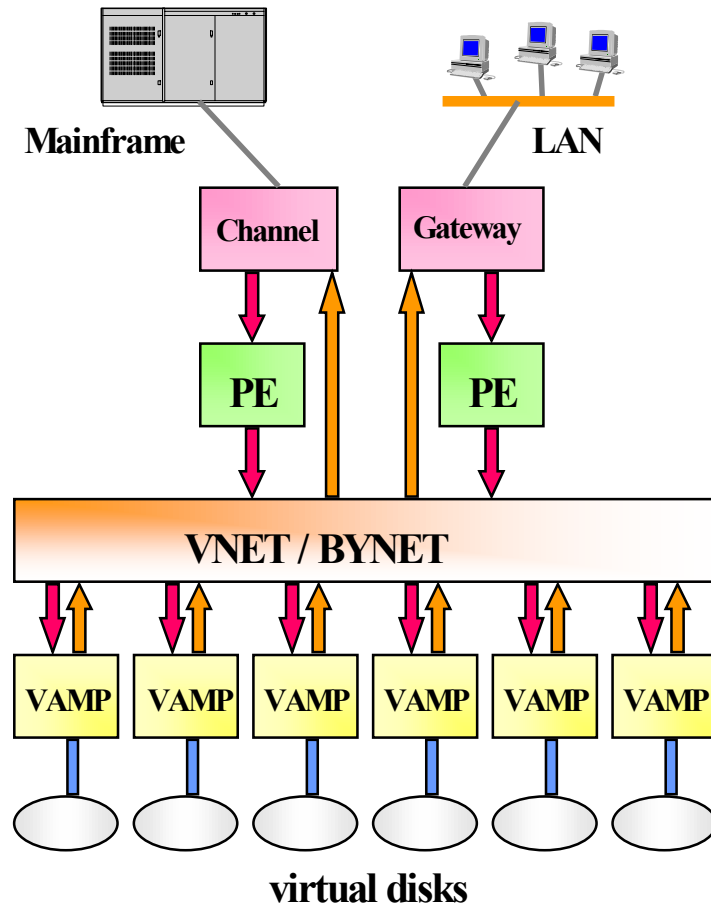
- **Shared Nothing** enables **hardware scalability**
- Teradata's **Parallel Everything** enables **software scalability** by eliminating single points of control at all levels
- **Massive Parallel Processing** (MPP)



**Degree of Parallelism = Number of Access Module Processors (AMPs),
Number of parallel units independent from tables or queries**

Teradata is “parallel aware” - Teradata always use all units of parallelism

Teradata Software Architecture



Teradata RDBMS = Parallel Shared Nothing Architecture

> PE: Parsing Engine

- Session control
- Parser, Optimizer & Dispatcher

> VNET / BYNET:

Communication between VPROCs

- Control- & Data messages
- Synchronization, Sort & Merge

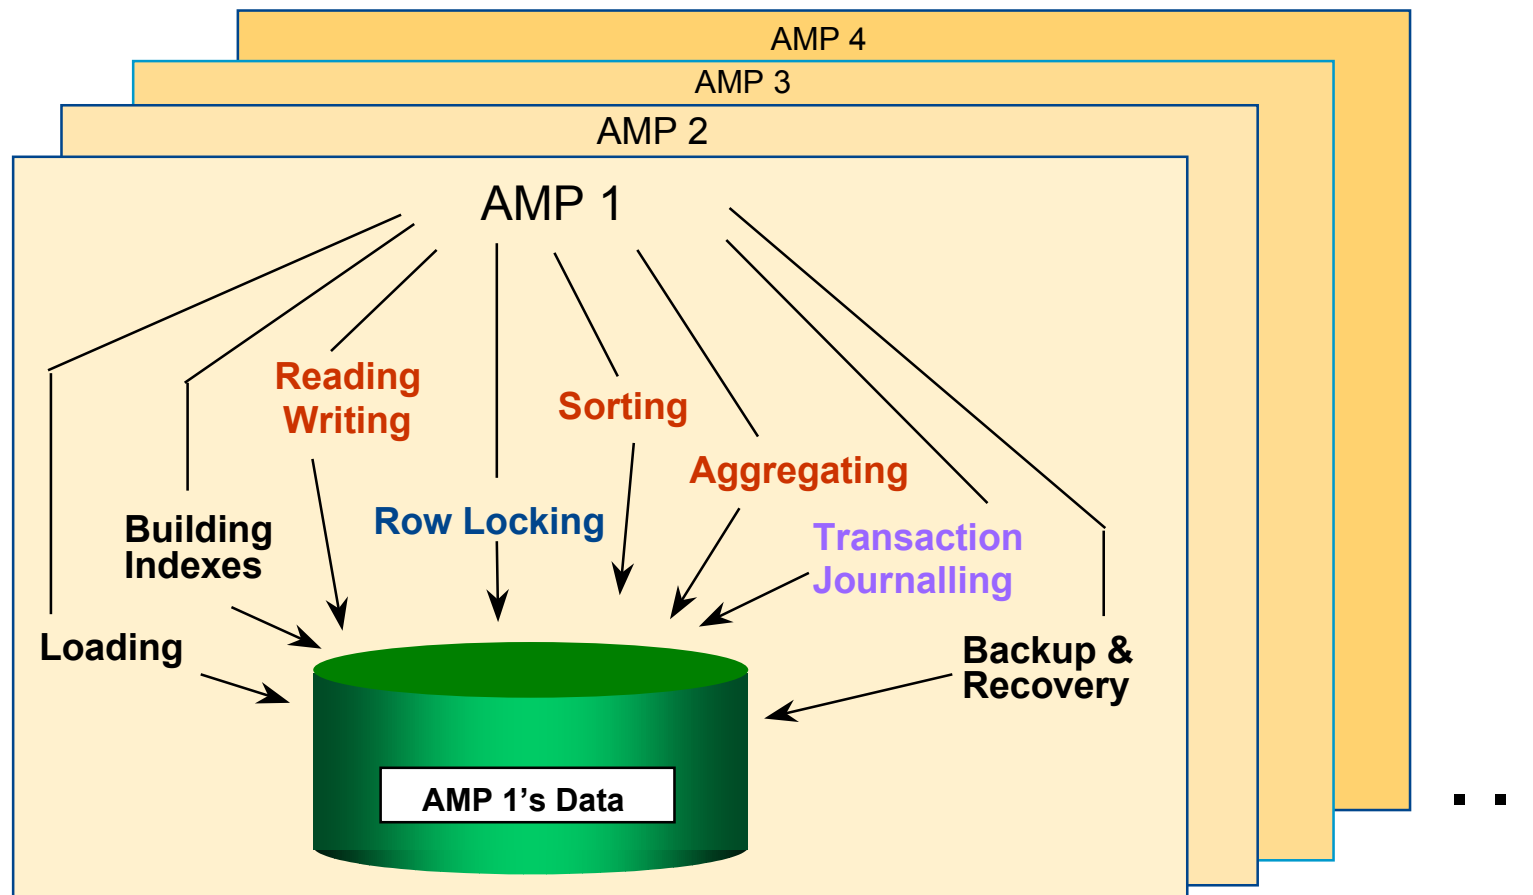
> VAMP: Virtual Access Module Processor

each AMP owns a portion of the database

- Data access, Insert, Update, Delete
- Cache management
- Journaling, Backup & Recovery
- Concurrency control & Locking

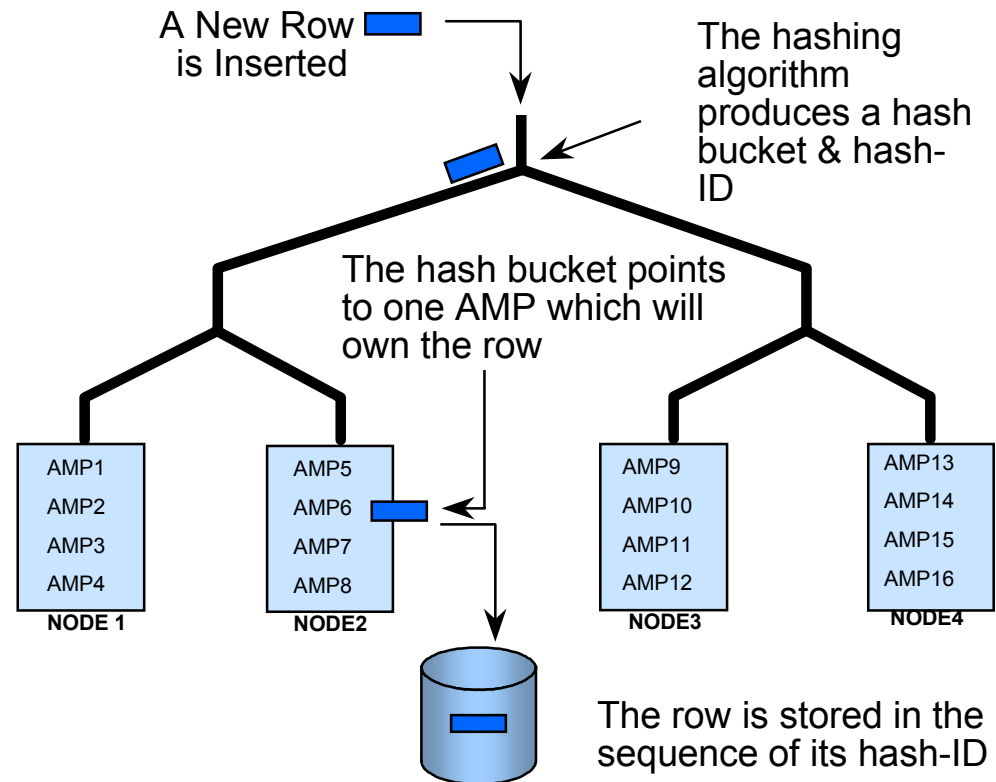
The Parallel Foundation

- Each parallel unit (AMP) owns and manages its own data



Balanced Data Placement

- All tables are partitioned on their primary index
 - > Hashing ensures an even spread of data
 - > Minimizes unbalanced work across parallel units



“Teradata is unique in “hashing” data directly to a physical address on disk, creating a potentially even distribution of data. This allows a balanced application of parallelism, and also avoids imbalance due to data skew.” -- Gartner Group

Teradata Data Management Illustration


- Random, automatic data distribution & placement
- Real-time, automatic data reorganization

*With Teradata
there is no
sense of data
ORDER,
therefore there
is no sense of
DISORDER,
eliminating the
need to
REORDER!*

AMP w/o Partitioning



AMP with Partitioning




■ SYSTEM TABLES
■ CUSTOMER
■ ORDERS
■ LINEITEM
■ PART
■ SUPPLIERS

*SELECT ... WHERE order_date
BETWEEN DATE '2002-10-1' AND DATE '2002-10-31'*


Teradata Bynet Interconnect

BYNET works much like a telephone network, where many callers can establish connections, including conference calls and broadcast connections

Point-to-Point Messaging



Broadcast Messaging



Specialised Services

for Teradata:

Message passing
Synchronisation
Merge


Other services:

TCP/IP

Bandwidth grows with number of nodes connected (up to 512 nodes)

Background: MPP Join Geography

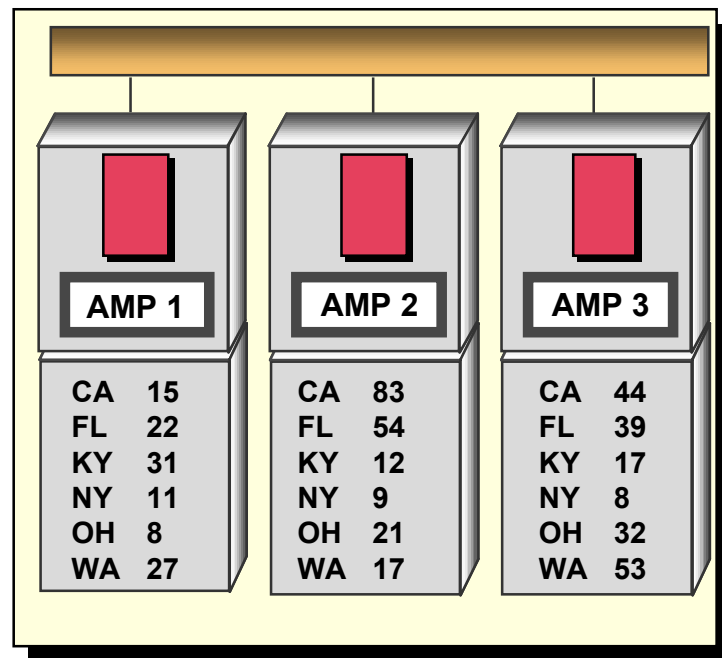
- Table Redistribution
 - Item rows are moved to the AMP where their related Orders are located
- Table Duplication
 - Smaller table is duplicated to each AMP
- AMP-local join
 - No movement: Both tables' related rows are already on the same AMP



**Prerequisite is a scalable and very fast
node interconnection - BYNET**

Fully-Parallelized Aggregations

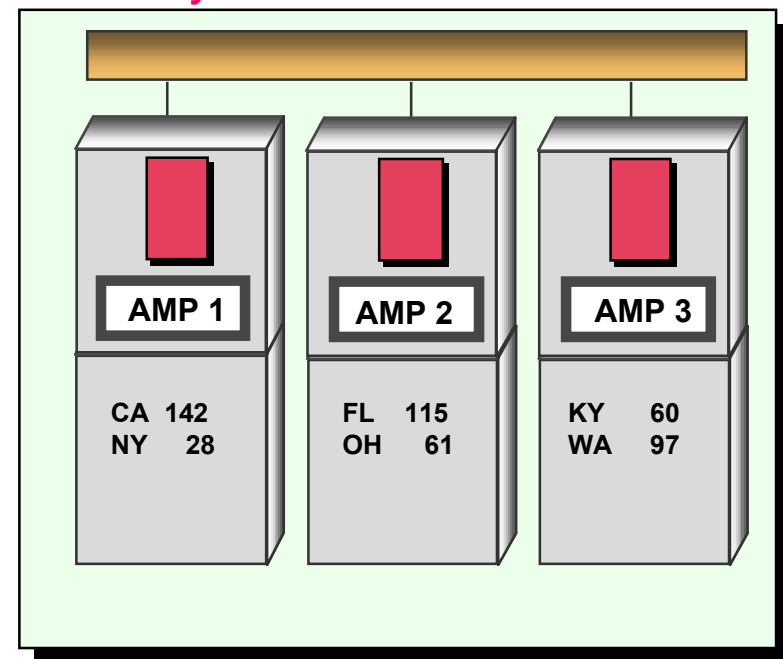
- Each AMP performs a local count, then executes a share of the global count



Step 1 - Local Sub-totals



A Count of the Number of Delinquent Customers by State



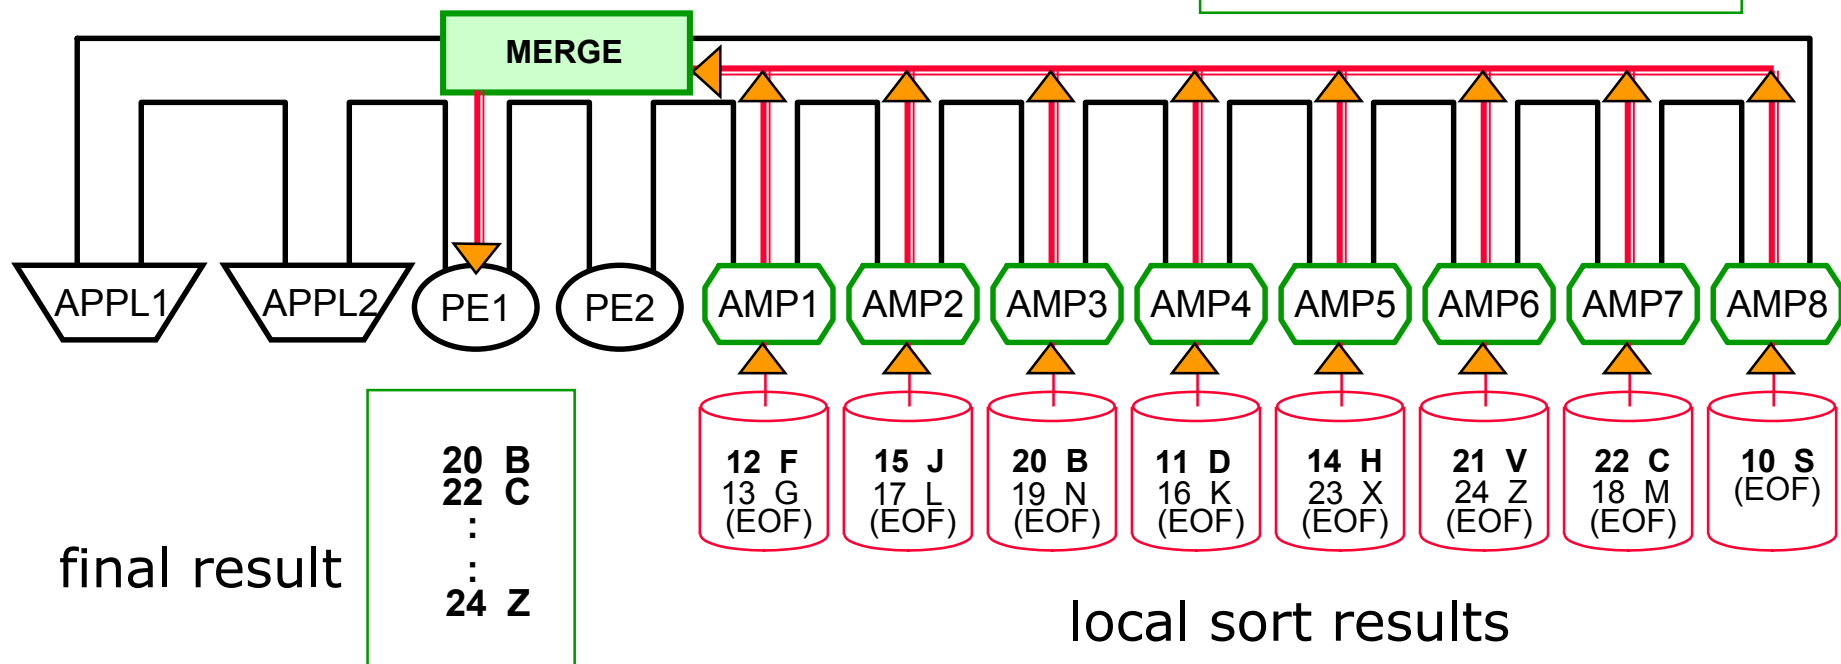
Step 2 - Global-Grand-total

ALL aggregations in parallel - No single node bottleneck

Fully-Parallelized Sorts

Each AMP performs a local sort, then BYNET executes a final sort/merge on-the-fly as the rows are returned to the user

```
SELECT NUMBER, LETTER
FROM SAMPLE
WHERE NUMBER > 9
ORDER BY LETTER
;
```




ALL sorts in parallel - No single node bottleneck

Teradata Optimizer: What does it Optimize?

- **Access Path:** Method of accessing each table
 - > Table Scan, Index Use, Bitmap Use
- **Join Method:** How pairs of table are joined
 - > Merge Join, Product Join, Hash Join
- **Join Geography:** How rows are relocated prior to the join
 - > Redistribute Rows, Duplicate Rows
- **Join Order:** Sequence of table joins
 - > 5 table look ahead, pick the cheapest

Teradata Optimizer: What it Needs to Know?

- **Environment information**
 - > Number of Nodes
 - > Number and type of CPUs
 - > Disk Array Information
 - > Interconnect Information
 - > Amount of Memory Available
 - > Number of Virtual AMPs
- **Statistics**
 - > Table Cardinality
 - > Column Demographics
 - > Collected by user
 - > Random Sampling




Database Space Management

- **Space allocation is entirely dynamic**
 - > No tablespaces or journal spaces or any pre-allocation
 - > Spool (temp) and tables share space pool, no fixed reserved allocations
- **If no cylinder free, combine partial cylinders**
 - > Dynamic and automatic
 - > Background compaction based on tunable threshold
- **Quotas control disk space utilization**
 - > Increase quota (trivial online command) to allow user to use more space

Teradata Scalability

Add more data, more users, and more subjects to your data warehouse with predictable performance




e.g. double number of nodes:

process **twice** the workload in same amount of time,
or process same workload in **half** amount of time

Availability (I)

Fault Tolerant Hardware


- **Node**
 - > Dual AC power
 - > Redundant power supplies & fans
 - > Hot plug components
- **Storage**
 - > RAID 1 protection
 - > Dual I/O for all paths to disk
 - > "Cliques" cross connection of nodes to disk arrays for fail-over
- **BYNET Interconnect**
 - > Dual networks
 - > Fault tolerant within each network
 - > Automatic network reconfiguration
- **Connectivity**
 - > Multiple Redundant Connections
- **Server Management**
 - > Optional Dual Configuration



Availability (II)


High Availability Software

- Virtual Processor Migration
 - > Move work to functional resources
 - > Node failure protection
 - > OS failure protection
- Parallel Recovery & Rollback
- On-Line Utilities
 - > Load, Export, Backup
 - > Purge
 - > Checkpoint / Restart
- Elimination of many operations
 - > Reorganization
 - > Index Rebuilds
- Fallback
 - > Covers even catastrophic failures




Teradata Performance

- Synchronized Table Scan
- Cylinder Read
- Managing the Flow of Work
 - Priority Scheduler
 - Dynamic Resource Prioritization
 - Teradata Dynamic Query Manager (TDQM)
 - Proactive Query Management
- Advanced Indexing
 - Join Indexes
 - Summary Tables
- Event-based Handling
 - Triggers



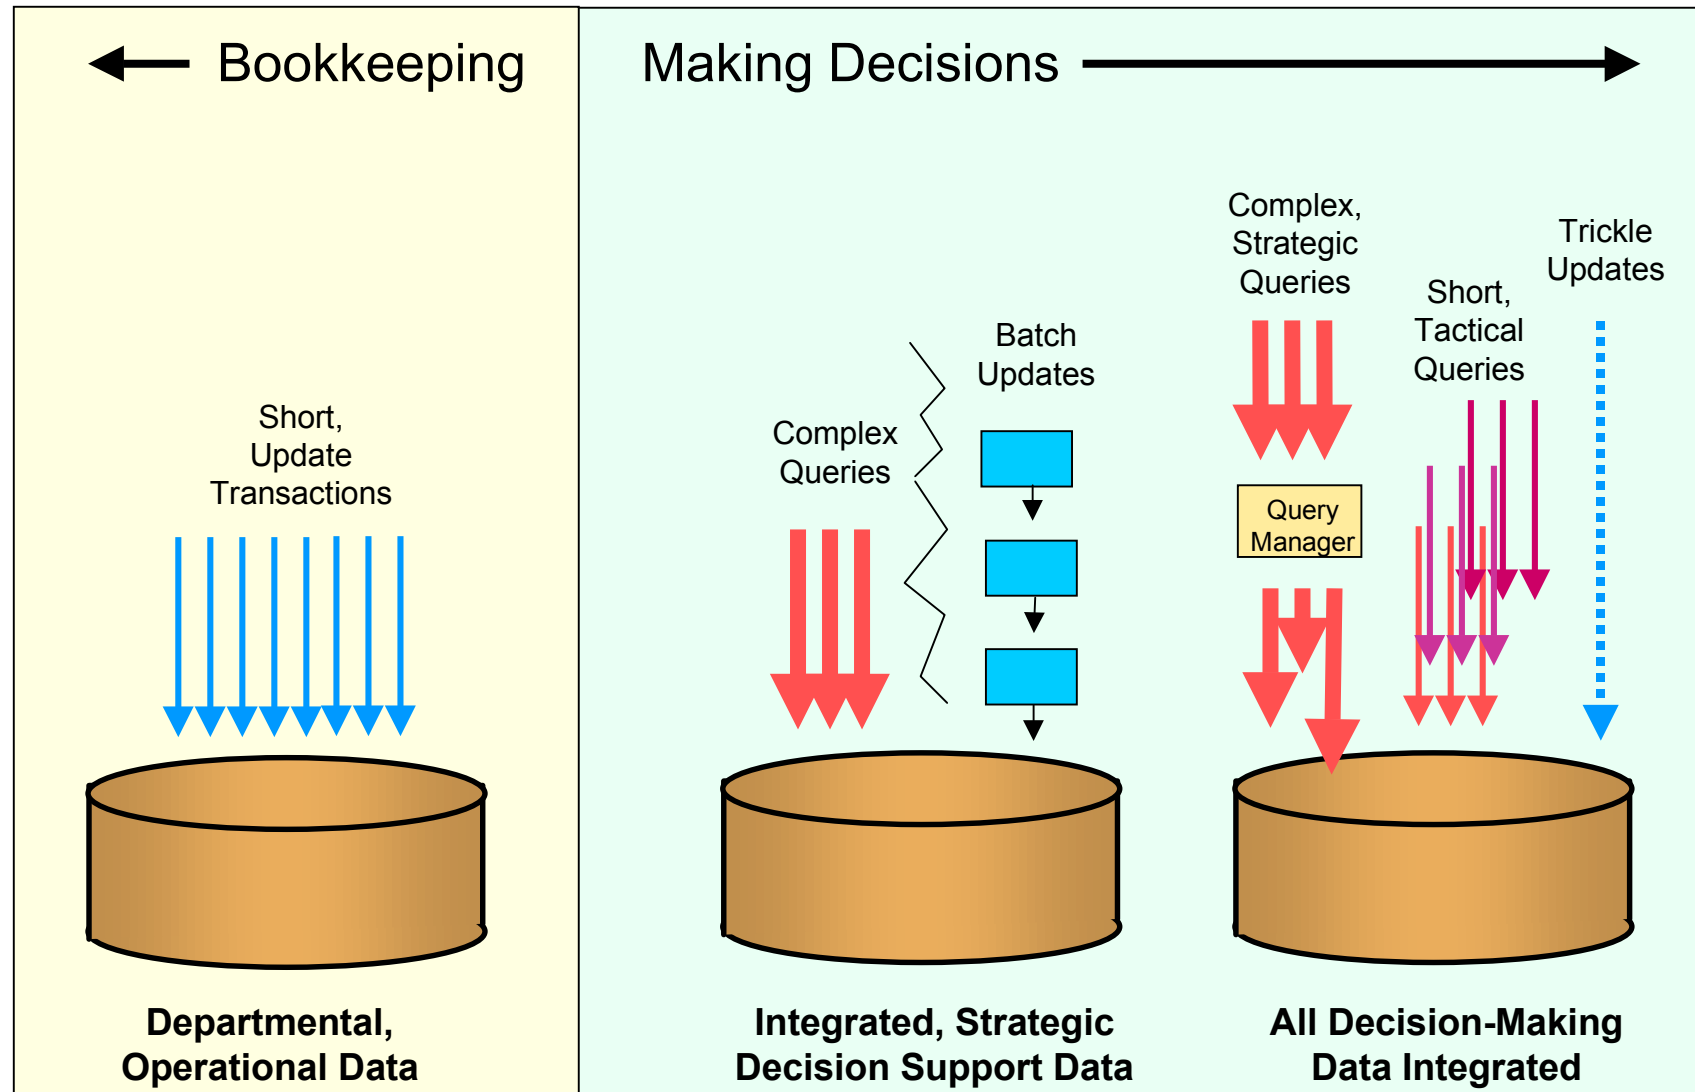
Synchronized Table Scan & Joins

- Allows multiple simultaneous scans & joins against a single table to share data blocks
- A new query joins the scan / join at the current scan point

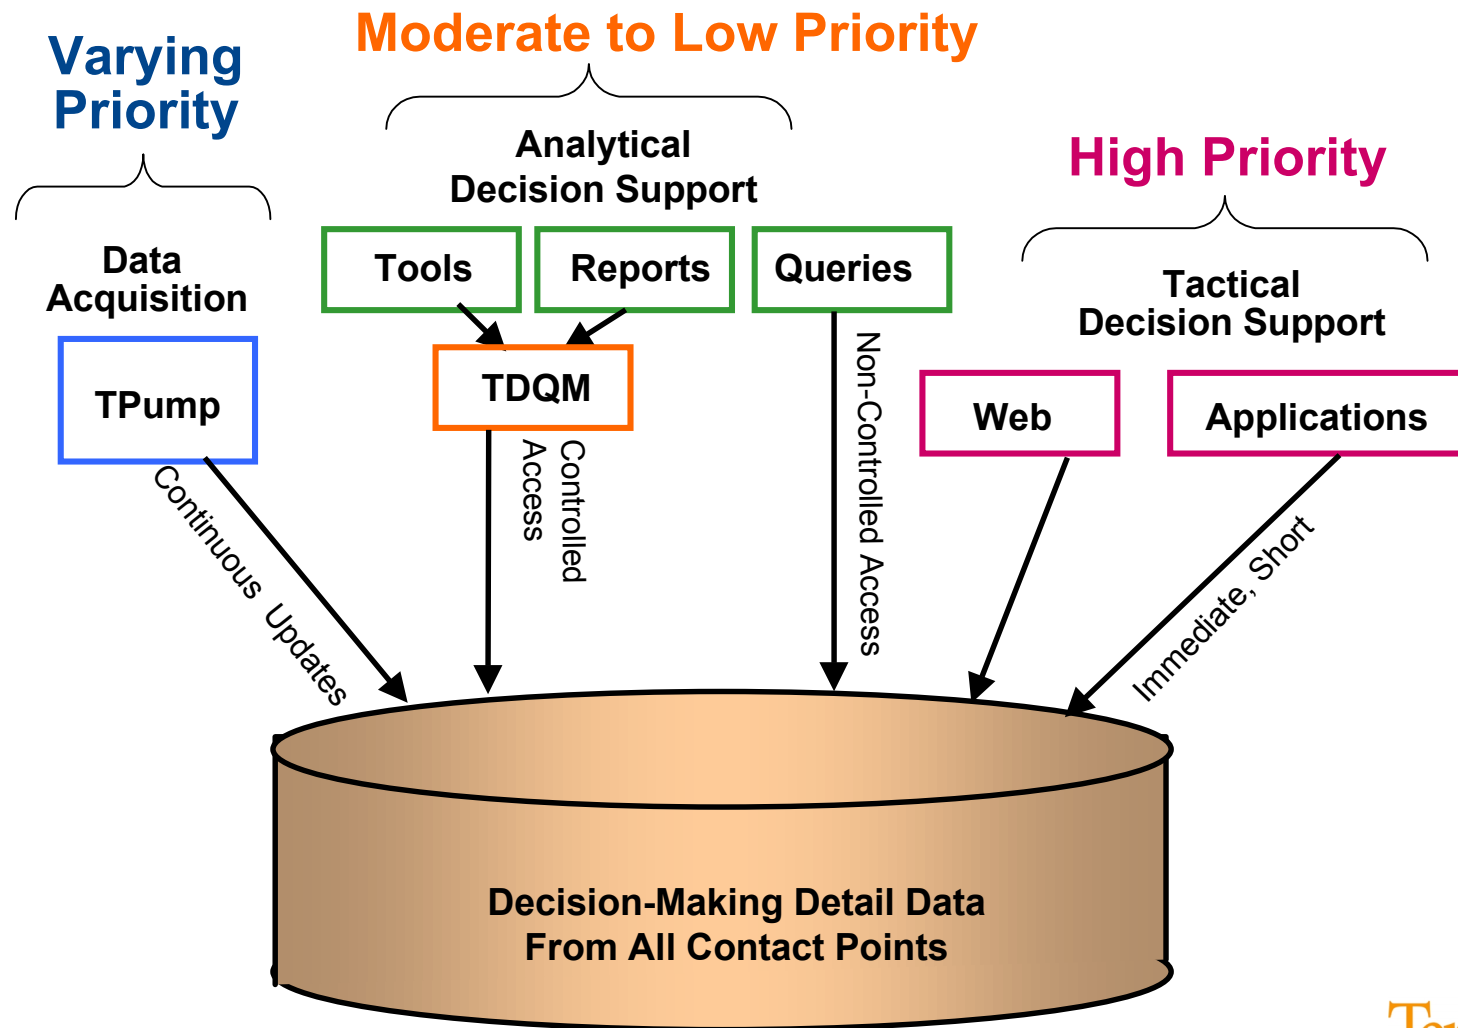


**Reduces significantly the query response time
in a multiuser environment**

Contrasting OLTP, Traditional, and Active Data Warehousing




Managing Diverse Work Inside Teradata: Priority Scheduler



Teradata Priority Scheduler

- Weights are distributed among Resource Partitions
 - RP0 100
 - RP1 80
 - RP2 20
- Weights are then allocated among Allocation Groups within
- Resource Partitions
 - RP0/L 5
 - RP0/M 10
 - RP0/H 20
 - RP0/R 40
 - RP1/L 2
 - RP1/M 10
 - RP1/H 50
 - RP1/R 100
 - RP2/L 10



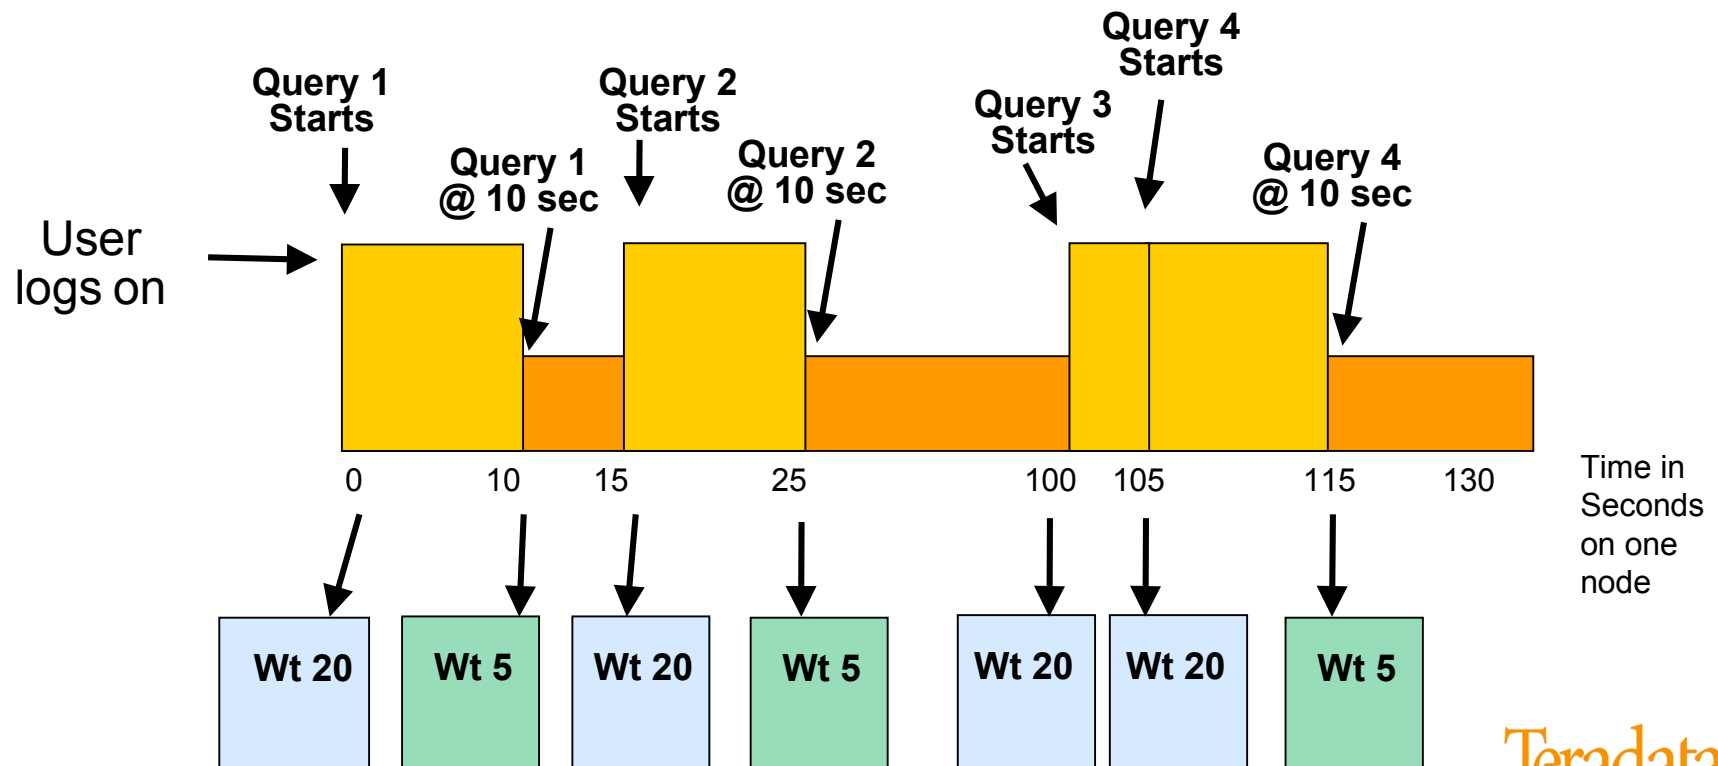
**Relative Weighting
and Allocating Resources**

Priority Scheduler Query Milestone

- Usage


- > Example of how it works follows:

- In this example, the Q milestone limit is 10 seconds
 - All queries will receive a weight of 20 until they consume 10 seconds of CPU on that node
 - After 10 seconds they will receive a weight of 5



Data Freshness

- Choose the load-utility for the type of application:
 - ➔ **Fastload, Multiload, TPump or Warehouse Builder**
 - Fastload for efficient load into empty table
 - Multiload for concurrent load into several tables
 - TPump trickle load utility with record-level locking facilitates continuous data loading while table still being accessed
 - Teradata Warehouse Builder Extract-Transform-Load within a single job



Agenda

- Teradata Warehouse Architecture
- **Teradata Tools & Utilities**
 - > Functional Overview
 - > Load/Unload Utilities
 - > Teradata Utility Pak
 - > DB Management Utilities
 - > Mainframe Connectivity
- Teradata Warehouse 7.0 Demo System (Demonstration)

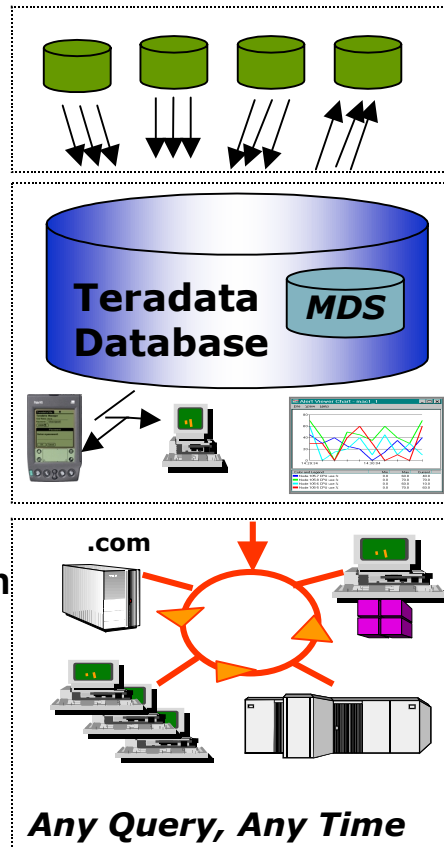
Teradata Tools & Utilities

Load/Unload

Teradata Warehouse Builder
Teradata TPump
Access Modules
FastLoad, MultiLoad & FastExport

Teradata Utility Pak

Teradata Administrator
Teradata SQL Assistant
Teradata SQL Assistant/Web Edition
Teradata MultiTool
BTEQ
Teradata CLI
ODBC Driver for Teradata
JDBC Driver for Teradata
OLE DB Provider for Teradata



Database Management

Teradata Manager
Teradata Dynamic Query Manager
Teradata Analyst
Teradata System Emulation Tool
Teradata Visual Explain
Teradata Index Wizard
Teradata Statistics Wizard
Teradata Meta Data Services


Mainframe Connectivity

Mainframe Channel Connect
TS/API, CICS, HUTCNS & IMS/DC

Data Integration Requirements


Data Warehouses require fast, fail-safe, robust utilities to load and manage data because:

- The amount and frequency of data coming into the warehouse will grow exponentially
- Loads, extracts, inserts, updates, and deletes will occur on a regular basis
- Batch windows shrink as companies grow to span multiple time zones
- Users want their data fast!



Parallel Load & Unload Utilities

- **Teradata Utilities are Fully Parallel**
- **Teradata Utilities Provide Automatic Checkpoint Restart Capabilities**
- **Supports seamless data movement from Mainframes, Solaris, UNIX MP-RAS, AIX, HP-UX and MS Windows/NT.**
- **Data Loads Directly from the Source into the Database**
 - > No file splitting!
 - > No intermediary file transfers!
 - > No manual data conversion!



Conventional Teradata Load / Unload Utilities

Conventional Utilities		Teradata Warehouse Builder Operators	Purpose
FastLoad	➤	Load	Load Data into empty tables
MultiLoad	➤	Update	Update, insert and deletes rows
FastExport	➤	Export	Extracts data from Teradata
TPump	➤	Stream	Continuously loads data into Teradata
BTEQ Import	➤	SQL Inserter	Loads small amount of data into Teradata
BTEQ Export	➤	SQL Selector	Extracts small amount of data from Teradata
OLE DB Access Module	➤	ODBC	Reads data from ODBC compliant sources


Parallel Load & Unload Utilities

Universal Access


OLE DB Access Module

Named Pipes Access Module

WebSphere Access Module



Continuous Load: **Teradata TPump**



Parallel Load Environment


Teradata Warehouse Builder with

Data Connector,

ODBC Operator,

Load, Export, Update,

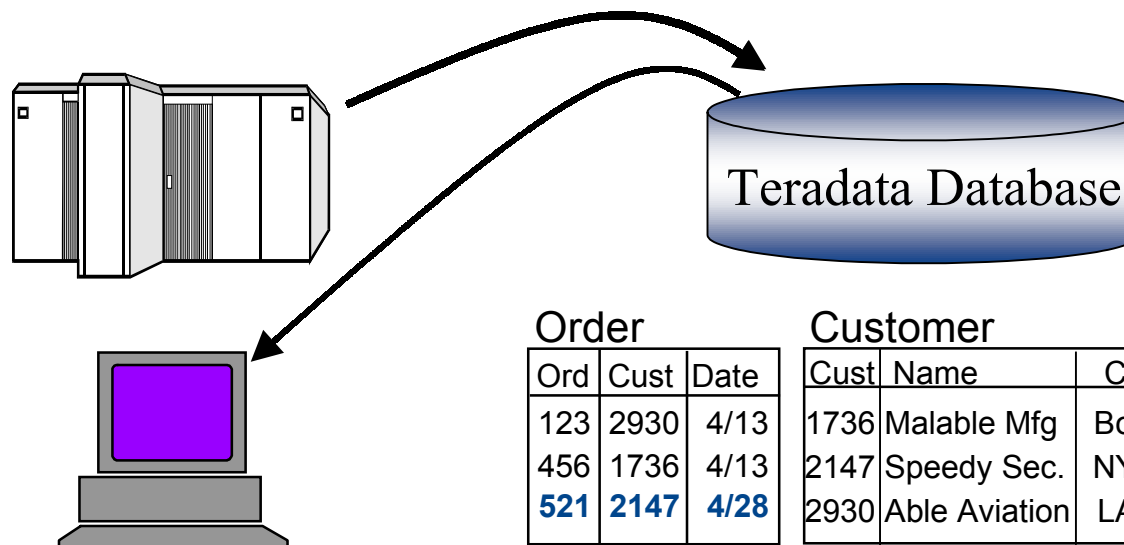
Stream Operators, and open API



Teradata Load & Unload Product Suite

Batch Load & Unload Utilities

- **FastLoad**
 - > High volume initial loads into empty tables
- **MultiLoad**
 - > High volume update, delete and insert of up to 5 tables within a single pass
- **FastExport**
 - > High volume extract of tables and views



Order

Ord	Cust	Date
123	2930	4/13
456	1736	4/13
521	2147	4/28

Customer


Cust	Name	City
1736	Malable Mfg	Boston
2147	Speedy Sec.	NYC
2930	Able Aviation	LA

Part

Part	Qty	Descr
345	925	8.5x11 Paper
360	935	8.5x14 Paper
421	120	#2 Pencil

Teradata TPump: Continuous Load

- Teradata load utility that allows continuous updates
 - > Continuous updates of data results in more accurate, timely data
- Updates without locking tables
 - > Enables concurrent loads to a single table
 - > Allows users to access the table during updates
- Alternative to MultiLoad for low-batch maintenance of large databases



Order

Ord	Cust	Date
123	2930	4/13
456	1736	4/13
521	2147	4/28

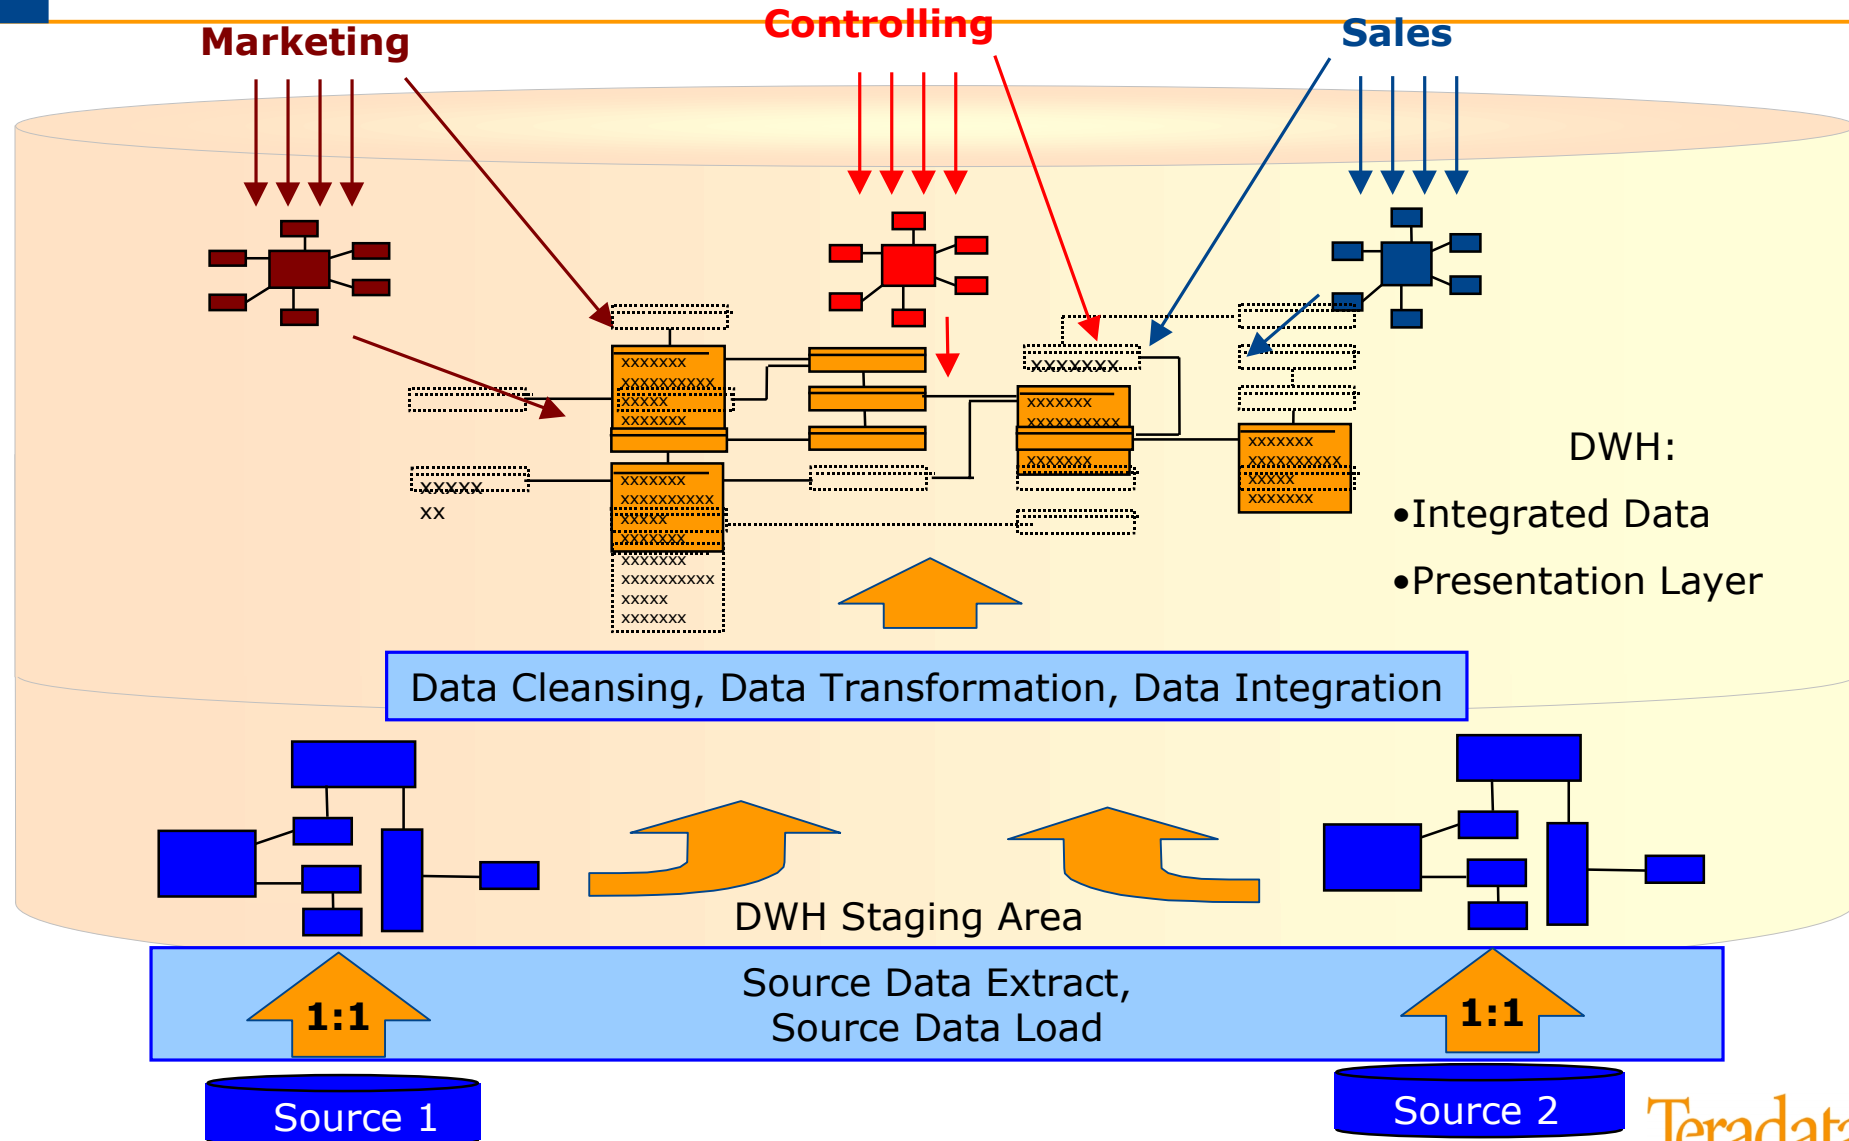
Customer

Cust	Name	City
1736	Malable Mfg	Boston
2147	Speedy Sec.	NYC
2930	Able Aviation	LA

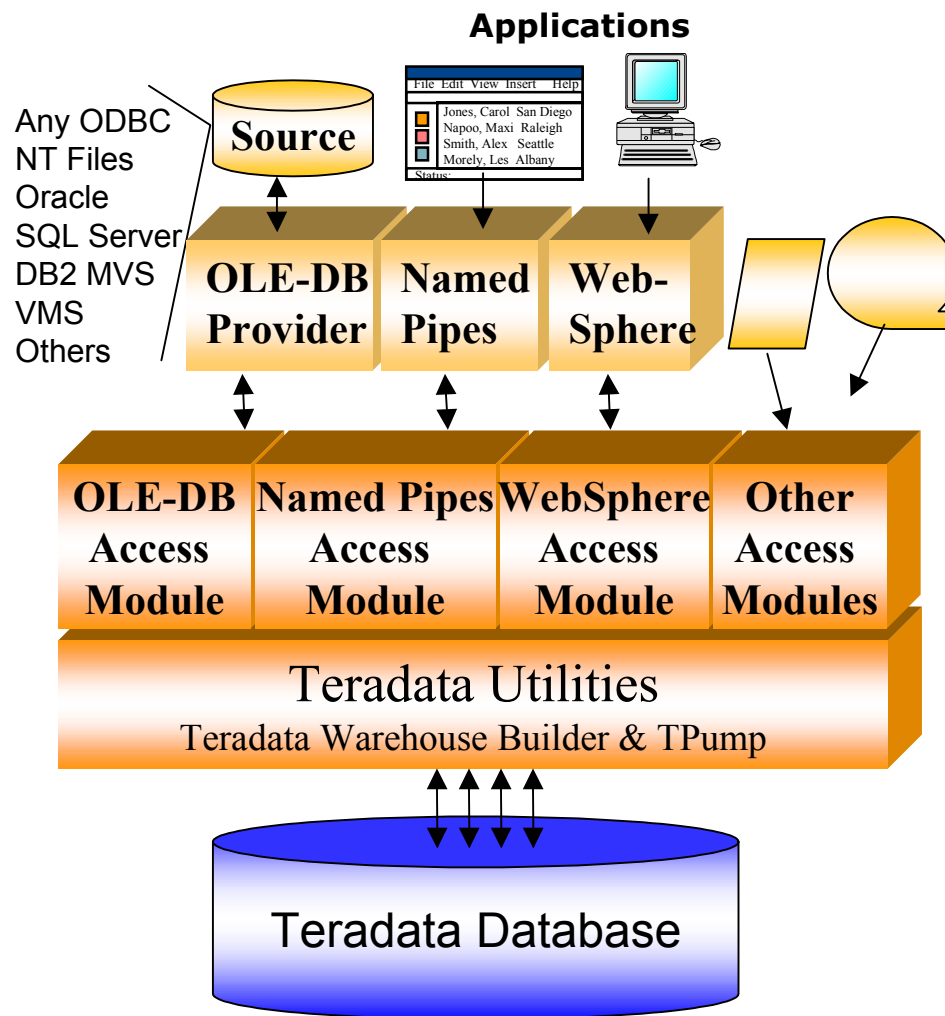
Part

Part	Qty	Descr
345	925	8.5x11 Paper
360	970	8.5x14 Paper
421	120	#2 Pencil

DWH Load Strategies – ELT



Teradata Load Access Modules



- Specialized interface to Client Utilities
- Data sources
 - > Named Pipes
 - Reads data from MP-RAS UNIX, Solaris or NT Named Pipes.
 - Enables user applications to continuously feed data to the Teradata database
 - > WebSphereMQ
 - Reads data from IBM WebSphere queues


Teradata Tools & Utilities

Load/Unload

Teradata Warehouse Builder
Teradata TPump
Access Modules
FastLoad, MultiLoad & FastExport

Teradata Utility Pak

Teradata Administrator
Teradata SQL Assistant
Teradata SQL Assistant/Web Edition
Teradata MultiTool
BTEQ
Teradata CLI
ODBC Driver for Teradata
JDBC Driver for Teradata
OLE DB Provider for Teradata



Database Management

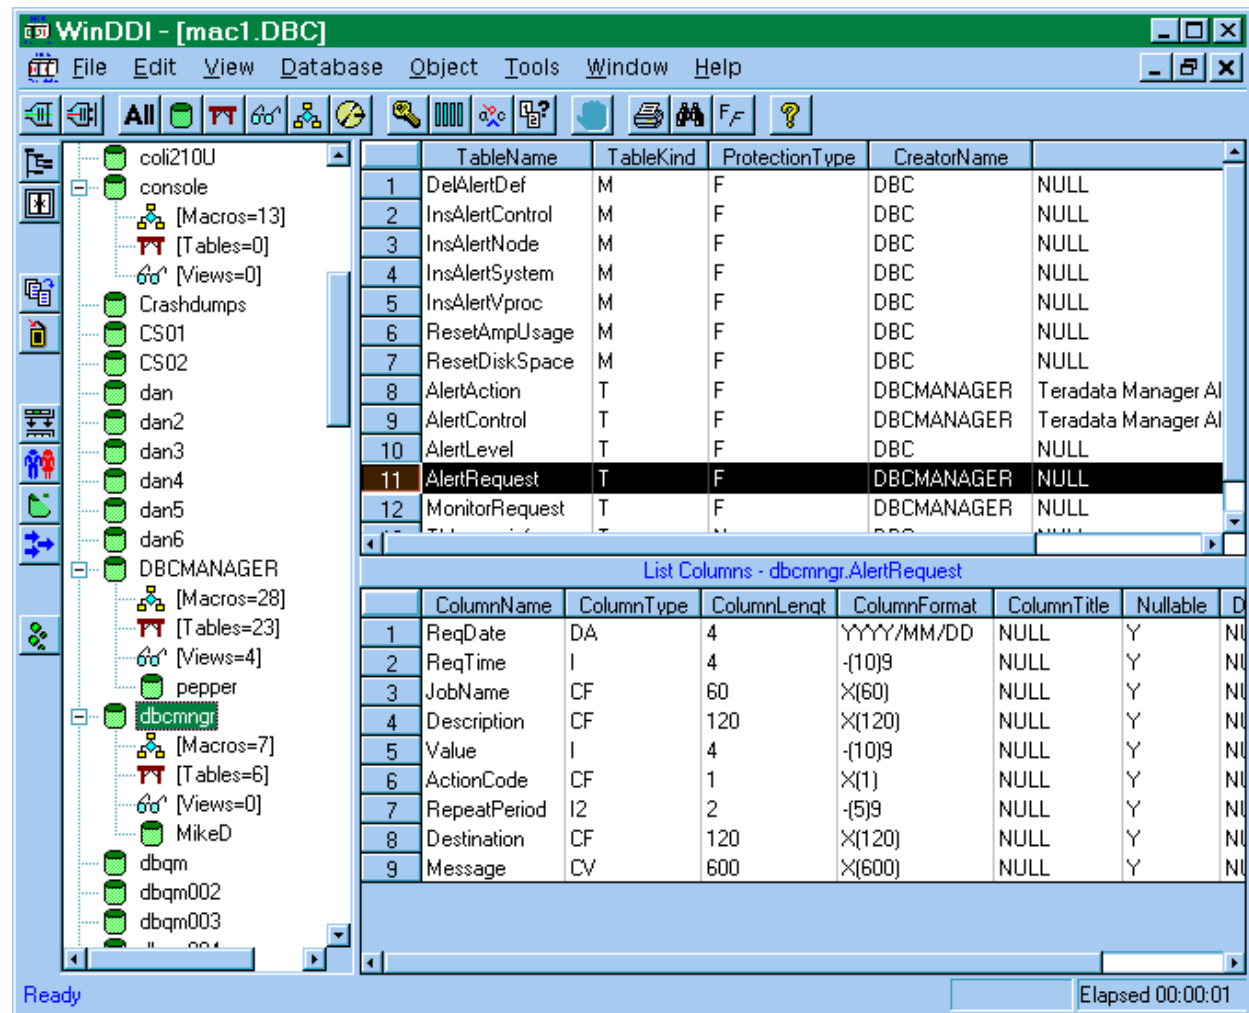
Teradata Manager
Teradata Dynamic Query Manager
Teradata Analyst
Teradata System Emulation Tool
Teradata Visual Explain
Teradata Index Wizard
Teradata Statistics Wizard
Teradata Meta Data Services

Mainframe Connectivity

Mainframe Channel Connect
TS/API, CICS, HUTCNS & IMS/DC



Teradata Administrator (WinDDI)

- Manage the database
- Perform typical database operations with a click of the mouse
- History of SQL operations



Teradata SQL Assistant (QueryMan)

- Windows-based SQL Query Tool
- Exports data from any ODBC enabled database to PC
- Retains historical records of submitted SQL with timings and status



Any Query, Any Database, Anytime

teradata
a division of  NCR

BTEQ: Basic Teradata Query

- **Batch interface for small loads**
 - > example: few row inserts or updates
 - > Transformations via SQL (ELT)
- **General-purpose Query and Reporting Tool**
 - > Report Formatting
 - > Ad Hoc Queries
- **Process BTEQ Commands for Database Administration**




Teradata Connectivity Tools

- **ODBC Driver for Teradata**
 - > Open Database Connectivity
 - > Open, standards-based Teradata RDBMS access
- **OLE DB Provider for Teradata**
 - > Object Linked Embedded
 - > Based on Microsoft industry standard
- **JDBC Driver for Teradata**
 - > Java Database Connectivity
 - > Teradata access for web-enabled Java applications
- **Teradata CLI**
 - > Call Level Interface
 - > High performance parallel native interface to the Teradata database


Open Database Connectivity: ODBC Driver for Teradata

- Based on industry standards defined by Microsoft.
- Maximizes application portability and staff expertise and training
- Available for Microsoft Windows 98/NT/2000/XP, MP-RAS, HP-UX, Solaris SPARC and AIX clients.



Object Linked Embedded: OLE DB Provider for Teradata

- Based on industry standards defined by Microsoft
- Native OLE DB Provider for Teradata enables high speed access to the Teradata database from any OLE DB compliant source



Agenda

- Teradata Warehouse Architecture
- Teradata Tools & Utilities
- **Teradata Warehouse 7.0 Demo System (Demonstration)**
 - > Setup & Documentation
 - > TD Administrator
 - > SQL Assistant
 - > Fastload
 - > Bteq

Fastload Sample (Projektpraktikum)


```
logon demotdat/dbc,password ;
database TUWIEN;
drop table EMPLOYEE_stage ;
drop table EMPLOYEE_ERR1 ;
drop table EMPLOYEE_ERR2 ;
CREATE SET TABLE EMPLOYEE_stage ,NO FALLBACK ,
      NO BEFORE JOURNAL,
      NO AFTER JOURNAL
      (
        EmpNo    VARCHAR(15) CHARACTER SET LATIN NOT CASESPECIFIC NOT NULL
      ,Name      VARCHAR(18) CHARACTER SET LATIN NOT CASESPECIFIC NOT NULL
      ,Address   VARCHAR(40) CHARACTER SET LATIN NOT CASESPECIFIC NOT NULL
      ,Phone     VARCHAR(35) CHARACTER SET LATIN NOT CASESPECIFIC NOT NULL
      ,DeptNo    VARCHAR(10)
      ,Salary    VARCHAR(15)
      ,YrsExp    VARCHAR(15)
      ,DOB       VARCHAR(10) /*DATE FORMAT 'YYYY-MM-DD'*/
      ,MedStat   VARCHAR(10) CHARACTER SET LATIN NOT CASESPECIFIC
      ,EdLev     VARCHAR(10)
      ,Note      VARCHAR(79) CHARACTER SET LATIN NOT CASESPECIFIC NOT NULL
      )
UNIQUE PRIMARY INDEX ( EmpNo ) ;
SET RECORD      VARTEXT ";"
      DISPLAY_ERRORS ;
DEFINE FILE =C:\TU_Wien\employee.txt ;
RECORD 2 ;
BEGIN LOADING EMPLOYEE_stage
      ERRORFILES EMPLOYEE_ERR1, EMPLOYEE_ERR2 ;
INSERT EMPLOYEE_stage.* ;
END LOADING ;
QUIT ;
```

Bteq Sample (Projektpraktikum)

```
logon demotdat/dbc,password ;
sel date, time ;
database TUWIEN;
help database TUWIEN;
show table employee_stage ;
drop table TUWIEN.EMPLOYEE;
CREATE SET TABLE TUWIEN.EMPLOYEE ,NO FALLBACK ,
      NO BEFORE JOURNAL,
      NO AFTER JOURNAL
      (
        EmpNo    VARCHAR(15) CHARACTER SET LATIN NOT CASESPECIFIC NOT NULL
      ,Name      VARCHAR(18) CHARACTER SET LATIN NOT CASESPECIFIC NOT NULL
      ,Address   VARCHAR(40) CHARACTER SET LATIN NOT CASESPECIFIC NOT NULL
      ,Phone     CHAR(15)      CHARACTER SET LATIN NOT CASESPECIFIC NOT NULL
      ,DeptNo    SMALLINT
      ,Salary    DECIMAL(16,2)
      ,YrsExp    BYTEINT
      ,DOB       DATE FORMAT 'YYYY-MM-DD'
      ,MedStat   CHAR(1) CHARACTER SET LATIN NOT CASESPECIFIC
      ,EdLev     BYTEINT
      ,Note      VARCHAR(79) CHARACTER SET LATIN NOT CASESPECIFIC NOT NULL
      )
UNIQUE PRIMARY INDEX ( EmpNo );
insert into employee
sel * from employee_stage ;
collect statistics on EMPLOYEE index ( EmpNo) ;
.quit
```

Thank you !

Teradata

a division of  NCR

You've never seen your business like this before.