History and Comparison of Relational Database Management Systems

TechnoCircle HVB Information Services

Contents

1. History of Relational Database Management Systems

2. Architecture Overview of important RDBMS

2.1 A Compilation of Generic Terms

2.2 Timeline of a Transaction

2.3 RDBMS Architecture Comparison

- Oracle
- DB/2
- Sybase
- SQL Server
- MySQL

3. Technology Comparison: Oracle, DB/2, Sybase,SQLServer

3.1 Server Security Model

3.2 Backup/Recovery and Logging

3.3 Data Consistency vs. Concurrency

3.4 High Availability Concepts

4. Vendor Comparison: Beyond Technology

4.1 Vendor Product Portfolios

- Oracle and Sybase

4.2 Platform Support / Market Shares

4.3 Cost Comparison

- SQL Server 2005
- Oracle 10g
- DB/2 8.2

4.4 Total Cost of Administration

- SQL Server vs. Oracle

1.1 Origin and Foundation

Q.: Why is the database industry dominated by US companies like

- IBM
- Oracle Corp
- Sybase, Inc.
- Teradata Corp. (NCR Corp.)
- Microsoft?

A.: There was in the beginning a strong competition between industry and academic research.

1.1 Origin and Foundation

- **19th century**: Us Gov. needed reports from large datasets
- **1890**: Herman *Hollerith*: first automatic information processing equipment => punchcard machine => census 1890 / 1900 were processed
- **1911**: IBM founded (*Hollerith* et.al.)
- control of industry production, tax calculation not possible without automatic data processing
- **1935**: Social Security Act: records of 26 mio individuals, special equipment by IBM
- Bureau of Census: bought UNIVAC I
- **1959**: 200 computers in the Pentagon
- **1960s**: remove hardware constraints from programmers, term: "database"

- Standardization efforts
  - computing => commercial market
  - techniques for: data access, data quality, security, control
  - **1960**: US DoD: Conference on Data system Languages (Codasyl)

    => Cobol (Common Business Oriented Language)
  - tapes => hard disks (serial => random)
  - **1961**: Charles *Bachmann* (GE): IDS (Integrated Data Store System)
  - **1960**s: Database Task Force Group (Codasyl): *Navigation* with computer languages in databases

1.1 Origin and Foundation

- **1968**: IBM: IMS (partly developed during the Apollo project)
- **1971**: formal standard: Codasyl approach to database management
  - hierarchical database model
  - network database model
  - => navigational databases (*Bachmann* 1973: The programmer as navigator)
- adoption in mainframe market outside IBM (Eckert-Mauchley Comp. Corp., Honeywell, Inc., Siemens AG, DEC, Prime Comp. Corp.)

- *Edgar F. Codd* (IBM) was dissatisfied ("taking the old line view, that the burden of finding information should be placed on users …")
- landmark paper: *"A relational Model of Data for large shared Data Banks" (see note)*
  - Indepence of Data from the Hardware- and Storage Implementation
  - automatic navigation to the data set (high level nonprocedural language for data access (Record => Set)
  - pointers => keys (primary, secondary)
  - theoretical proposal, no practical design or implementation

1.1 Origin and Foundation 6/11

- impact was not fully recognized (even by *Codd*)
- IBM: IMS was preferred (sole strategic) product (invest, infrastructure, success, profit)
- relational model had to survive in and outside IBM
- *Codd* initiated public debate with *Bachmann* (Codasyl)
- 1970s: two projects for development of relational products:
  - Army, Air Force, Navy, NSF, UC Berkeley funded: *Ingres*
  - IBM funded: *System R* (1973)
    - begin of competition!

1.1 Origin and Foundation

System R: (see notes)

- high level, non navigational data independent user interface
- 1974/75 prototype
- 1978/79 fully functional, multi user version => SQL

Ingres (Interactive Graphics and Retrieval System)

- handling of geographic data
- 1974: Prototype, own query language: QUEL
- several revisions, testing through user community => feedback
- rapid spread because of DEC hardware at universities

1.1 Origin and Foundation

Diffusion and Commercialization (Ingres)

- Source free available (1000 copies in circulation)
- *Stonebraker* founded Ingres Corp. (1994 bought by CA)
- *Robert Epstein* (Chief Programmer) => Britton-Lee, Inc. => Sybase, Inc. (ideas from Ingres, Public Agencies as first customers)
- *Held/Youseffi* => Tandem Comp., Inc. (predecessor of Nonstop SQL)
- *Held*: up to 1998 Sen. Vice Pres. Engineering, Oracle
- *Hawthorn/Ubell* => Britton-Lee => cofounder of Illustra Inform. Technol., => *Informix*
- …

1.1 Origin and Foundation

Epstein: "What came from Ingres was the experience of having built a prototype … to say what parts need to be done differently."

- hard competition between Ingres and System R
- System R: 15 people
- Ingres 1973-1979: 30 people (never more than six)

- SQL persisted (would have been possible years earlier)

- *Larry Ellison* (IBM): read publications of the System R group => Oracle
  - sold SQL compatible product <u>before</u> IBM
- 1980:
  - IBM developed SQL/DS => mainframe market
  - second generation of relational systems
  - *Codd*: ACM Turing Award
- 1982: End of Ingres Project
- 1985: continues in Postgres (UC Berkeley), object relational

Lessons learned

- strong governmental influence at dev. of rel. systems (tasks like census)
- funding speeds up commercialization (competition System R <=> Ingres)
- free publication of research results => technology transfer in community
- scientists from industry laboratories founded their own enterprises

1.2 History of Oracle (*)

1975-1979 1st Rel.: Oracle 2	- Market Demand: RDBMS - Larry Ellison (Marketing) , Bob Miner, Ed Oates (Code): Software Development Laboratories (SDL) => Implementation of a RDBMS => Oracle 2 (see notes) - less than one year development time
1979 1981 VAX Factor 1983	SDL => RSI (Relational Software, Inc.) market pressure: fast avail., cheap product, - PDP-11 (widely used VAX platform), subqueries, joins, no transact. - Oracle 2.3 : complete RDBMS -RSI => Oracle Corp.

1.2 History of Oracle

1983 Oracle 3 Portability	- Rewrite in C (White Smith C) - many basic design decisions (influence up to date) (see excurs 1) - not very stable !
1985	- hard competition in database industry, begin of tool development (see excursus 2) - Oracle Marketing (see excursus 3) - many systems only remakes of old hierarchical systems - Codd: 12 Rules for RDBMS (Appendix A)

1.2 History of Oracle 3/16

1984 Version 4	- rewrite in K&R C - Portability on almost any platform (flex. FS and IPC facility) (Note) (see Excursus 4) - Views, Named Accounts - Read Consistency
1985/86 Vers. 5.1	- encrypted passwords - subqueries - precompiler (Cobol, Fortran IV, PL/1) ■ SQL embedding in 3GL programs

- bad quality until Version 8
- OCI was the choice (Library Calls for exec SQL in 3GL prg.)

- SQL*Net:

  - C/S Arch.
    - system, platform independence
    - DB Server also Client (Database Link)
    - Query Only Distrib. Database Support (see note)
    - Gateways for other DB engines

1.2 History of Oracle 5/16

	■ growing spread of PCs => mismatch ■Server: DB + Appl., PC: only Terminal + Interface ■Two Task Arch. IPC => TCP/IP was possible (see note) ■ better networks => more effic. implementation of SQL*Net ■n-Tier architectures were possible - first tries with clusters
1987	- standardized applications were needed - Founding of Applications Division at Oracle (Business Mgmt. Software)

1.2 History of Oracle

late 1980s Problems: Perf./ Scaling	- Oracle: good for rapid dev . of OLTP Applications - no scaling, few sessions - no precompiled queries, Parsing Overhead - RBO: good for OLTP on well normalized schemas bad for DSS - growing number of users aggravated performance problems
1988 - 1991 Version 6	- true row level locking (additional charge) (see notes) - Online Backups - Shared Pool (minimization of parsing overhead)

1.2 History of Oracle 7/16

	- change from block level logging to field level logging - keep rollback segments in the database ■ recovery of undo structures via redologs (see notes) ■ Read Consistency - huge memory consumption when many clients => MTS (see notes) - no Upgrade Path from Oracle 5 to Oracle 6

1.2 History of Oracle 8/16

1992 Version 7	- Oracle had to rescue non-tech educated users from their errors - data integrity: declarative constraints (up to now only unique keys and NULL-value checking) - PL/SQL storage in database (procedural language + SQL) - triggers: implement data rules (insert into a => update b) - stored procedures: implement process rules - =>radical changes in application design - emergence of standard software (database independent design) - role concept

1.2 History of Oracle

	- Cost Based Optimizer: not perfect - Rule Based Optimizer since Version 6 no longer improved
1993 1994 1995 1996	- Business Applications in Client/Server Mode - Port to Windows - Distributed Transactions (2PC) - Parallel Query => Data WareHouse - porting of all Tools to Windows - graphical EM: not usable (see note)

1.2 History of Oracle

1997 Version 8	- LOBs (storage of unstructured data) - Partitions (ease of administration) - Materialized Views (Benchmark Opportunism) (see note) - Function Based Indexes (error tolerance) - Bitmapped Indexes: Performance Boost for DWH - Oracle Discoverer: Ad Hoc End User Query Tool for BI - Support for object oriented Developing, new datatypes - completely revised development suite (native Java): ■ Jdeveloper - Oracle Application Server 4.0 (middle tier)

1.2 History of Oracle

1998 Version 8i	- native Java Runtime Environment - SQLJ (SQL in Java Code) - interMedia: managing of Multimedia Content - Repository for Designer (Modeling and Application Generation Tool) - move away from C/S Paradigma: Oracle Applications run in browser - Port of RDBMS and Application Server to Linux - RAW Iron Initiative: Flop (Pack. of Server+cust. OS+Oracle Env.) - XML Support

1.2 History of Oracle

1999	- Oracle WebDB (manage DB driven Websites) => Oracle Portal
	- Applications 11 => Applications 11i
2000	- Internet Filesystem (iFS) (see note)
	- Oracle AS9i Portal Technology (Content Mgmt, Portlets, Deployment Options)
	- OAS9i Wireless
	- Webcache Technology (Part of OAS)
	- Preview Cache Fusion (RAC)

2001 Version 9i	- DEC could not enter the mainframe market - research in ■ SMP- architectures (cache coherency) ■ cluster architectures (Shared Disks) also with Oracle Shared Instances (Instance:DB = N:1) (see note) - Real Application Cluster (Oracle Cluster Technology)

1.2 History of Oracle 14/16

2001	- native XML Support - Jdeveloper: J2EE- and XML-Support - Table Compression (Block Level) (Index Compr. in Vers. 5 and 8) - Flashback Query - OEM becomes usable (scalable architecture, repository)
2004 Oracle 10g	-development was always customer driven, but now: -enormous improvements in manageability and self-tuning of the database (see note) - Automatic Workload Repository (AWR) - Automatic Database Diagnostic Monitor (ADDM) - Automatic Space Managament (ASM)

	- Selftuning must be licensed! - Selftuning is valuable for small businesses without dedicated DBAs - Oracle can afford not to satisfy small customers at moderate prices and let big customers pay extra charges - Product Quality: ■ new features often half hearted and error prone (see note) ■ initial release of each version is unusable in the real world ■ this is in 10g different

1.2 History of Oracle - Summary 16/16

1979	Oracle 2	first commercially available RDBMS
1983	Oracle 3	Portability
1984	Oracle 4	Transactional Integrity, Read Consistency
1986	Oracle 5	Client/Server
1988	Oracle 6	Scalability
1992	Oracle 7	Data Integrity
1997	Oracle 8	Partitioning
2001	Oracle 9	Cluster technology comes of age
2004	Oracle 10	Manageability, Selftuning

1.3 History of Sybase

1984 1986	- Robert Epstein (UCB), Mark Hoffmann (Univ. of Arizona) => Robert Epstein Consulting => Sybase (System+Database) - Sun: major Partner - SQL Server V1.0 (Beta) and DB Library on SunOS: ■ Client/Server Database for OLTP (Appl. Logic in Server) ■ pioneer with Stored Procedures, Triggers
1987	- SQL Server (Beta) on VMS - 100 employees, 150 customers

1.3 History of Sybase 2/13

1988	- second major Partner: Microsoft - SQL Server for OS/2 (joint dev. of Sybase, Microsoft, Aston-Tate) - Version 3.0: Stored Procedures, Triggers, CBO ▪ first market Version - Sybase: Unix/VMS, Microsoft: Windows, OS/2
1989	- Open Client / Open Server: API for Comm. between Clients and different Data Sources
1990	- Sybase SQL Server 4.0: text, image datatypes - Vision of "Total Solution C/S Company" : aquis. of SQL Solutions - Integration of IBM MVS mainframes in "C/S World"

1.3 History of Sybase 3/13

1991	- Version 4.2: international language support - enters Chinese Market - Powersoft introduces Powerbuilder 1.0
1992	- Version 4.8: SMP support (engines), union, new datatypes - Version 4.9: multibyte Character Sets - Version 4.9.1: extraordinary stable - Version 4.9.2: log-based Replication

1993 System X	-Backup Server, Monitor Server, - RI Constraints, Cursors, Identity Columns, system roles - thresholds, Auditing, encrypt. Password Storing, sybsystemprocs DB - CT Library replaces DB Library (C/S API) ■ but: severe Installation- and Stabilityproblems - Replication Server (Version 10): Open Replication Technology allows replication regardless of Database - Vertical Industry Solutions Group ■ Healthcare, Financial Services, Oil & Gas, Telecom

1.3 History of Sybase

1994	- Microsoft quits Partnership (reseller for OS/2 and Win NT) ■ gets copy of Sourcecode Version 4.2 Sybase SQL Server ■ Microsoft SQL Server now an independent Product - both comp.: SQL Server and Transact SQL (T-SQL) (Trademark still today owned by Sybase)
1995	- Merger with Powersoft (leading supplier of developm. tools for 4GL Applications) - with Powersoft came Watson: small RDBMS and Enterprise Synchronization Solution on sev. OS for mobile and embedded products ■ SQL Anywhere Studio => Adaptive Server Anwhere ■ undisputed mobile DB market leader

	- Version 11.0: ■ Improvement in Quality ■ Table Partitioning ■ Memory Partitioning (named Caches) ■ ANSI Transaction Isol. Level 0 (Dirty Reads) ■ sp_sysmon, online database, Server config File
1996	- dominant provider in mobile DBMS Market ■ ASA: market share 22% (2000: 68%) - Strategy: Focus on OLTP, DataWarehouses, Internet Applications

1.3 History of Sybase

1997	- Version 11.5: ■ Sybase SQL Server renamed to Adaptive Server Enterprise ■ case expression ■ CIS (Proxy tables) ■ XP Server, Historical Server ■ dbcc checkstorage, user defined Roles ■ Query Parallelism, Resource Limits ■ log. Process Manager - Jaguar: 1999 EA-Server (Web Appl. Server for >30 DB)

1998	- Version 11.9 "EARL" (Eagerly Awaited Row Level Locking) ■ DOL on row level / page level (see note) ■ lock timeouts , lock table statement ■ standby_access database, reorg command ■ ANSI Isol. Level 2 ■ license monitor ■ Optimizer Statistics Redesign - 11.0.3.3 : Port to Linux

1.3 History of Sybase 9/13

1999	- ASA for Handhelds - Aqu. Home Financial networks, Comb. with Financ. Server Unit ■ => own subsidiary: Finacial Fusion (leader in fin. Services) ■ 2000: 68% Wall Street comp., 95% Fortune 100 with Sybase ■ market share in Financial Sector: 58% - Aqu. of Data Warehose Network => Sybase Industry Warehose Studio - ASE Version 12.0: ■ exec immediate, identity_gap, license manager ■ abstract query plans, alter table, quiesce database ■ optional licensable features: Java in ASE, Adv. Sec., DTM, HA

1.3 History of Sybase

2001	- ASE 12.5
	■ union views, login trigger, dynamic server reconfiguration
	■ larger page sizes, XQL-Parser
	■ ddlgen
	■ new licensable Options:
	    ■ LDAP, Filesystem access with Proxy Tables
	    ■ builtin Jave Enterprise Beans Server
	- Customer Base: 90% of Investment Banks, 60% dep. banks, NYSE

2002	- Version 12.5.0.1: ■ ASE Replicator (lightweight Data Replication) ■ sybmigrate ■ free Developer Edition - Version 12.5.0.2: Port on MacOS
2003	- Version 12.5.0.3: ■ MDA monitoring Tables ■ multiple Temp DBs ■ statistics sampling - Version 12.5.1: job scheduler, dynamic data caches

	Version 12.5.3: cross platform transfer with dump/load
2005 ASE 15	- revised partitioning support -computed columns, function indexes - Adv. System Metrics, auto update Statistics - native conn. to MQSeries - completely rewritten CBO - Unicode Data Support - native Encryption in the Database - selective load

1.3 History of Sybase – Summary

1984	Foundation
1998	SQL Server: 1sc C/S Database for distr. Appl. and mobile Dev. Mgmt.
1990	Open Replication Technology – Database independent
1991	early activities in China: today market leader
1995	IT Infrastructure Standard in Wall Street => Dominance Fin.Service Market
2003	Sybase IQ: worldwide largest DWH at Nielson Media
2004	Sybase iAnywhere: Marketleader in Mobile Devices

1.4 History of SQL Server – Milestones 1/2

1987	Sybase launches SQL Server for Unix
1988	Microsoft, Sybase and Aston-Tate port SQL Server to OS/2
1989	SQL Server 1.0 for OS/2
1990	SQL Server 1.1 supports Win 3.0 Clients Aston – Tate quits Development of SQWL Server
1991	Microsoft and IBM quit joint Deveopment of OS/2
1992	Microsoft SQL Server 4.2 for 16-Bit OS/2 1.3 is released
1992	Port of SQL Server to Win NT (MS and Sybase)

1.4 History of SQL Server – Milestones 2/2

1993	Win NT 3.1 is released
1993	Microsoft and Sybase release SQL Server 4.2 for Windows NT
1994	Quit of joint Development of SQL Server by sybase and Microsoft. Sybase develops Unix variant, Microsoft Windows variant separately.
1998	SQL Server 7.0
2000	SQL Server 2000
2005	SQL Server 2005

1.5 Overview DB/2 History: Foundations 1/2

System/R Project	- IBM Implementation of the relational Model
Project Aris	- Row Level Locking
Rstar Project	- Rel. Model for Distributed Environments
Starburst Project	- Optimization Strategies in the Relational Model
Garlic Project	- Data Management in diverse Systems

1.5 Overview DB/2 History: Milestones 2/2

1980	Database Integration in System/38 (first Implementation of System / R)
1982	SQL/DS (Mainframe VM, VSE)
1983	DB/2 (Database 2 on VMS)
1987	Database Manager in OS/2 (first on distr. Systems)
1988	SQL/400 for AS/400
1993	DB/2 for AIX
1994	DB/2 for HP-UX, Solaris
1995	DB/2 for Windows
1999	DB/2 for Linux

1.6 Landmarks and Cross Relations 1/3

1970	Codd: concept RDBMS			
1976	System/R: RDBMS prototype			
1977		Larry Ellison: Oracle		
1979		Oracle Corp. 1st comm. Database (SQL)		
1983	DB/2 for MVS			

1.6 Landmarks and Cross Relations 2/3

1988		Oracle 6	Sybase – Microsoft Partnership	
1989				MS SQL 1.0
1993	DB/2 for AIX		End of Partnership	
1995	DB/2 for Windows		ASE 11.0 Merger Sybase Powersoft	MS SQL 6.0
1996			ASE 11.5	MS SQL 6.5
1998		Oracle 8	ASE 11.9.2	MS SQL 7.0

1.6 Landmarks and Cross Relations 3/3

2000	DB/2 V7	Oracle 8i	ASE 12.0	MS SQL 2000
2001		Oracle 9i		
2003			Sybase IQ	
2004		Oracle 10g		
2005			ASE 15	MS SQL 2005
			Replication Server 15	

2. Architecture Overview of important RDBMS

2.1 A Compilation of Generic Terms

2.2 Timeline of a Transaction

2.3 RDBMS Architecture Comparison

- Oracle
- DB/2
- Sybase
- SQL Server
- MySQL

2.1 Generic Terms: Logical Components 1/10

Generic	*DB/2*	*Informix*	*MySQL*	*Oracle*	*Sybase*	*MSSQL*
Background Processes and Memory Structures						
Instance	Instance	Instance	Server	Instance	Server	Server
Persistent Objects on Storage (Figure A)						
Database	Database	Database	Database	Database	Database	Database
Relation (Grouping of related information, Figure B)						
Table	Table	Table	Table	Table	Table	Table
special purpose table for quicker lookups						
Index	Index	Index	Index	Index	Index	Index

Gen. Terms: 2/10

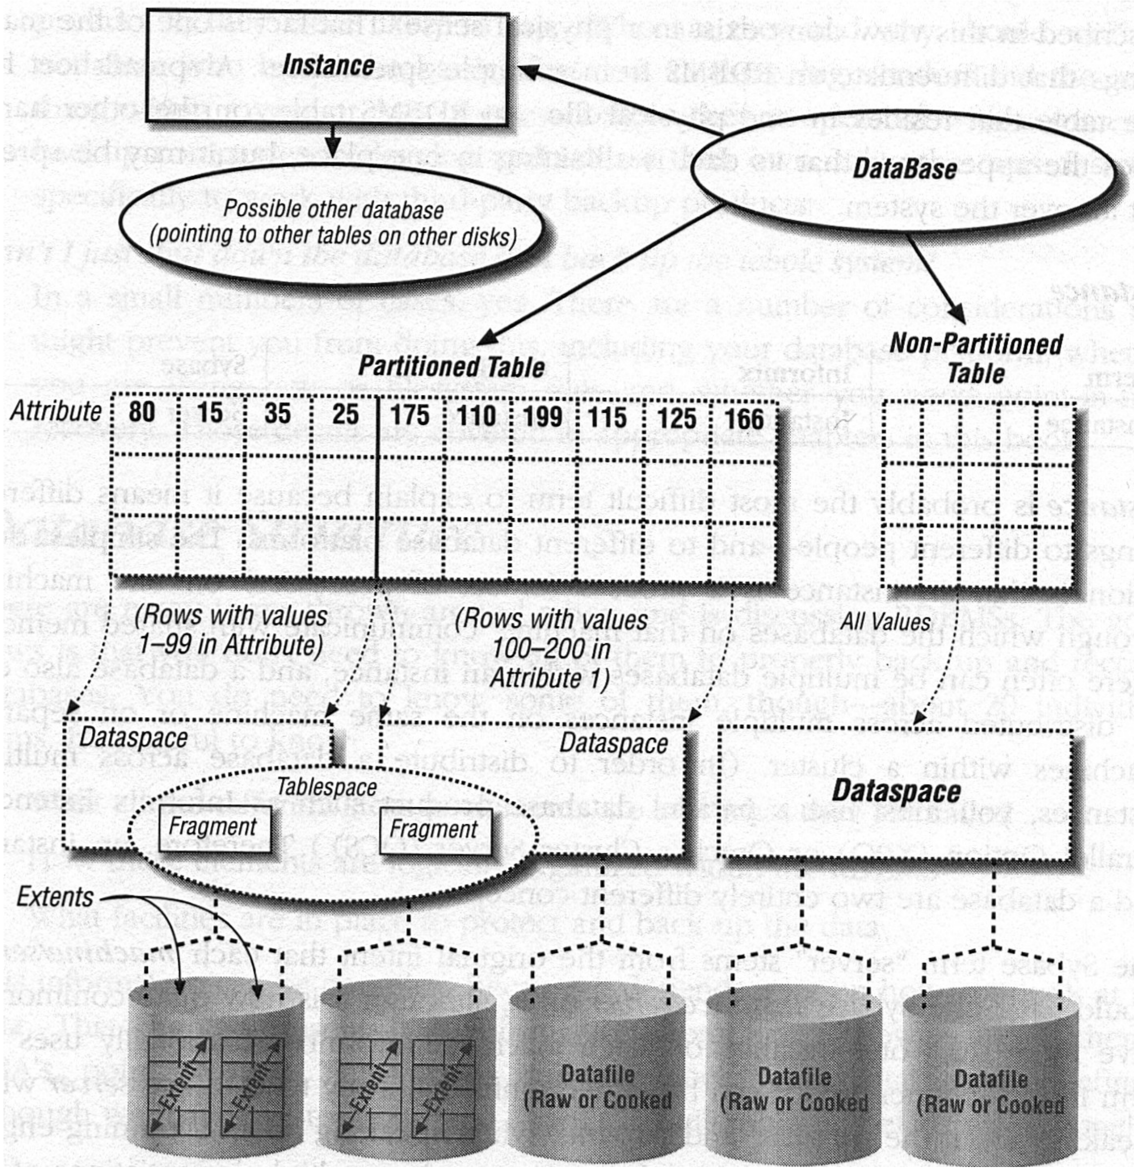

Figure A

Graphical Representation

of a Database

2.1 Generic Terms - Logical Components

Generic	*DB/2*	*Informix*	*MySQL*	*Oracle*	*Sybase*	*MSSQL*
Data which do not fit in a "normal" table						
BLOB space		**(smart) BLOBspace**		BLOB, CLOB	image data type	varchar (max) varbinary (max)
any type of table						
Object	Object	Object	Object	Object	Object	Object
Collection of related Attributes (Figure B)						
Row = Tuple	Row	Row	Row	Row	Row	Row

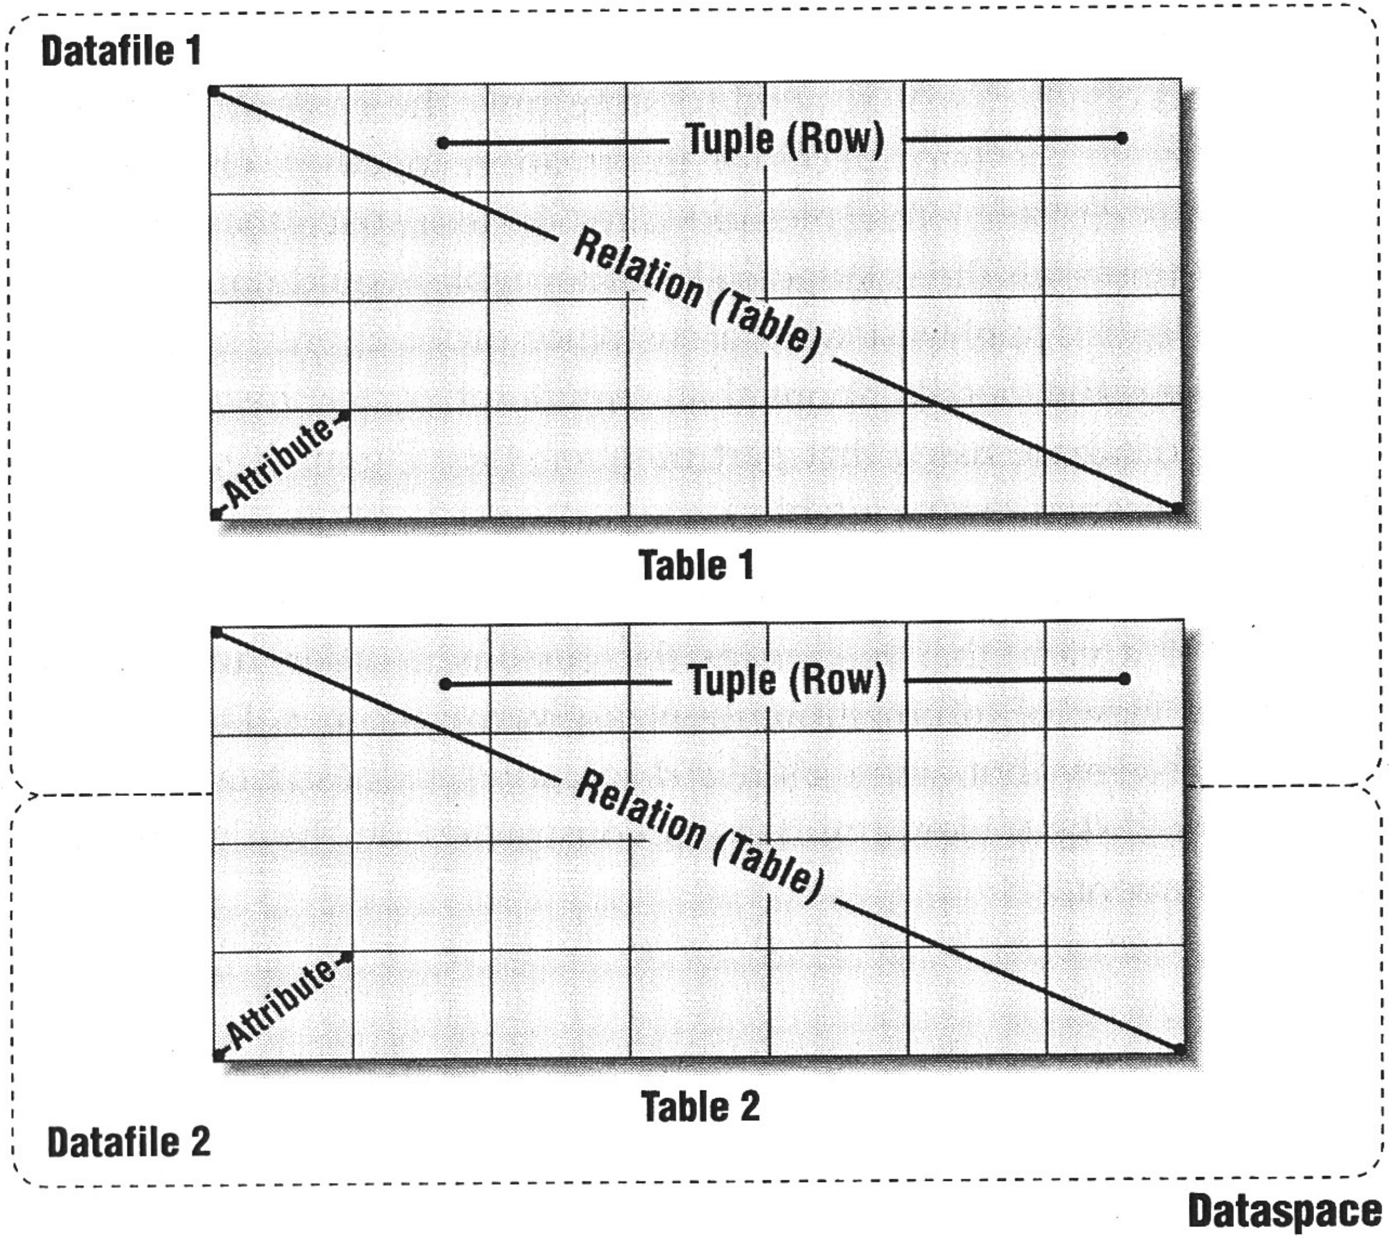

Gen.Terms 4/10

Figure B

Table Layout

2.1 Generic Terms: Logical Components 5/10

Generic	*DB/2*	*Informix*	*MySQL*	*Oracle*	*Sybase*	*MSSQL*
Basic Element of Data within a Table (Figure B)						
Attribute	Attribute	Attribute	Attribute	Attribute	Attribute	Attribute
Activity, which changes one or more Attributes						
Transaction	**Transaction**	**Transaction**	**Transaction**	**Transaction**	**Transaction**	**Transaction**
Point in Time where everything is on Disk						
Checkpoint	**Checkpoint**	**Checkpoint**	**Checkpoint**	**Checkpoint**	**Checkpoint**	**Checkpoint**

2.1 Generic Terms: Physical Components 6/10

Generic	*DB/2*	*Informix*	*MySQL (Inno)*	*Oracle*	*Sybase*	*MSSQL*
Basic Building Block						
Page/Block	Page	Page		Block	Page	Page
Structure for storing data (raw or cooked files)						
Datafile	Container	Chunk	Datafile	Datafile	Device	Datafile/ Filegroup
logically contiguous number of pages (1)						
Extent	Extent	Extent (2)	N/A	Extent	Extent	Extent (3)

2.1 Generic Terms: Physical Components 7/10

Generic	*DB/2*	*Informix*	*MySQL*	*Oracle*	*Sybase*	*MSSQL*
collection of logically contiguous extents (1)						
Fragment	N/A	Tblspace	N/A	N/A	Allocation Unit (2)	Allocation Unit
Space occupied by a single Table (Figure C)						
Tablespace	Object	Tblspace	**Tablespace**	Segment	Disk fragment	N/A
structures where tables go in (4, Figure C)						
Dataspace	**Tablespace (3)**	**Dbspace**		**Tablespace**	**Segment**	**N/A**

Terms: 8/10

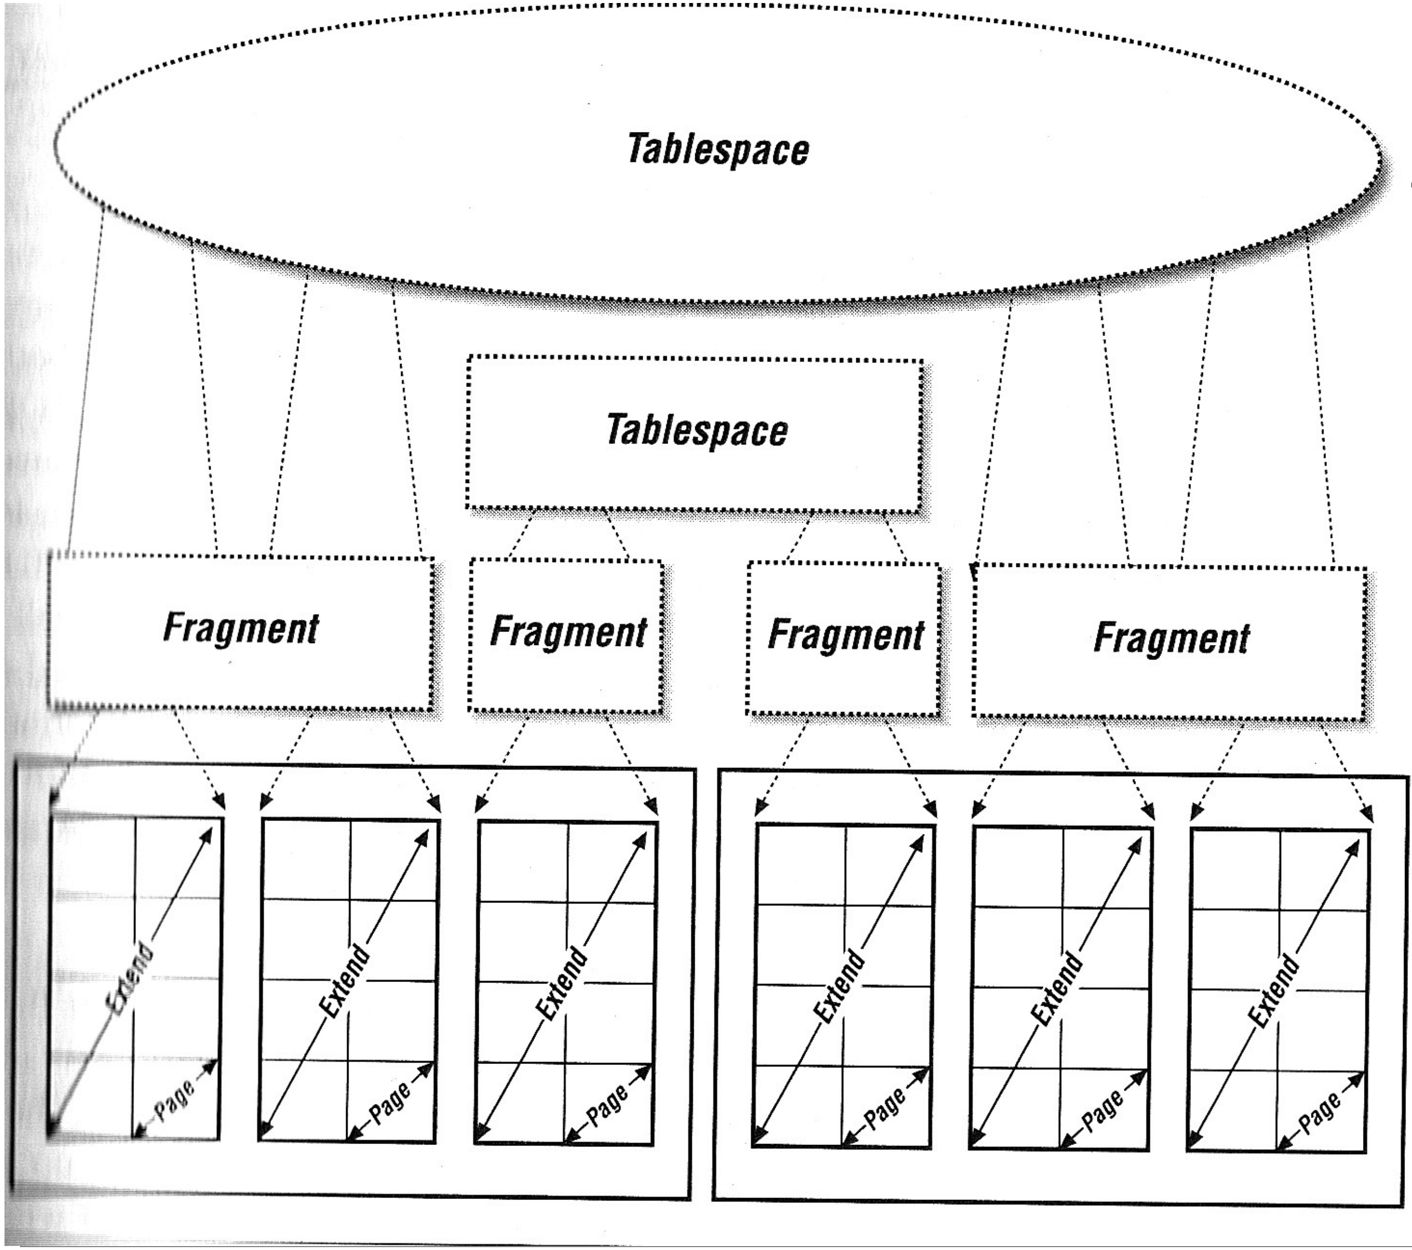

Figure C

Tablespace Layout

2.1 Generic Terms: Physical Components 9/10

Generic	*DB/2*	*Informix*	*MySQL*	*Oracle*	*Sybase*	*MSSQL*
structure for distributing a table across multiple dataspaces						
Partition	Partition	Fragment	N/A	Partition	Partition	Partition
Inventory of all database parts						
Master Database	control files	**sysmaster onconfig file rootdb**	**Information Schema (1)**	control file	master database	master database
storage place of the "before image" of a transaction						
Rollback Log	Logfiles	Physical Log	Undo (2) Segments	Undo Segments	**Transaction Log**	**Transaction Log**

2.1 Generic Terms: Physical Components 10/10

Generic	*DB/2*	*Informix*	*MySQL*	*Oracle*	*Sybase*	*MSSQL*
log place for transaction and page changes						
Transaction Log	Logfiles	Logical Log	Binary Log	Redolog	**Transaction Log**	**Transaction Log**

2.2 General Transaction Properties - ACID -

Atomicity	either all of a transaction happens or nothing
Consistency	Database is brought from one consistent state to another
Isolation	The effects of a transaction may not be visible to other transactions until the transaction is committed.
Durability	After commit the transaction is permanent.

2.2 General Transaction Properties - Timeline -

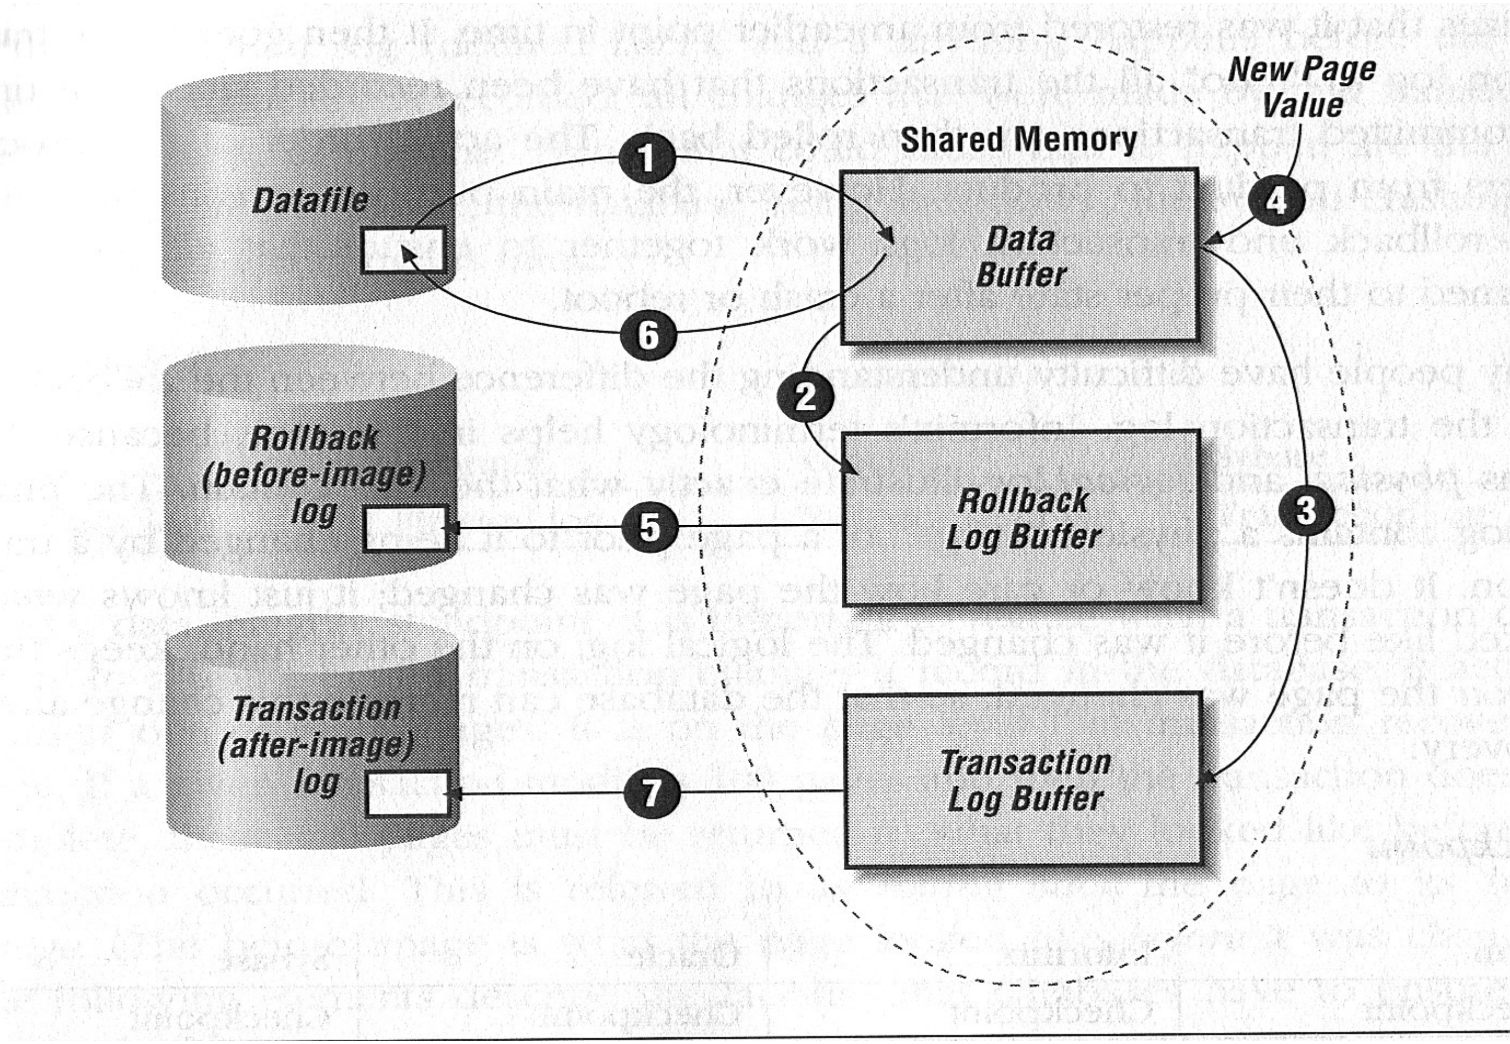

2.3 RDBMS Architecture Comparison - Oracle -

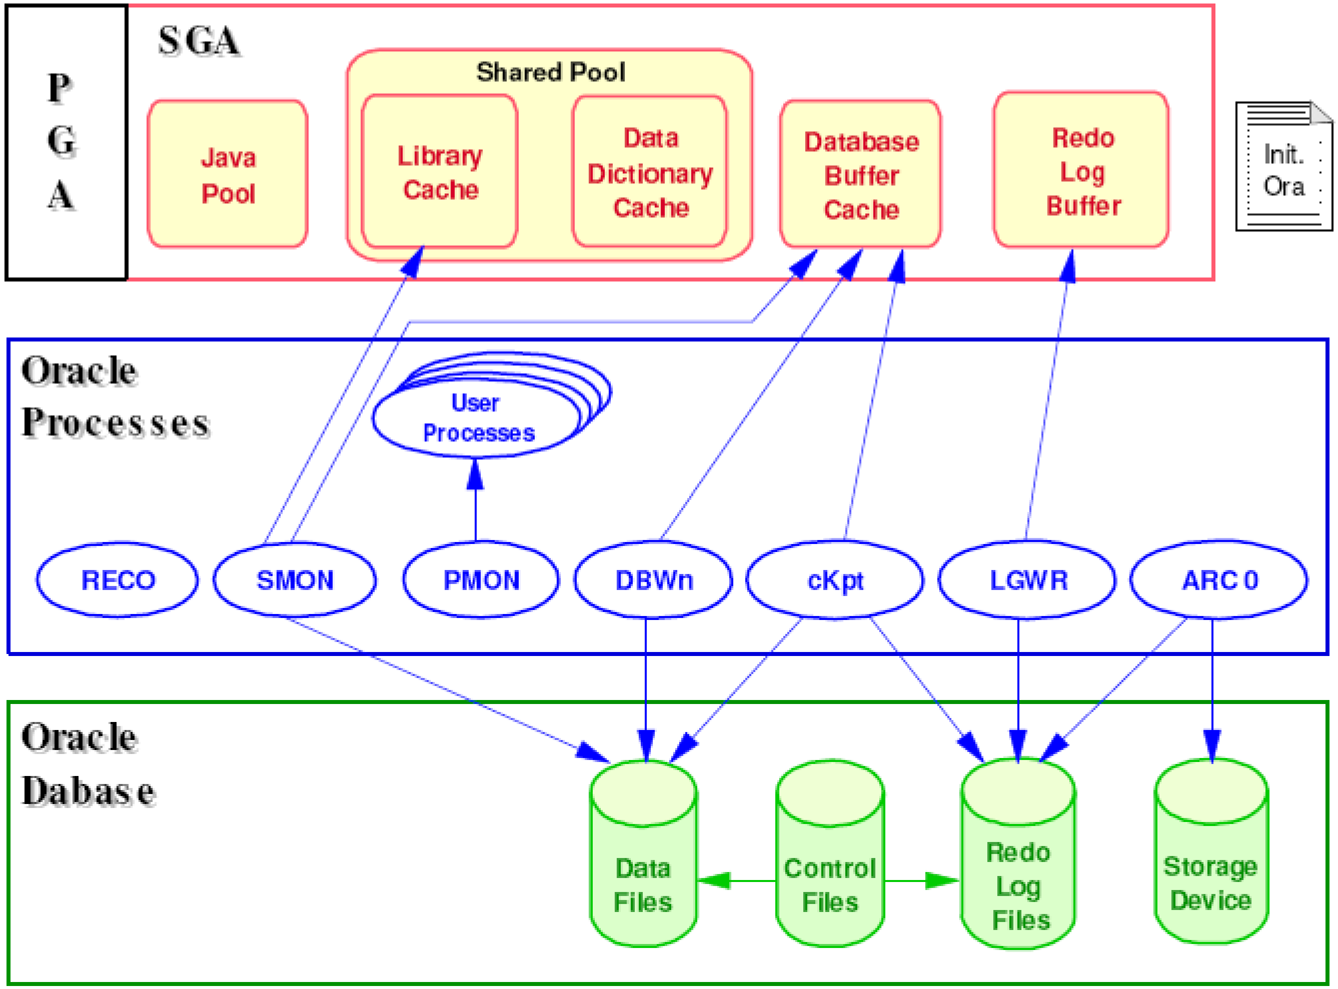

2.3 RDBMS Architecture Comparison - Oracle -

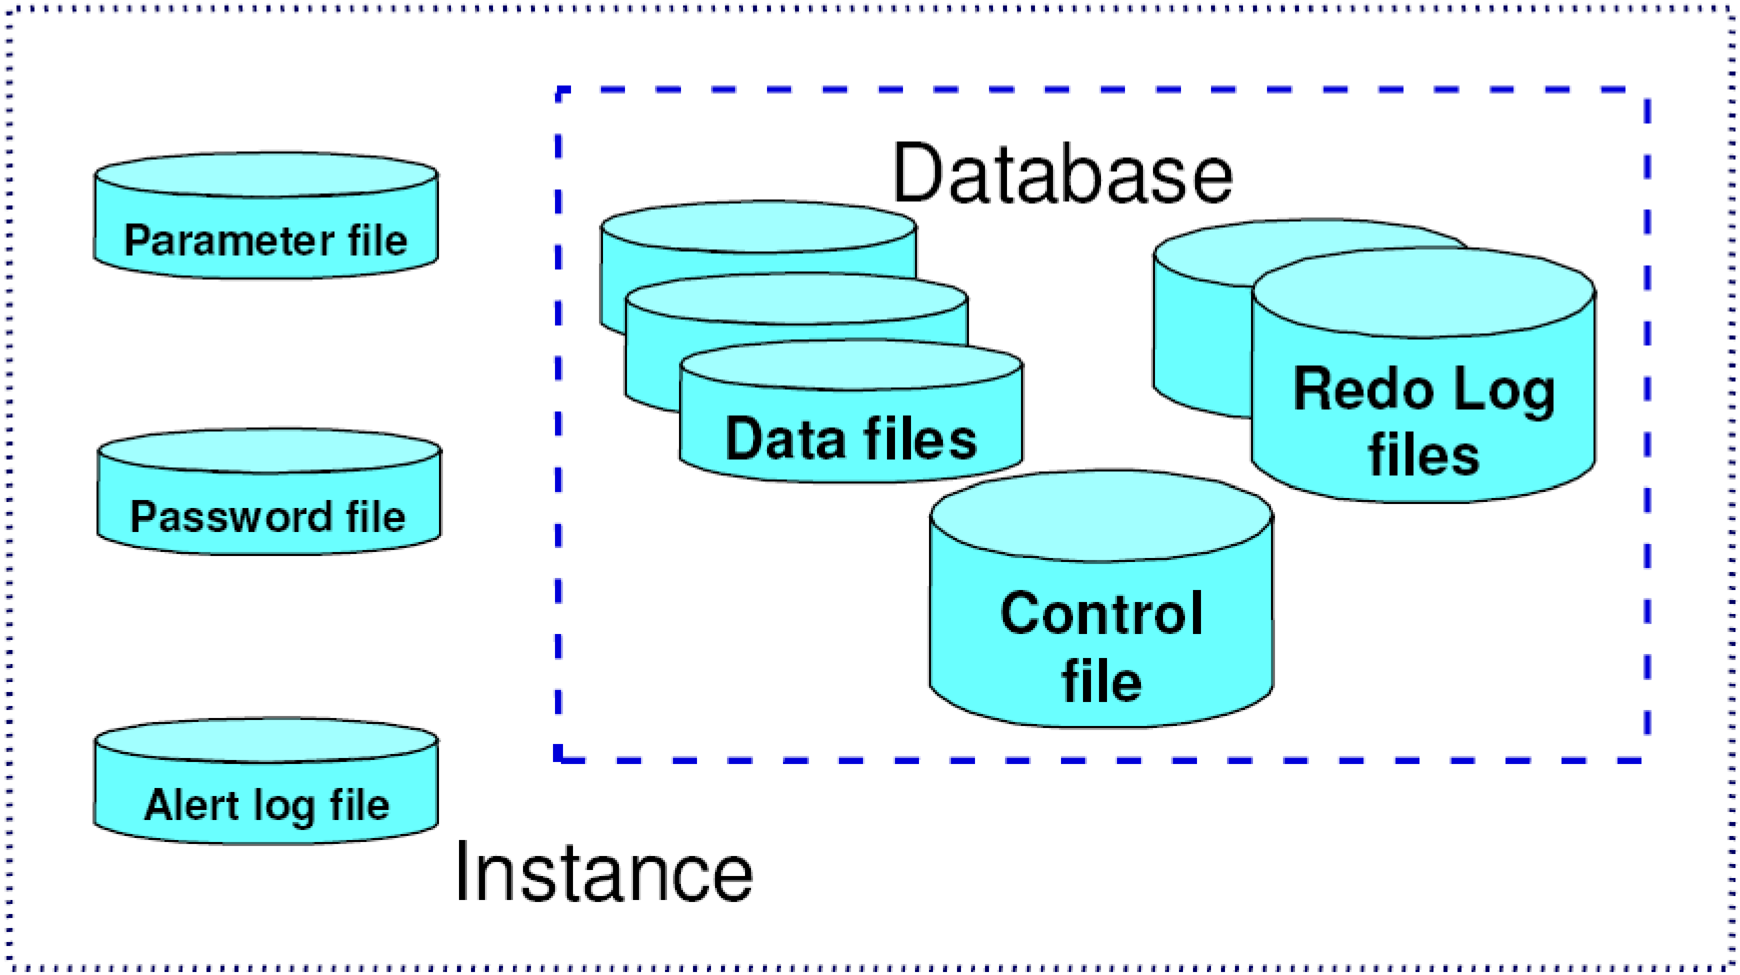

2.3 RDBMS Architecture Comparison - DB2/UDB -

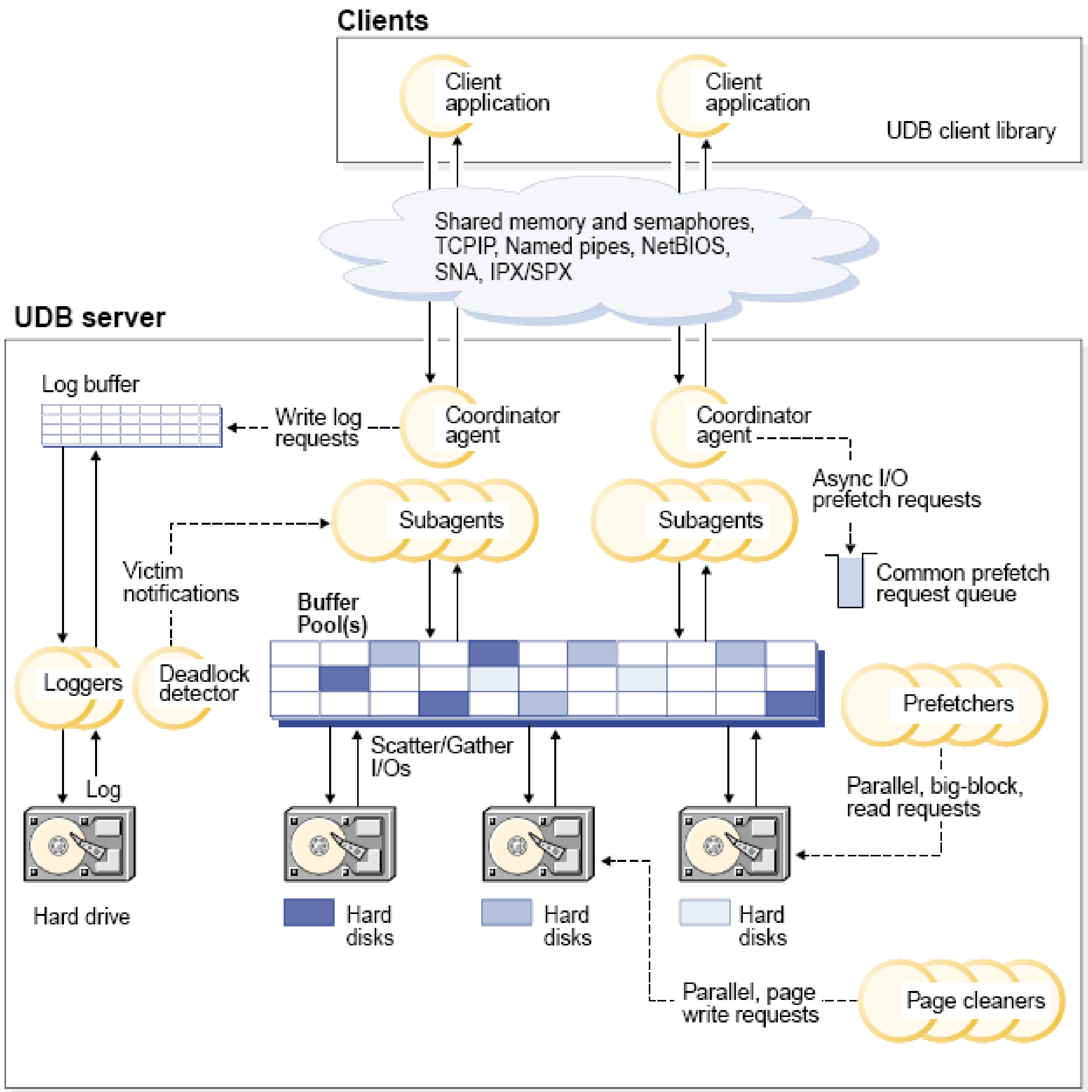

2.3 RDBMS Architecture Comparison - DB2/UDB -

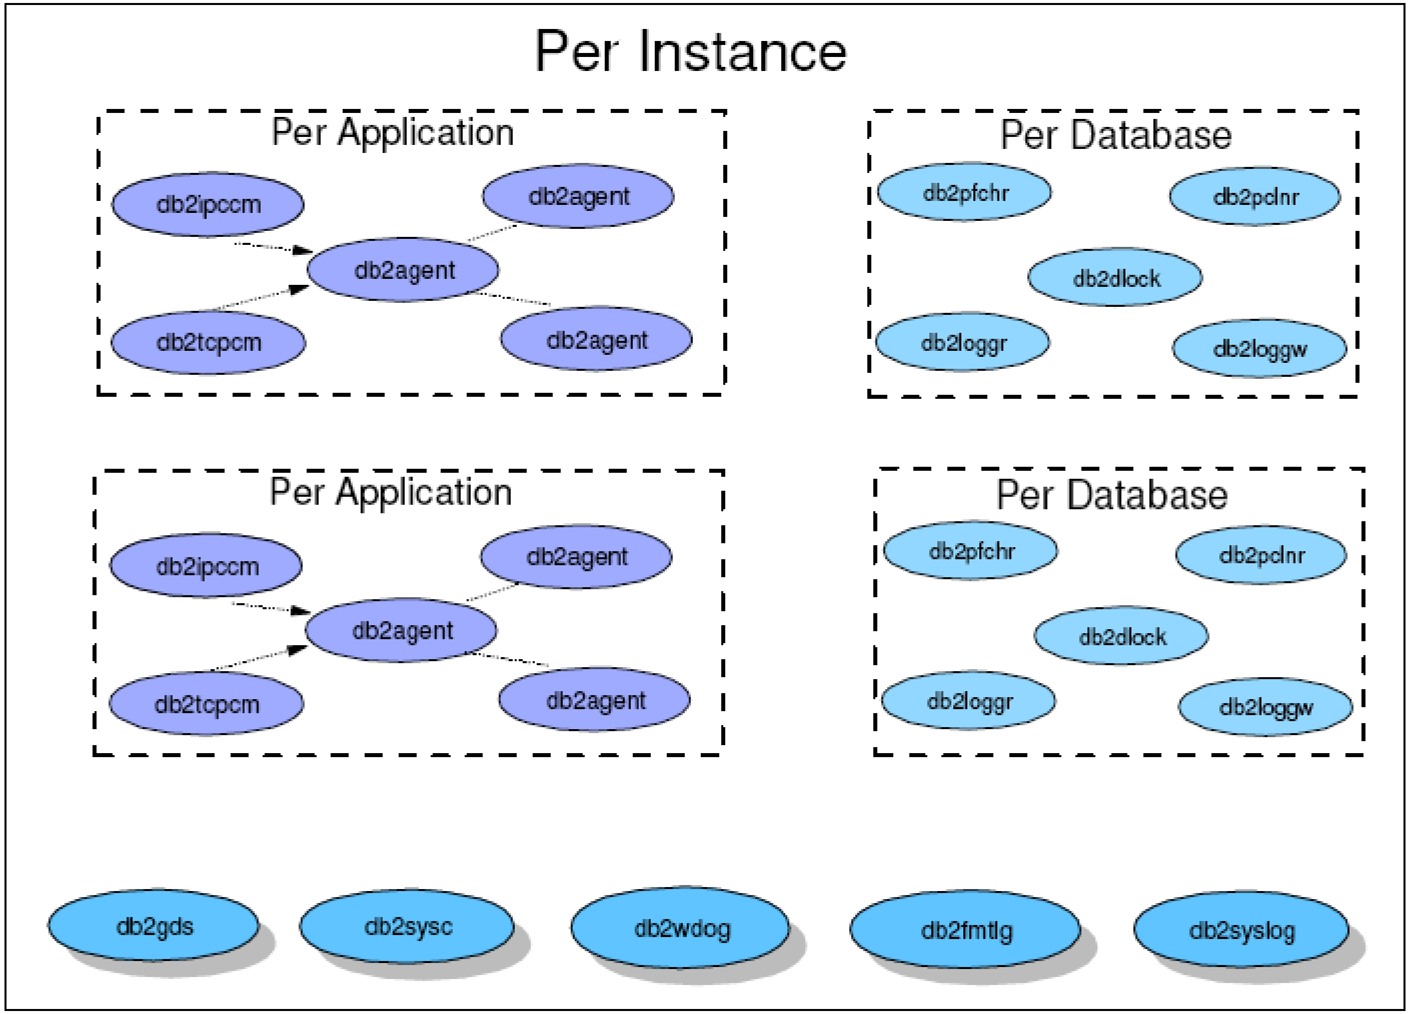

2.3 RDBMS Architecture Comparison - DB2/UDB -

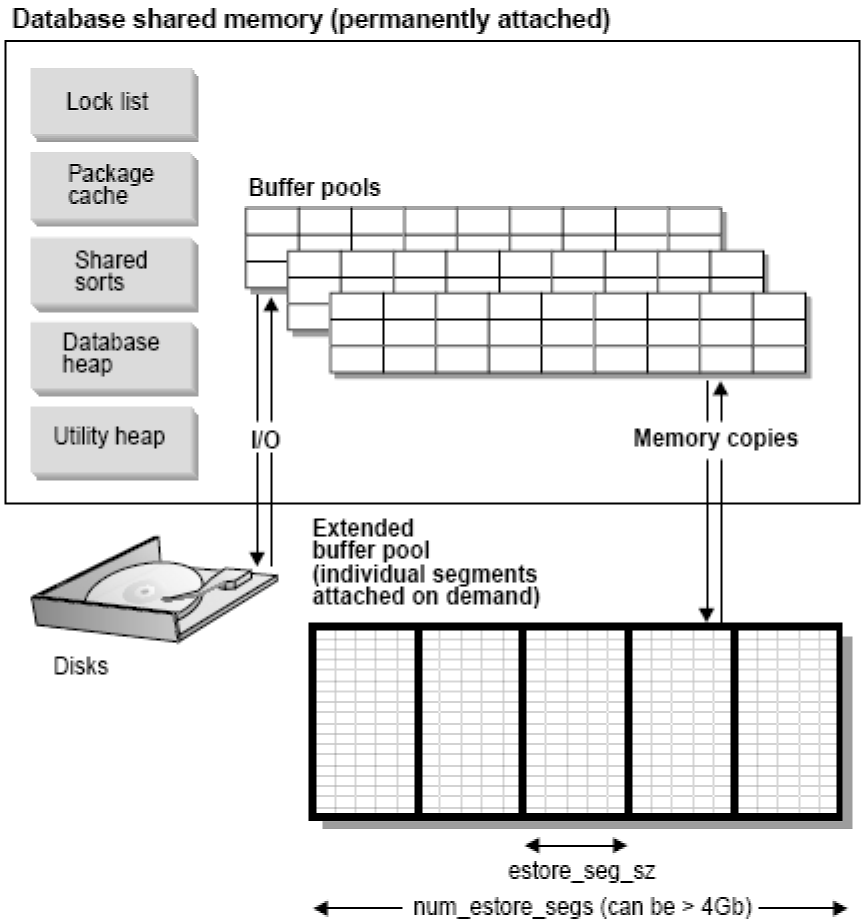

2.3 RDBMS Architecture Comparison - DB2/UDB -

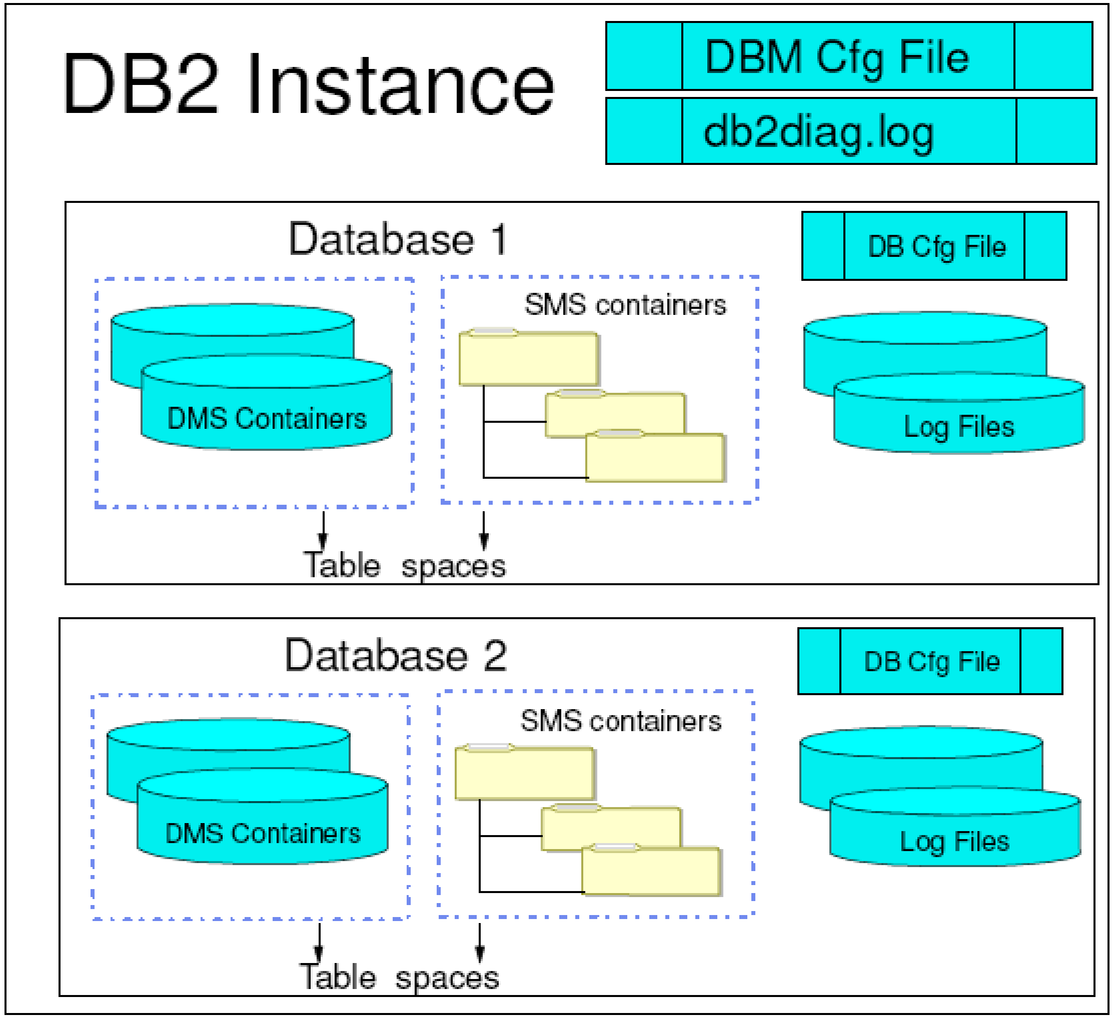

2.3 RDBMS Architecture Comparison - Sybase ASE -

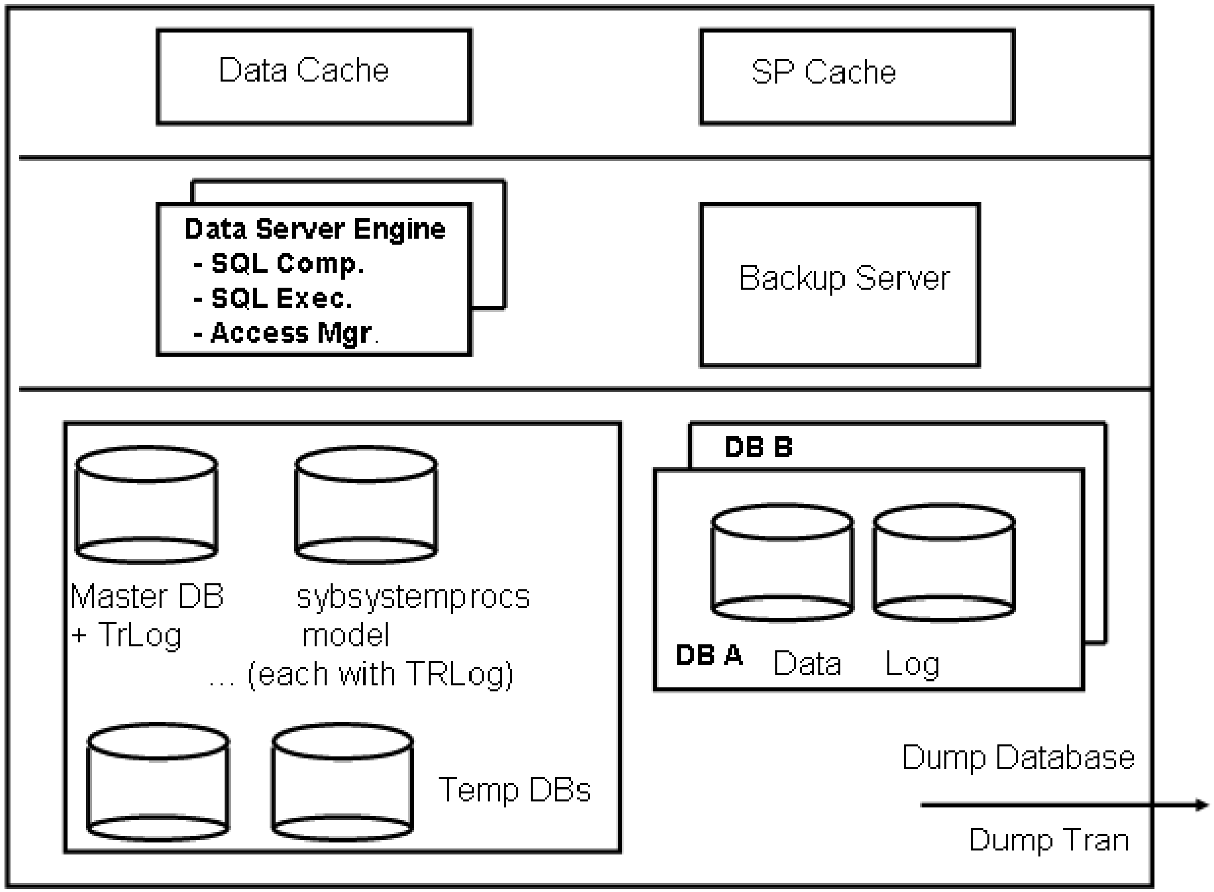

2.3 RDBMS Architecture Comparison - SQL Server -

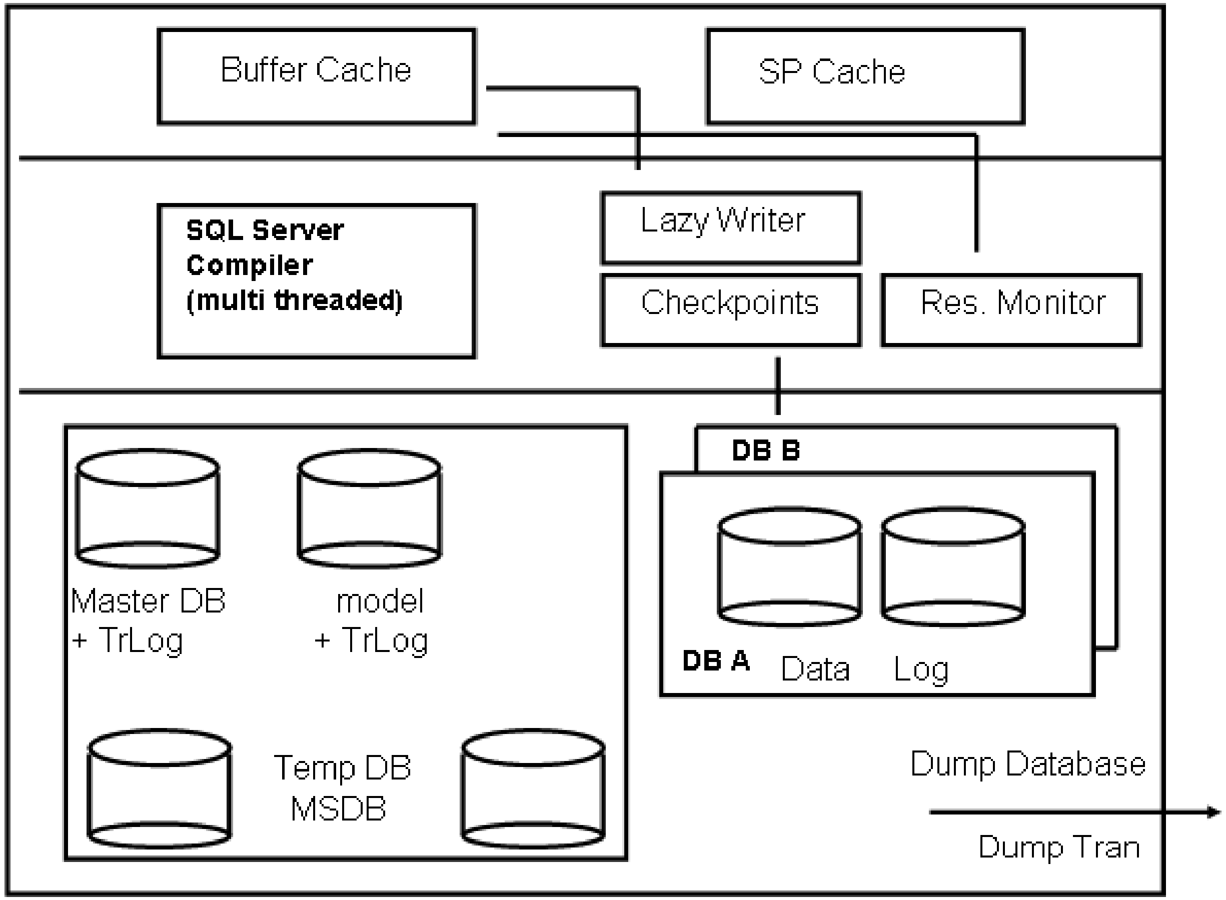

2.3 RDBMS Architecture Comparison - MySQL -

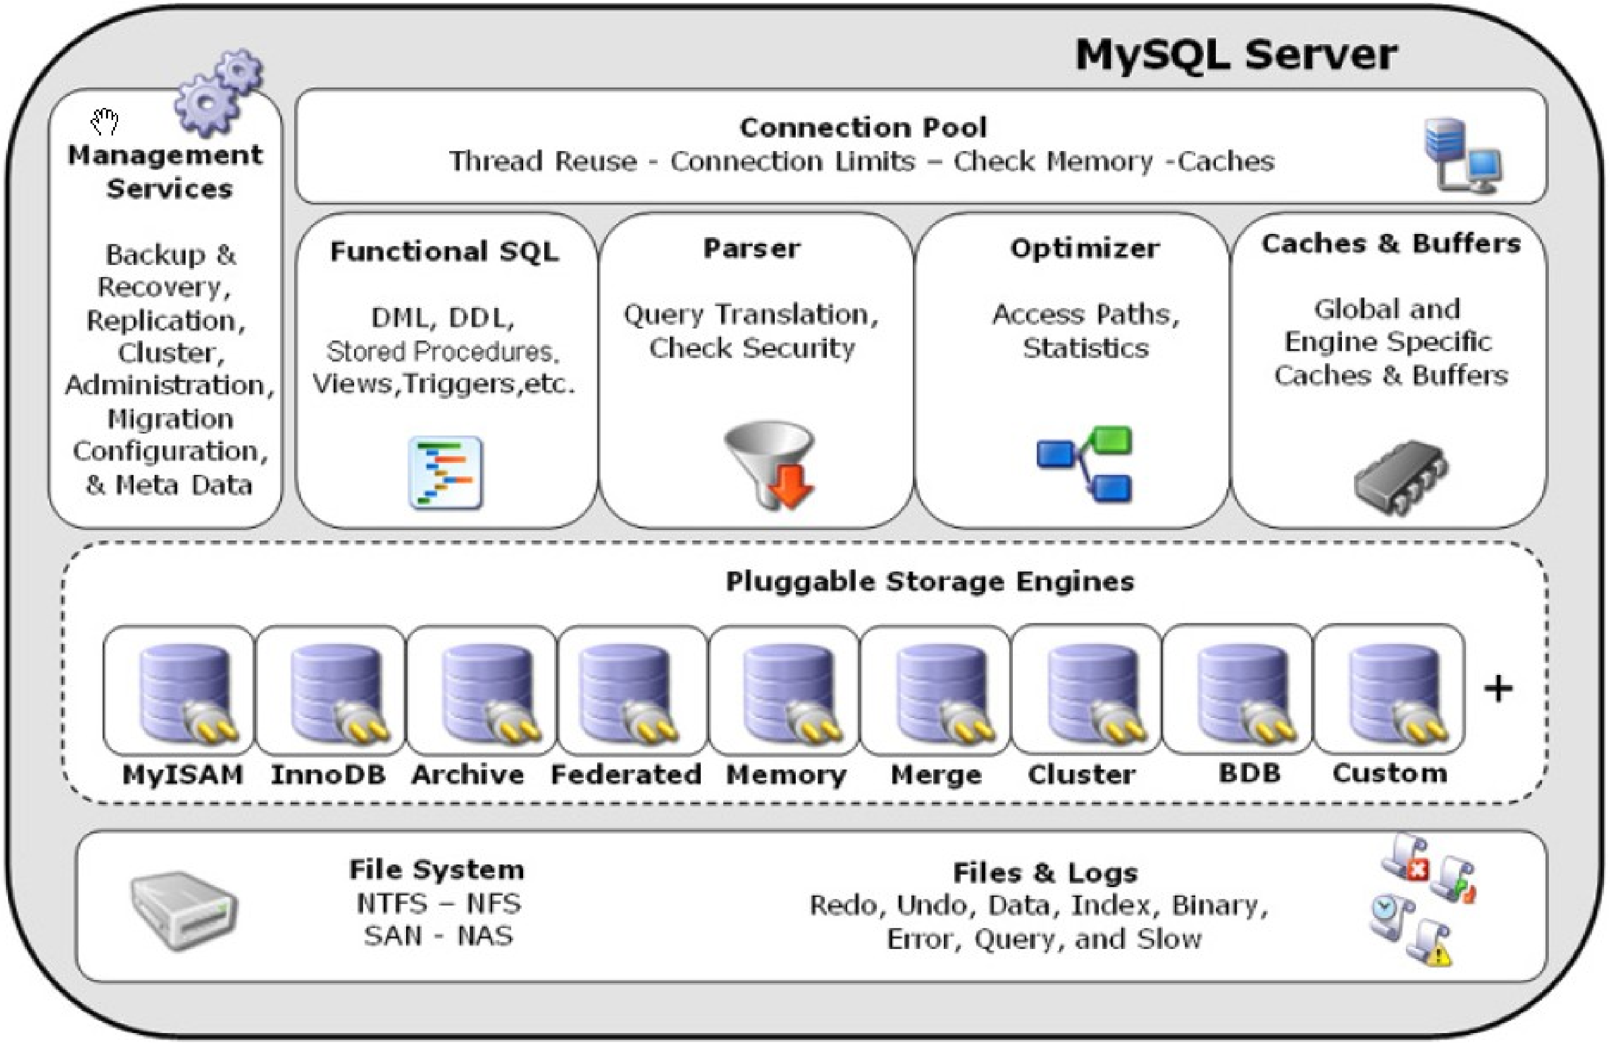

2.3 RDBMS Architecture Comparison - MySQL -

The MySQL Pluggable Storage Engines

- MyISAM

  - Table Locking
  - no Transactions

- InnoDB

  - ACID compatible
  - row level locking
  - Tablespaces

2.3 RDBMS Architecture Comparison - MySQL -

- default Isolation Level: Read Committed
- multi version concurrency control (readers do not block writers and vice versa)

- Archive

- compressed data -> ILM

- Federated

- links to tables on another server

3. Technology Comparison: Oracle, DB/2,Sybase,SQLServer

3.1 Server Security Model

3.2 Backup/Recovery and Logging

3.3 Data Consistency vs. Concurrency

3.4 High Availability Concepts

3.1 Server Security Model

Feature	Definition	Oracle	Sybase	SQL Server	DB/2
Schema	- collection of logical structures of data - classification of objects (namespace)	rudimentary	-	new in 2005	yes
- Login Account - DB User	- right to access the database server - security domain and ■ Quotas ■ Privileges ■ Resource Limits	N / A User	Login DB User	Login DB User	

3.1 Server Security Model 2/5

Feature	Definition	Oracle	Sybase	SQL Server	DB/2
- Groups	- collection of similar users	N/A	up to 11.5	via Windows Groups	
-Privileges (Authorization)	right to run a particular SQL statement or to access a database object	system pr. object pr.	system pr. object pr.	statement lev. pr. object pr.	database pr. object pr.
-Roles	named (protected) group of related privileges	yes	yes	yes	no

3.1 Server Security Model 3/5

Authentication: Identity Verification

Oracle	- Operating System (externally) - Network (globally) ■ 3rd Party (DCE, Kerberos) ■ PKI: SSL, OCI, Wallets, LDAP ■ Remote: Radius - Database: (password) - Multitier: (with OCI)

Authentication: Identity Verification

Sybase	- ASE (password) - external: ■ Kerberos (Network) ■ LDAP ■ PAM (OS)
SQL Server	- Windows Authentication (Integrated Security, NTLM or Kerberos) - SQL Server Authentication

3.1 Server Security Model 5/5

Authentication: Identity Verification

DB/2	- Server: comparison with OS - Server_ Encrypt - Client: comparison to valid credentials on client node - Kerberos - Krb_Server_Encrypt - add. Security Plugins may be installed (LDAP, PKI, …)

3.2 Backup/Recovery and Logging: Oracle 1/6

Logical Backup

- import/export (proprietary format), impdp/expdp (XML)
- SQL Loader for large quantities of data

Physical Backup

- noarchivelog Mode: redologs will be overwritten in a circular fashion
  - no PITR possible
  - protection from instance failure
  - most recent changes available only for crash or instance recovery
- archivelog mode: redologs will be duplicated to an archivelog
  - complete PITR from instance or media failure possible
  - all changes are permanently saved in archivelogs

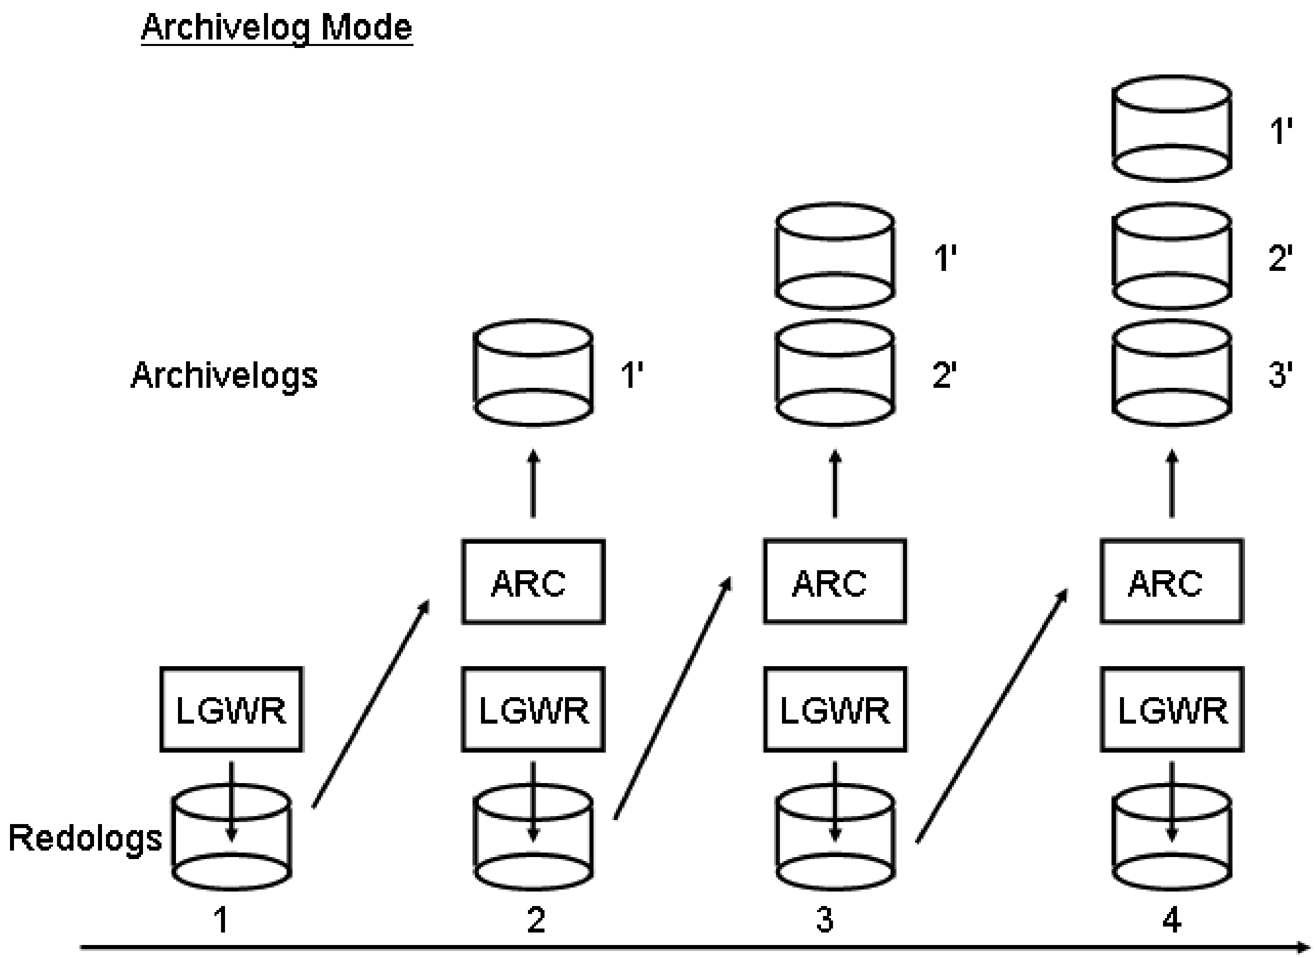

Archivelog Mode
Archivelogs
Redologs
LGWR
ARC
1'
2'
3'
1
2
3
4

3.2 Backup/Recovery and Logging: Oracle 3/6

- Offline Backup: copy of DB files while Database ist down
- Online Backup:
  - User Managed B&R:
    - begin .. end backup + OS commands
    - recover Statement
  - RMAN + (opt. Recovery Catalog) + API for Integration with 3rd – Party Media Manager

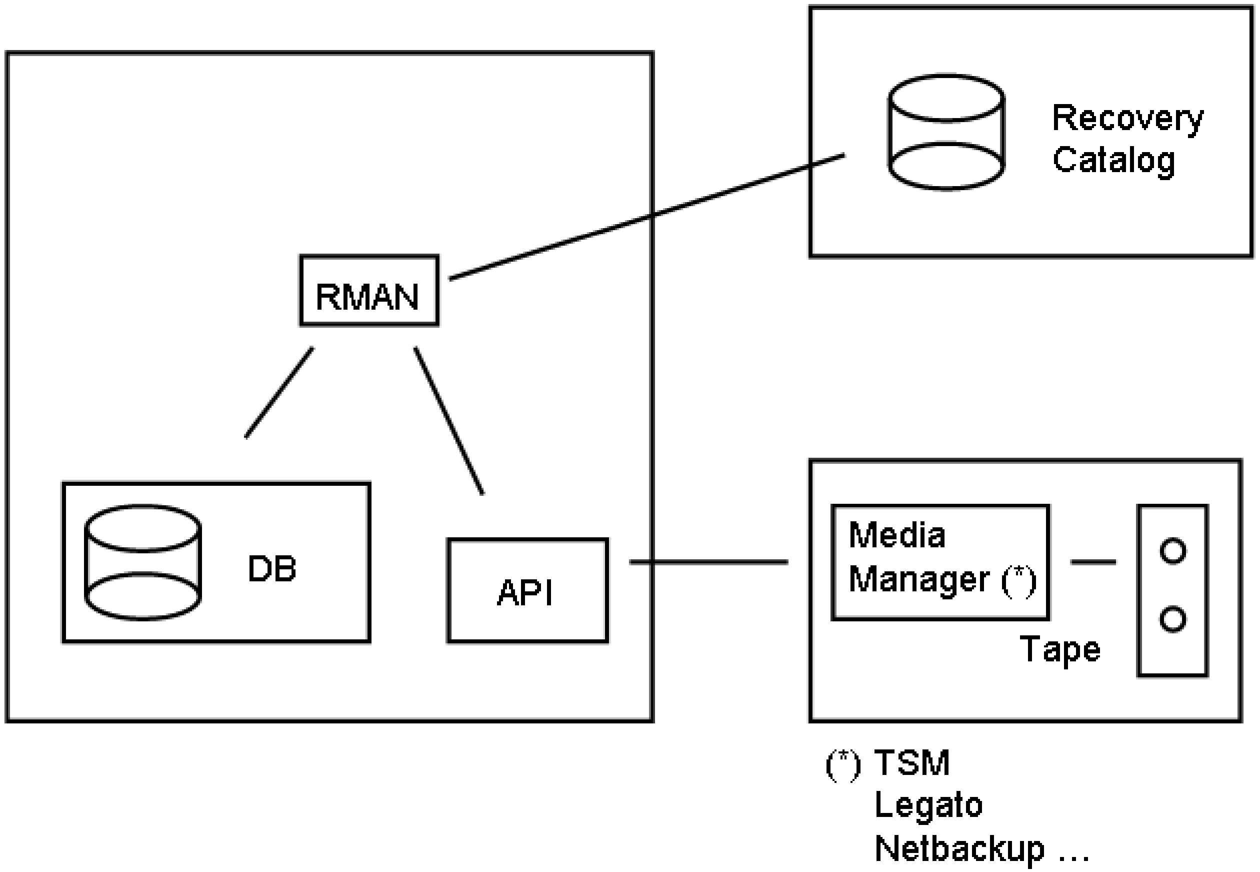
Recovery
Catalog
RMAN
DB
API
Media
Manager (*)
Tape
(*) TSM
Legato
Netbackup …

3.2 Backup/Recovery and Logging: Oracle 5/6

Feature	RMAN	User managed	Export
closed DB Backup	x	x	-
open DB Backup	x (no begin .. end Backup)	x (begin .. end Backup)	x (Undo required for consistency)
incremental Backup	x	-	-
corrupt Block Detection	x	-	x
automatic Backup	x	-	x (full, user, table) - Level

3.2 Backup/Recovery and Logging: Oracle 6/6

Feature	RMAN	User managed	Export
Backup catalog	x	-	-
Backup to Media Manager	x (API)	x	x
Backup init-File, Password – File	x	x	-
OS independent Language	x	-	x

3.2 Backup/Recovery and Logging: Sybase 1/3

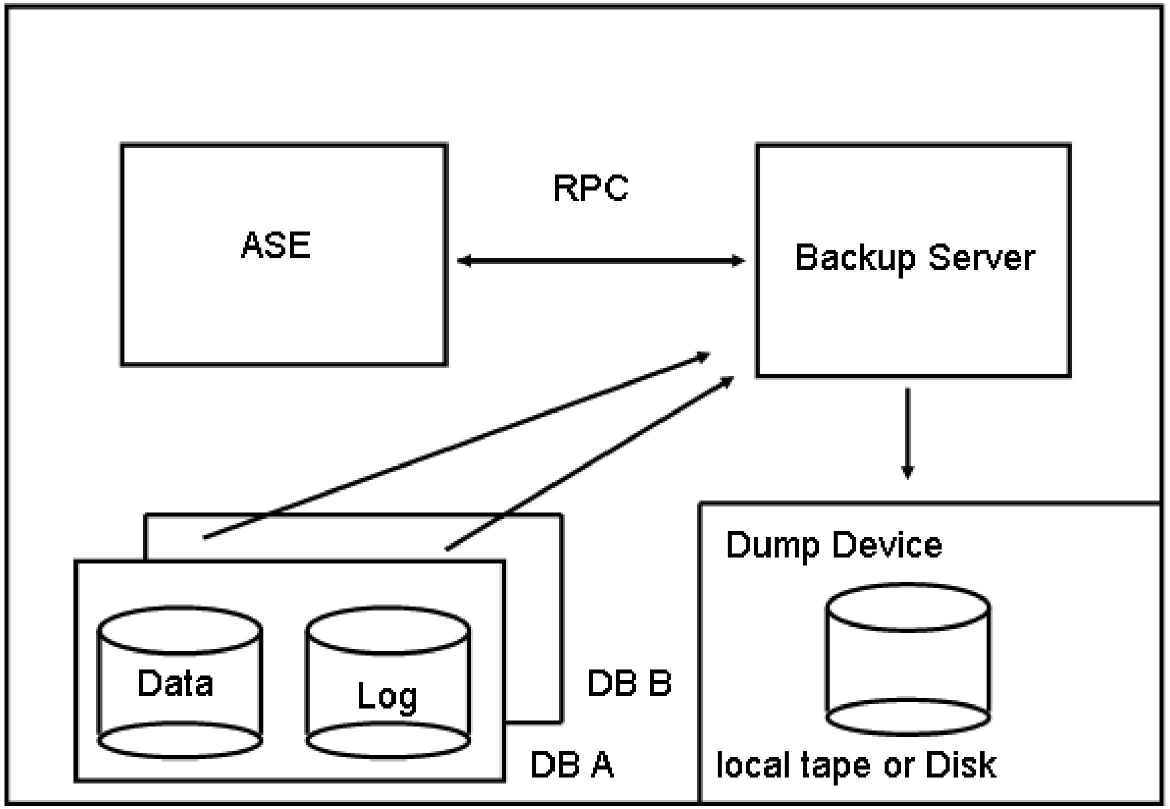

3.2 Backup/Recovery and Logging: Sybase 2/3

Physical Backup

- same Logging Model as SQL Server (per database one Transaction Log)
- dump database (only whole Database)
- dump transaction (thereafter Log will be truncated)
- threshold procedures have to be implemented
- database options may be set:
    - abort transaction on log full
    - truncate log on checkpoint (removal of committed transactions)
- Media Manager Integration through SQL*Backtrack (BMC) or Veritas Netbackup (Veritas)

3.2 Backup/Recovery and Logging: Sybase 3/3

Logical Backups

- bcp (for data, only at object level)
- ddlgen (for structures)

3.2 Backup/Recovery and Logging: SQL Server

Physical Backup

- auto growing tran log

  may be defined otherwise

  stored threshold procedure

- full backup (backup database)

  (local tape or disk)

- full differential backup

  (backup database with

  differential) based on full backup

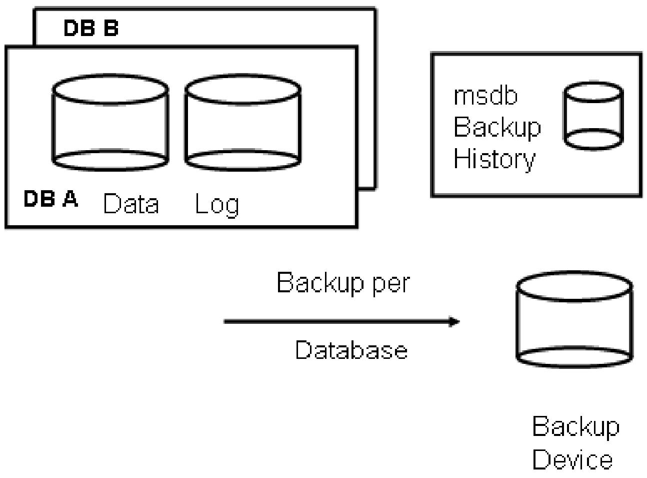

3.2 Backup/Recovery and Logging: SQL Server 2/3

- partial Backup: Backup of write intensive Portions of a Database
- partial differential backup
- File backup, Filegroup Backup: backup of a Database at the File Level (faster Recovery)
- transaction Log backup: default autogrow enabled (backup log) (after Backup it will be truncated)
- API for split Mirror hot Backup available (with suspension of Database Activity)
- no integration of media management software

3.2 Backup/Recovery and Logging: SQL Server 3/3

Logical Backup

- bcp
- DTS (Data Transformation Service)
- new in 2005: SSIS (SQL Server Integration Service) – an ETL Tool

3.2 Backup/Recovery and Logging: DB/2 1/4

Logical Backup

- import/export (sql select in various formats)
  - Formats:
    - IXF (Integration Exchange Format) (see note)
    - WSF (Worksheet File Format) (IBM spreadsheet format)
    - DEL Delimited ASCII (seq. file with row and column del.)
    - ASC Undelimited ASCII (seq. file with row del.)
- load for large quantities of data

3.2 Backup/Recovery and Logging: DB/2 2/4

Logging

- circular Logging
  - Default
  - only logs with current transaction data (active logs) are kept
  - only full and Offline Backups ("Version Recovery")
  - protects from Instance Failure
- retained Logging
  - PITR possible
  - protects from Instance and Media Failure
  - Log Archiving via Userexit (call of external Program: diskcopy, TSM, tape, any XBSA compliant utility)

3.2 Backup/Recovery and Logging: DB/2 3/4

Backup per Database

Log Files

- Transaction Recording

Recovery History File

- Summary of Backup Information to determine Recovery Options

Tablespace Change History File

- determine Log Files for Recovery of a Tablespace

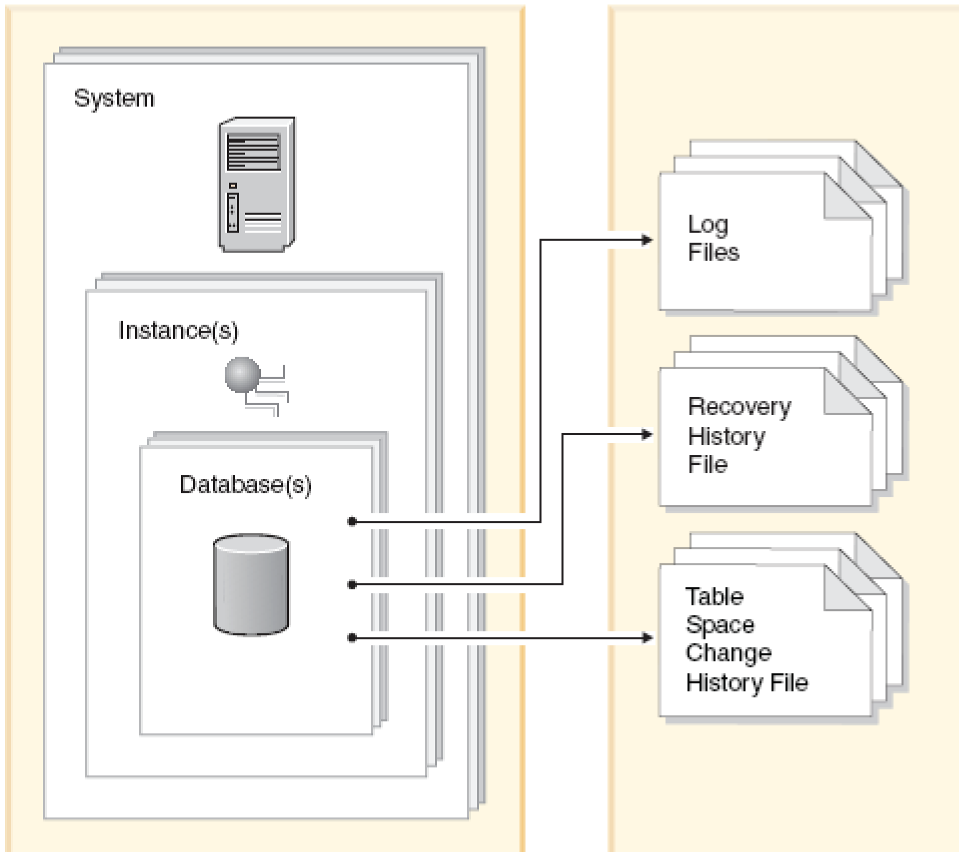

3.2 Backup/Recovery and Logging: DB/2 4/4

Physical Backup

- retained (archive) Logging:
  - db2 backup <database> (Program)
  - db2backup (API)
  - incremental Backups (cumulative, delta) possible
  - split Mirror Backups

-circular Logging:
  - OS Backup only offline (Integration with TSM possible)

3.3 Data Consistency vs. Concurrency 1/27

Fundamental problem of Multi User Databases

How to ensure Data Consistency and maintain Concurrency at the same Time?

RDBMS handle this differently:

- Oracle: Multi Versioning

- Non Oracle (Sybase, DB/2, SQL Server): Read locking

3.3 Data Consistency vs. Concurrency

Multi Versioning

- read consistent queries: consistent results with respect to a given point in time

- non – blocking queries: queries are never blocked by writers of data

- Oracle is able to simultaneously materialize multiple versions of the data.

3.3 Data Consistency vs. Concurrency 3/27

Demo Example

- cursor does not fetch data
- delete command places data in undo segment

```
SQL> create table t
  2  as
  3  select * from all_users;

Table created.

SQL> variable x refcursor
SQL> begin
  2  open :x for select * from t;
  3  end;
  4  /

PL/SQL procedure successfully completed.

SQL> delete from t;

8 rows deleted.

SQL> commit;

Commit complete.

SQL> print x

USERNAME                          USER_ID CREATED
------------------------------ ---------- --------
SYSMAN                                 31 26.06.06
PERFSTAT                               25 26.06.06
DIP                                    19 26.06.06
TSMSYS                                 21 26.06.06
DBSNMP                                 24 26.06.06
SYS                                     0 26.06.06
SYSTEM                                  5 26.06.06
OUTLN                                  11 26.06.06

8 rows selected.

SQL>
```

3.3 Data Consistency vs. Concurrency 4/27

Transaction Isolation Levels

Purpose of a Transaction: Take DB from one consistent State to another.

In a Multiuser DB different Transactions are carried out simultaneously.

The Sensitivity of a Transaction to Changes made by others is determined via *Transaction Isolation Levels*.

Isolation Levels are defined in Terms of Phenomena.

3.3 Data Consistency vs. Concurrency

Phenomenon	Explanation
Dirty Read	read of uncommitted data no data integrity
Nonrepeatable Read	read of the same row at times t1,t2 may yield different results
Phantom Read	query at times t1,t2 (t2 > t1) may yield additional rows at t2

3.3 Data Consistency vs. Concurrency

Transaction Isolation Levels (x: permitted, -: not permitted)

Isolation Level	Dirty Read	Nonrepeatable Read	Phantom Read
L0: Read Uncommitted	x	x	x
L1: Read Committed	-	x	x
L2: Repeatable Read	-	-	x
L3: Serializable	-	-	-

Locking and Default Isolation Levels (Non-Oracle)

- all locks are stored in memory

- locking schemes: row (page), table

- Concurrency and Overhead grow from table level locking to row level locking.

3.3 Data Consistency vs. Concurrency

Sybase	- Allpages Locking (data + index) - Datapages Locking - Datarows Locking - Lock Escalation may occur (page => table, row => table) - shared (for read) and exclusive locks - Default Isolation Level: L1
DB/2	- row Level Locking - Lock Escalation may occur. - Default Isolation Level: L1

SQL Server	- Row Level Locking - Lock Escalation may occur (Row => Table) - Default Isolation Level : L1 - new in 2005: Snapshot Isolation (attempt to mimic Oracle Behaviour, uses tempdb)

3.3 Data Consistency vs. Concurrency 10/27

Locking and Isolation Levels (Oracle)

- Locks are stored in Data Blocks (no lock escalation).
- Locks only when data are modified.
- Oracle *always* uses non-blocking reads (Multi Versioning).
- *Writes do not block Reads and vice versa.*
- L1 and L3 are supported (L1 default)
- Intention of SQL Standard:
  - L0: to mimic non blocking reads
  - L1: does not give consistent results
  - L2: to guarantee read consistent results

3.3 Data Consistency vs. Concurrency 11/27

However, in Oracle:

- L1: because all queries are read consistent, it is equivalent to L2

- L0: The "spirit" of L0 is in Oracle always achieved (all reads are non – blocking)

- Additionally L4 (Read Only) Level is supported (L3 + no data modification).

  - Transaction sees only changes, that were committed when transaction began, but IUD are not allowed.

Example: Take a look at a Transaction in all Isolation Levels and compare Oracle- and Non-Oracle Behaviour.

Row	Table Accounts		
	Account (PK)	Balance (t1)	Balance (t8)
1	123	500.00	100.00
2	456	240.25	240.25
3	781	0.00	0.00
4	987	100.00	500.00
Sum		840.25	840.25

2 simultaneous Sessions:

- Session 1: select sum(balance) from account;
- Session 2: transfer 400.00 from Account 123 to Account 798.

Caption:

S: Session	O: Oracle like Database
R: Row	NO: Non – Oracle like Database
L: Level	
t: time	

3.3 Data Consistency vs. Concurrency 14/27

Read Uncommitted (L0) (Oracle does not support L0, so L1 is used.)

Time	S1 (NO) : L0	S1 (O) : L1	Transfer
T1	R1=500.00, Sum=500,00	R1=500.00 Sum=500.00	
T2	R2=240.25 Sum=740.25	R2=240.25 Sum=740.25	
T3			Update R1, excl. Lock on R1 set R1=100.00
T4	R3=0.00 Sum=740.25	R3=0.00 Sum=740.25	

Time	S1 (NO) : L0	S1 (O) : L1	Transfer
T5			Update R4, excl. Lock on R4 set R4=500.00
T6	R4=500.00 Sum=1240.25	R4=100.00 Sum=840.25 (see Note)	
T7			Commit
T8	S=1240.25 Wrong!	Sum=840.25	

Read Committed (L1)

Time	S1 (NO) : L1	S1 (O) : L1	Transfer
T1	R1=500.00 Sum=500.00	R1=500.00 Sum=500.00	
T2	R2=240.25 Sum=740.25	R2=240.25 Sum=740.25	
T3			Update R1; excl. Lock Set R1=100.00
T4	R3=0.00 Sum=740.25	R3=0.00 Sum=740.25	

Time	S1 (NO) : L1	S1 (O) : L1	Transfer
T5			Update R4; excl. Lock R4=500.00
T6	try to read R4 (is locked) S1 blocks, Query stops	R4=100.00 (from Undo) Sum=840.25	
T7			Commit;
T8	R4=500.00 S=1240.25 Wrong!	-	
The same Isolation Levels yield different answers!			

3.3 Data Consistency vs. Concurrency 18/27

Repeatable Read (L2) (not supported/needed in Oracle, uses L1)

Time	S1 (NO) : L2	S1 (O) : L1	Transfer NO	Transfer O
T1	R1=500.00 Sum=500.00 Lock R1	R1=500.00 Sum=500.00		
T2	R2=240.25 Sum=740.25 Lock R2	R2=240.25 Sum=740.25		
T3			try to update R1 (is locked) Update suspended	Update R1; excl. Lock Set R1=100.00
T4	R3=0.00 Sum=740.25 Lock R3	R3=0.00 Sum=740.25		

3.3 Data Consistency vs. Concurrency

Time	S1 (NO) : L2	S1 (O) : L1	Transfer NO	Transfer O
T5				Update R4; excl. Lock Set R4=500.00
T6	R4=100.00 Sum=840.25 Lock R4	R4=100.00 (from Undo) Sum=840.25		
T7				Commit;
T8	Commit; release all Locks			

Time	S1 (NO) : L2	S1 (O) : L1	Transfer NO	Transfer O
T9			Update R1 Set R1=100.00 excl. Lock	
T10			Update R4 Set R4=500.00 excl. Lock	
T11			Commit;	

Remarks

- Correct Answer also in NO.
- Cost: Transactions had to be executed sequentially!
- Readers of Data will block Writers of Data!
- Spurious Errors still possible: see next Page.

Time	S1 (NO) L2	Transfer 50.00 from R4 to R1
T1	R1=500.00 Sum=500.00 Lock R1	
T2	R2=240.25 Sum=740.25 Lock R2	
T3		Update R4=50.00; excl. Lock
T4	R3=0.00 Sum=740.25	
T5		try to Update R1 (is locked) Update suspended
T6	try to read R4 (is locked) Query suspended	
Classical Deadlock One Transaction will be chosen as victim, killed and rolled back.		

Porting Issues

In DB with shared Read Locks and L2: Lost Updates cannot happen.

But:

In DB without shared Read Locks: Lost Update Problem!

Lost Update Problem (Database without Shared Read Locks)

T1: S1: read R1

T2: S2: read R1

T3: S1: Update R1; Commit;

T4: S2: Update R1; Commit:

⇒ Result from T3 is lost!

How to avoid: Manual Locking (optimistic or pessimistic)

optimistic Locking: defer locking as late as possible (Application Design!)

pessimistic Locking: select for update (lock for write, not for read)

3.3 Data Consistency vs. Concurrency 25/27

Serializable (L3)

Each Transaction appears to be the only Transaction in the Database at a given Time.

Oracle:

- L3 extends the Read Consistency from Statement to Transaction Level (using Undo).

- Price: ORA-8177: can't serialize access for this transaction

  (when updating data which has changed since the begin of the transaction)

Non Oracle:

- locking may yield Deadlocks, Blocking, …

3.3 Data Consistency vs. Concurrency 26/27

Read Only (L4) (Oracle Only)

- as L3 but no modifications allowed => no ORA-8177
- in other systems: L2
- because in Oracle without Locking: ORA-1555 might happen.

3.3 Data Consistency vs. Concurrency 27/27

Isolation Level	Implementation	Writes Block Reads	Reads Block Writes	Deadlock Sensitive Reads	Incorrect Query Results	Lost Updates	Lock Escalation or Limits
L0	NO	n	n	n	y	y	y
L1	NO	y	n	n	y	y	y
	O	n	n	n	n	n (*)	n
L2	NO	y	y	y	n	n	y
L3	NO	y	y	y	n	n	y
	O	n	n	n	n	n	n

3.4 High Availability Concepts: Oracle 1/14

- Real Application Cluster
- Flashback
- Data Guard
- Replication (Streams)

(a) Real Application Cluster (RAC) and Oracle ClusterWare (OCW)

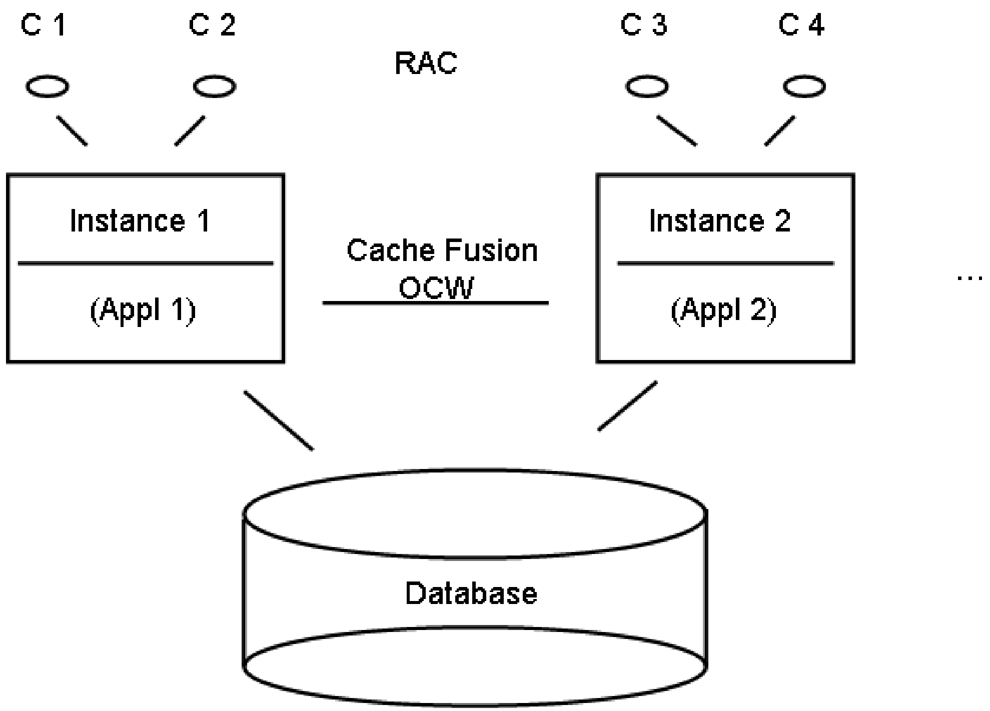

High Availability and horicontal Scaling

3.4 High Availability Concepts: Oracle

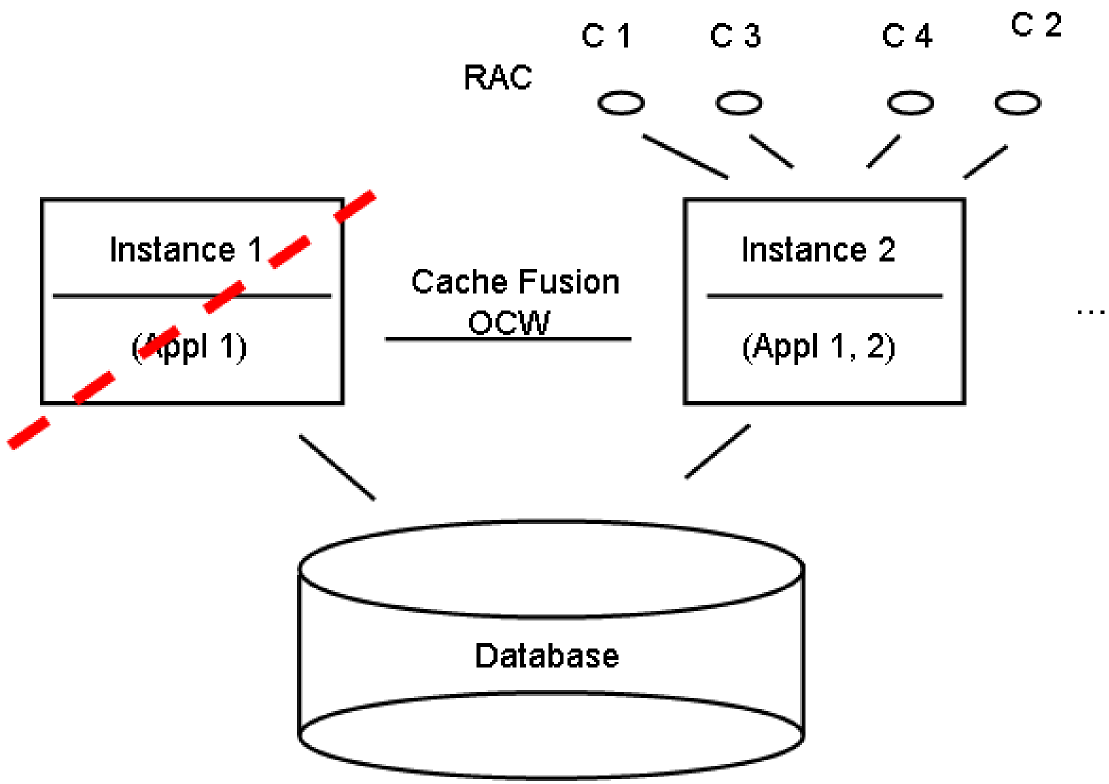

3.4 High Availability Concepts: Oracle 4/14

Oracle Clusterware

- comes from Tru-Cluster (DEC)

- OCW + ASM:
  - Ressource Monitoring
  - Ressource Management
  - Cluster Control
  - Failover APIs
  - Cluster Membership
  - LVM

3.4 High Availability Concepts: Oracle 5/14

OCW:

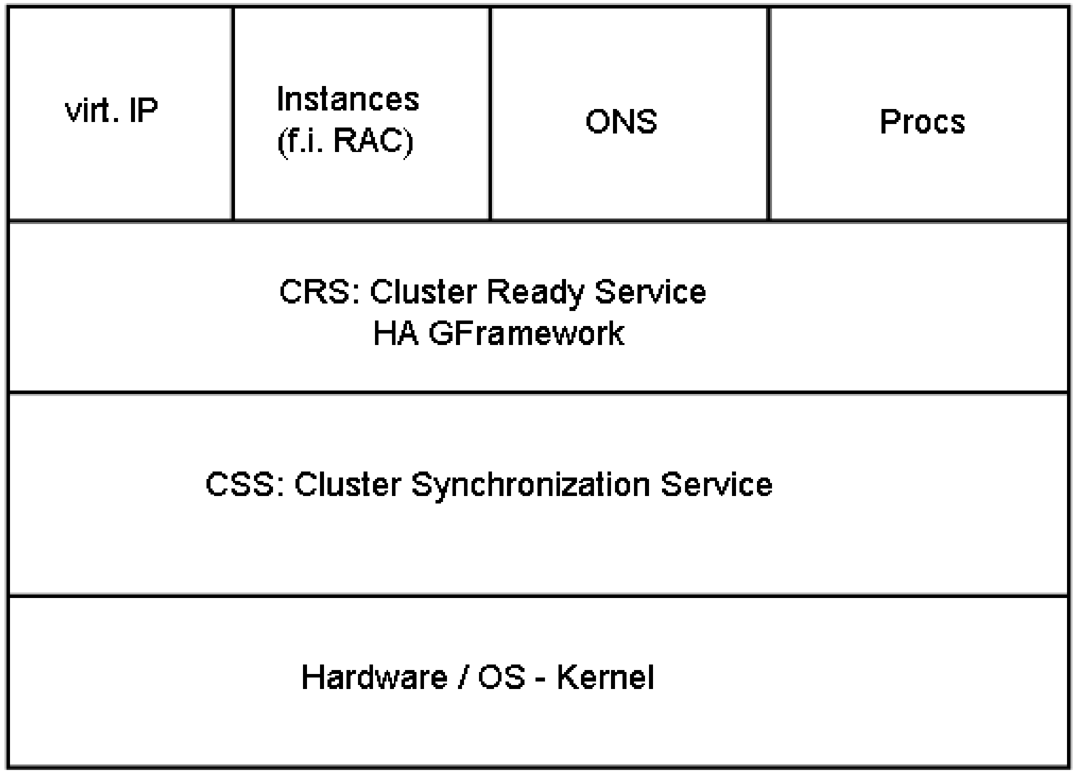

HVB Information Services Member of UniCredit Group

3.4 High Availability Concepts: Oracle

additional features of OCW:

- Cluster for any Application
- Loadbalancing, Workloadmanagement
- Fast Application Notification (FAN) for subscribed Clients
- Transparent Application Failover (TAF)

(b) Flashback – Protection from logical Errors

Feature	Purpose
Flashback Versions Query	to reproduce historical Operations on Database Objects
Flashback Transaction Query	generate Undo-SQL for historical Operations
Flashback Table	set Database Objects back in Time
Flashback Drop	restore Database Objects after dropping them
Flashback Database	set whole Database back in Time

(c) DataGuard (Standby Database)

Type of Standby Database	Purpose /Protection from
Physical Standby Database	- User Errors, logical Corruptions - Disaster Recovery - Fast Site and Database Failover - Planned Switchover for Maintenance - Divergence for Reporting and Testing (Flashback) - taking Backups / Load ! - Read Only Capability

3.4 High Availability Concepts: Oracle 9/14

Type of Standby Database	Purpose /Protection from
Logical Standby Database	- Disaster Recovery - Read/Write Capability for normal Operation - additional Objects may be created for Reporting - Rolling Database Upgrades

Maximum Availability Architecture: RAC + DataGuard

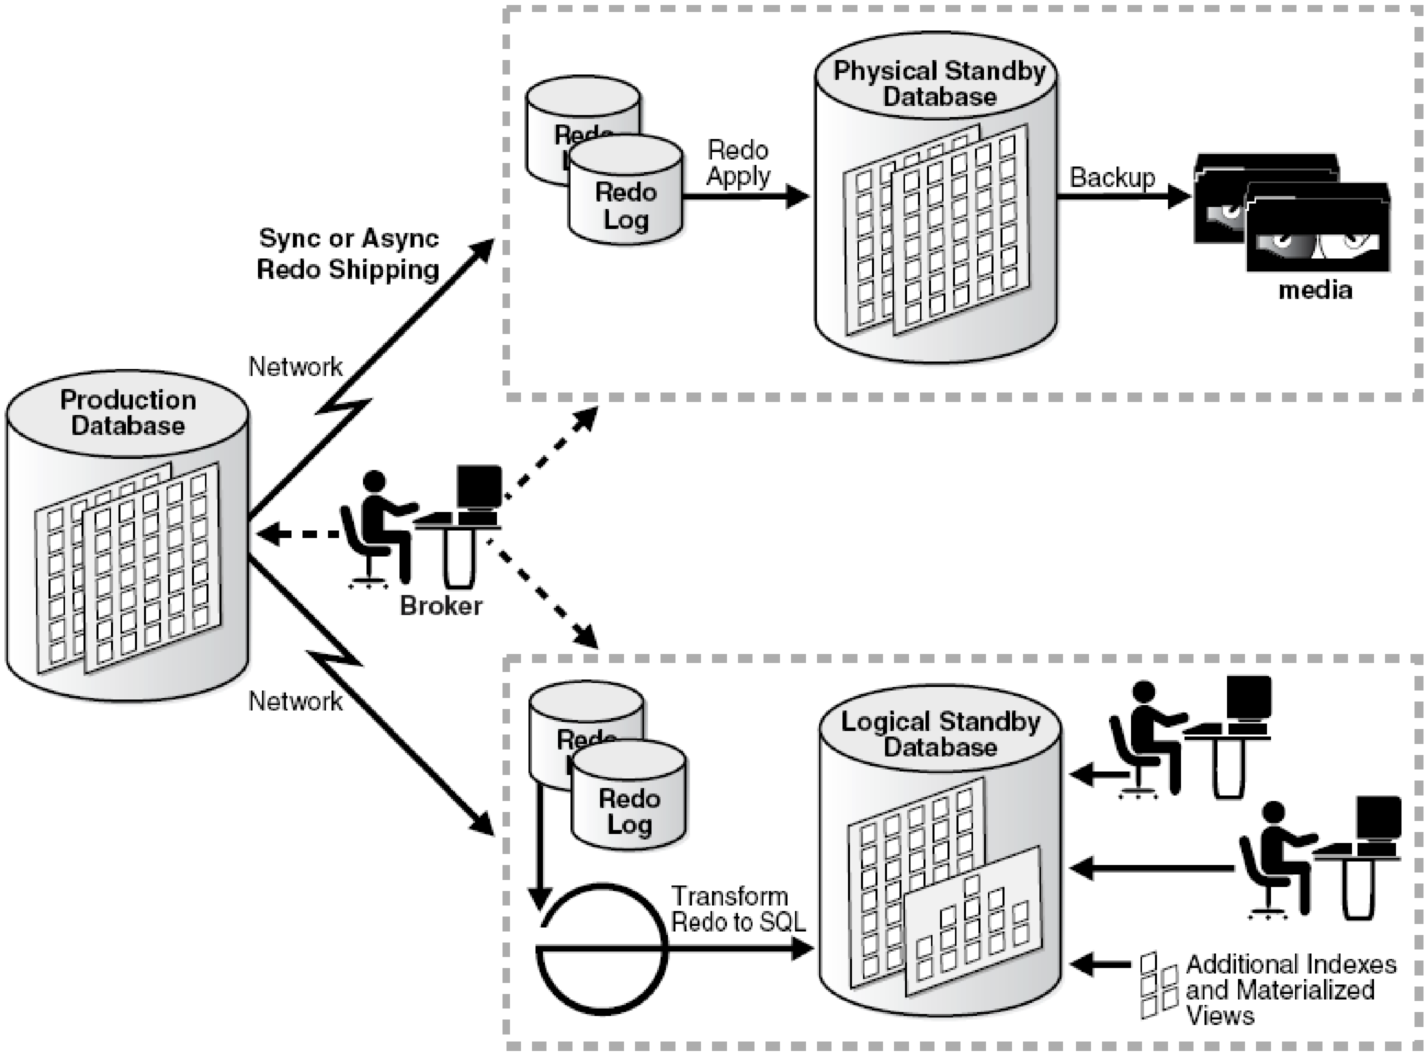

Physical Standby
Database
Redo
Log
Redo
Apply
Backup
media
Sync or Async
Redo Shipping
Network
Production
Database
Broker
Network
Redo
Log
Logical Standby
Database
Transform
Redo to SQL
Additional Indexes
and Materialized
Views

3.4 High Availability Concepts: Oracle 11/14

(d) Replication: Oracle Streams

- highest Degree of Flexibility

- highly customizable

- heterogeneous Platform Support

- high Configuration and Administration Flexibility

1. capture Messages at a Source Database
2. stage Messages in Queues
3. propagate Messages to another Queue
4. consume Messages
5. apply Changes in Destination Database

3.4 High Availability Concepts: Oracle 13/14

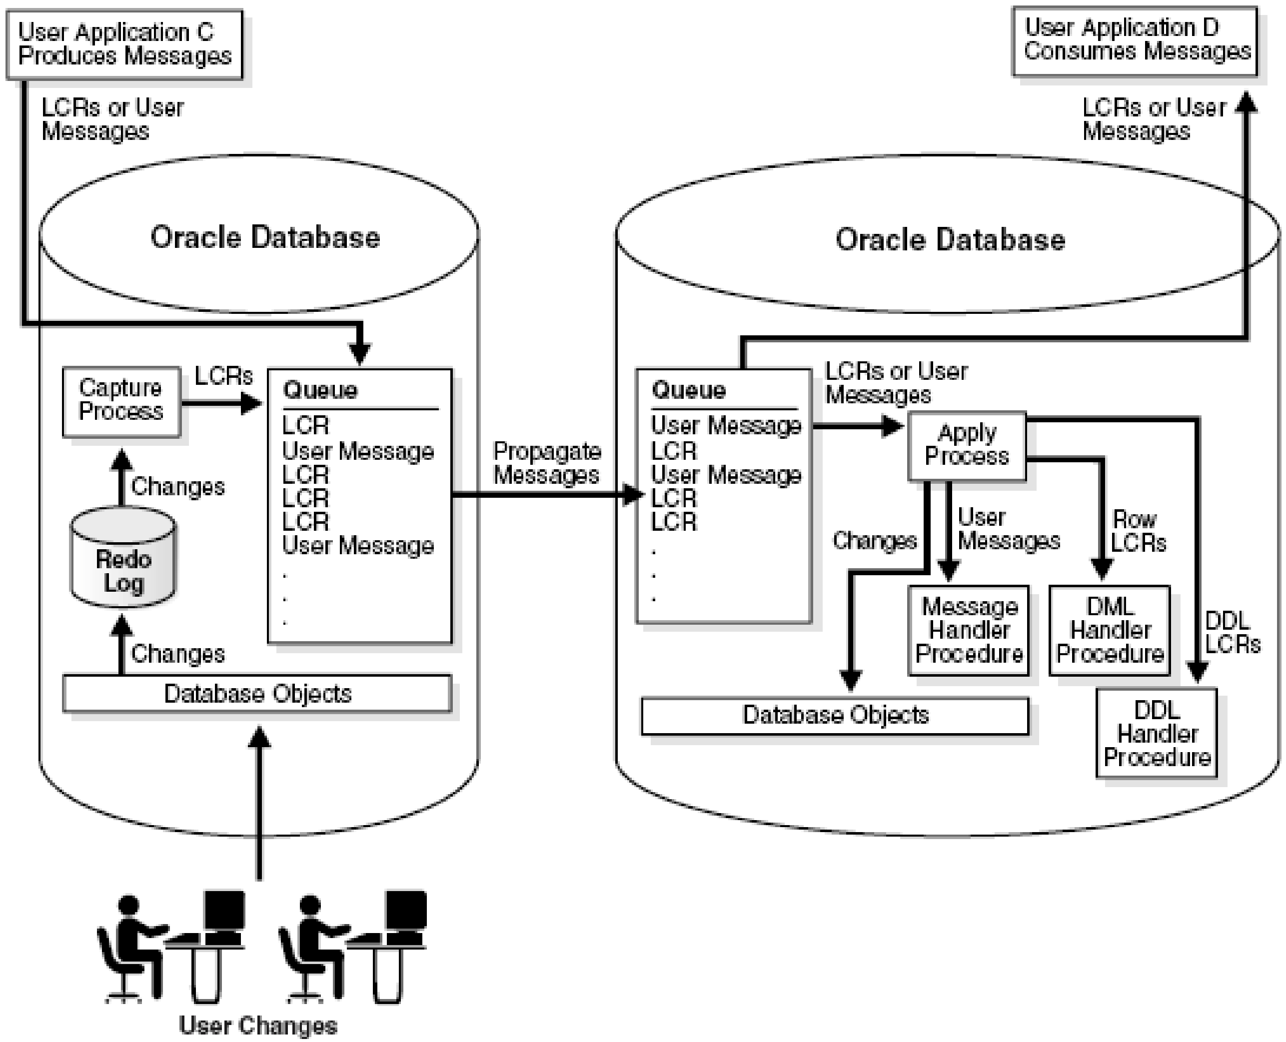

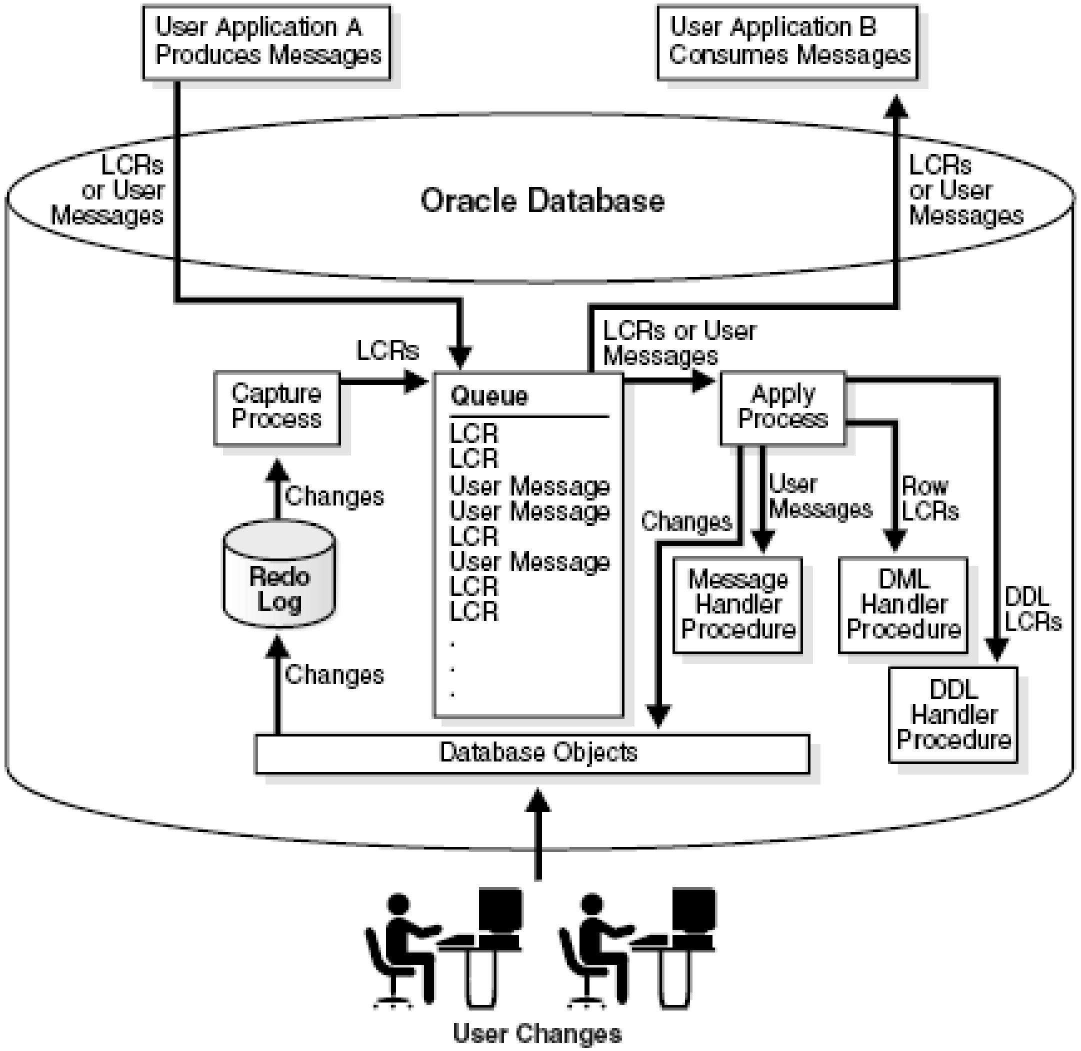

HVB Information Services
Member of UniCredit Group

3.4 High Availability Concepts: Sybase 1/9

(a) Dump/Load (Poor Man's Solution)

Primary	Secondary
Dump DB	
	Load DB
Dump Transaction Log	
	Load Transaction Log
Dump Transaction Log	
	Load Transaction Log
Fail	

3.4 High Availability Concepts: Sybase 2/9

(a) Dump/Load (Poor Man's Solution)

Primary	Secondary
Dump Transaction Log (if possible)	
	Load Transaction Log (if available)
	Online DB
	Point Clients here (manually or automatically)

- Drawback: secondary unavailable prior Online DB

3.4 High Availability Concepts: Sybase 3/9

(b) Dump/Load (Poor Man's Solution) with Standby Access

Primary	Secondary
Dump DB	
	Load DB
Dump Transaction Log with standby_access	
	Load Transaction Log Online DB for standby_access
Dump Transaction Log with standby_access	

3.4 High Availability Concepts: Sybase 4/9

(b) Dump/Load (Poor Man's Solution) with Standby Access

Primary	Secondary
	Load Transaction Log Online DB for standby_access
Fail Dump Transaction Log (if possible)	
	Load Transaction Log (if available)
	Online DB

(b) Dump/Load (Poor Man's Solution) with Standby Access

Primary	Secondary
	Point Clients here (automatically or manually)

Advantage: Secondary available for non change Queries during normal Operation

3.4 High Availability Concepts: Sybase 6/9

(c) Database Device Mirroring: Protection against Disk Failures

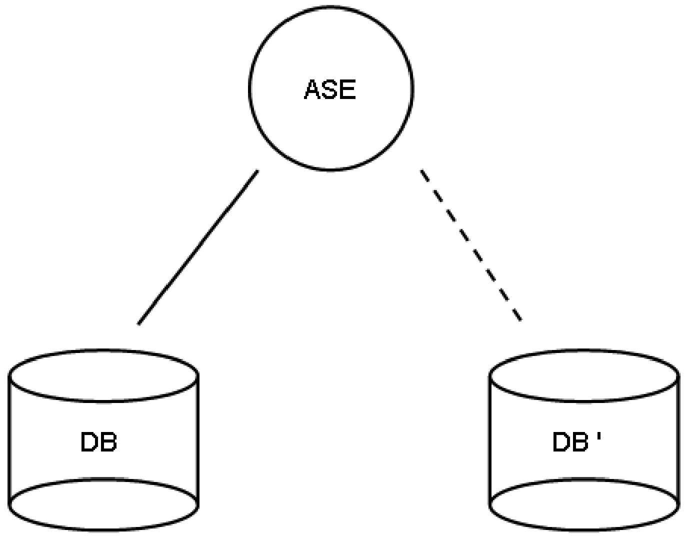

(d) Replication Server (or ASE Replicator)

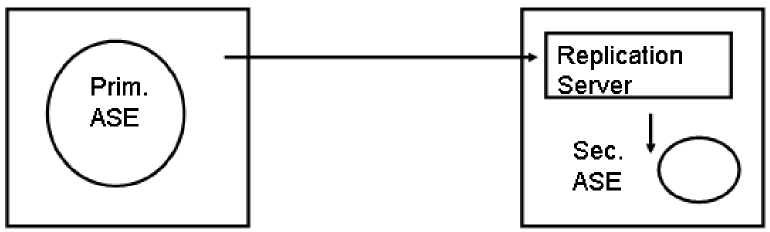

- Publish/Subscribe Model
- Warm Standby Model (Database Granularity, also DDL Replication)
- Latency Period
- automatic Failover
- automatic Client Failover (includes Application Change)

(e) Open Switch

- no application change resp. client redirect

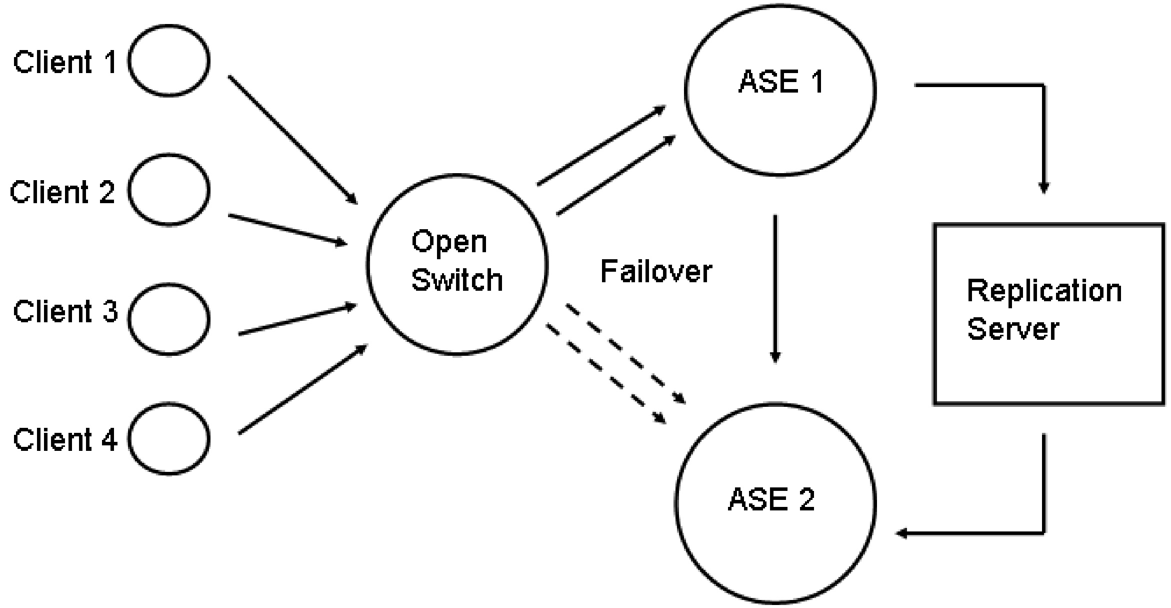

3.4 High Availability Concepts: Sybase 9/9

(f) HA Companion Server

- depends on Cluster Technology
    - SunCluster
    - HACMP
    - MC/ServiceGuard
    - VeritasCluster
- ASE 2 in Standby Mode or with own Databases
- Client Redirect automatic through network interfaces

3.4 High Availability Concepts: SQL Server 1/4

(a) Database Mirroring

- per Database
- possible in conjunction with Microsoft Cluster Service
- mirror Database available for limited reporting Needs
- optional "Witness" for automatic Failover
- Client redirection is accomplished through local Caching of the Server Names

3.4 High Availability Concepts: SQL Server 2/4

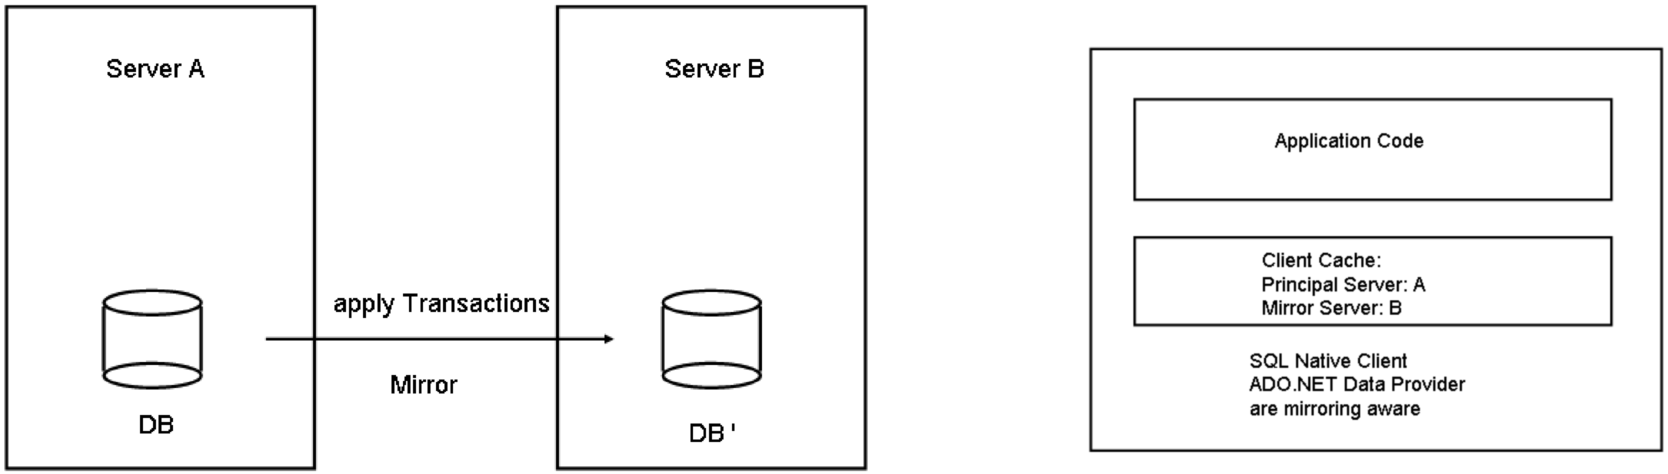

3.4 High Availability Concepts: SQL Server 3/4

(b) Database Snapshots

- read only Databases as exact copy of a User Database in a specific Point of Time

  - at Page Level
  - Copy – on – Write (only modified Blocks will be copied)
  - for Reporting
  - for Recovery Purposes

3.4 High Availability Concepts: SQL Server 4/4

(c) Replication

- copy Databases and Database Objects from one Database to another and synchronize
- transactional
- also as multi-master variant (peer – to – peer)
- Replication with non SQL Server Databases is supported

(d) Log Shipping

- same Principle as Oracle Standby Database

3.4 High Availability Concepts: DB/2 1/7

(a) HADR (High Availability Disaster Recovery)

- takeover command for role switch
- automatic client reroute (client has both connect information)

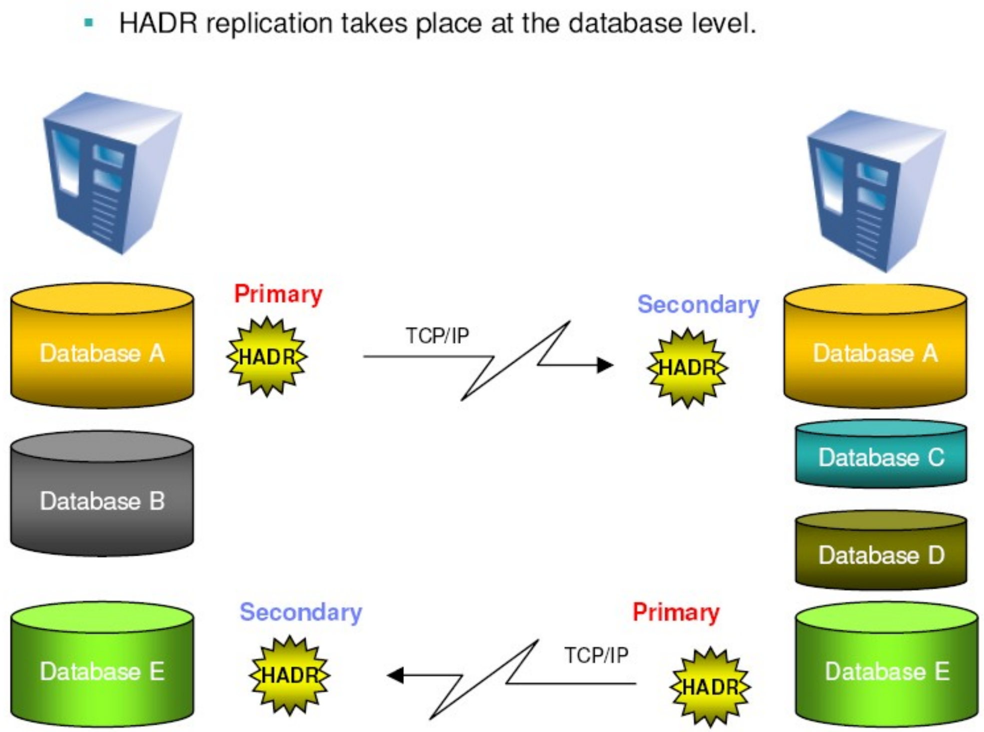

HVB Information Services Member of UniCredit Group

Running HADR

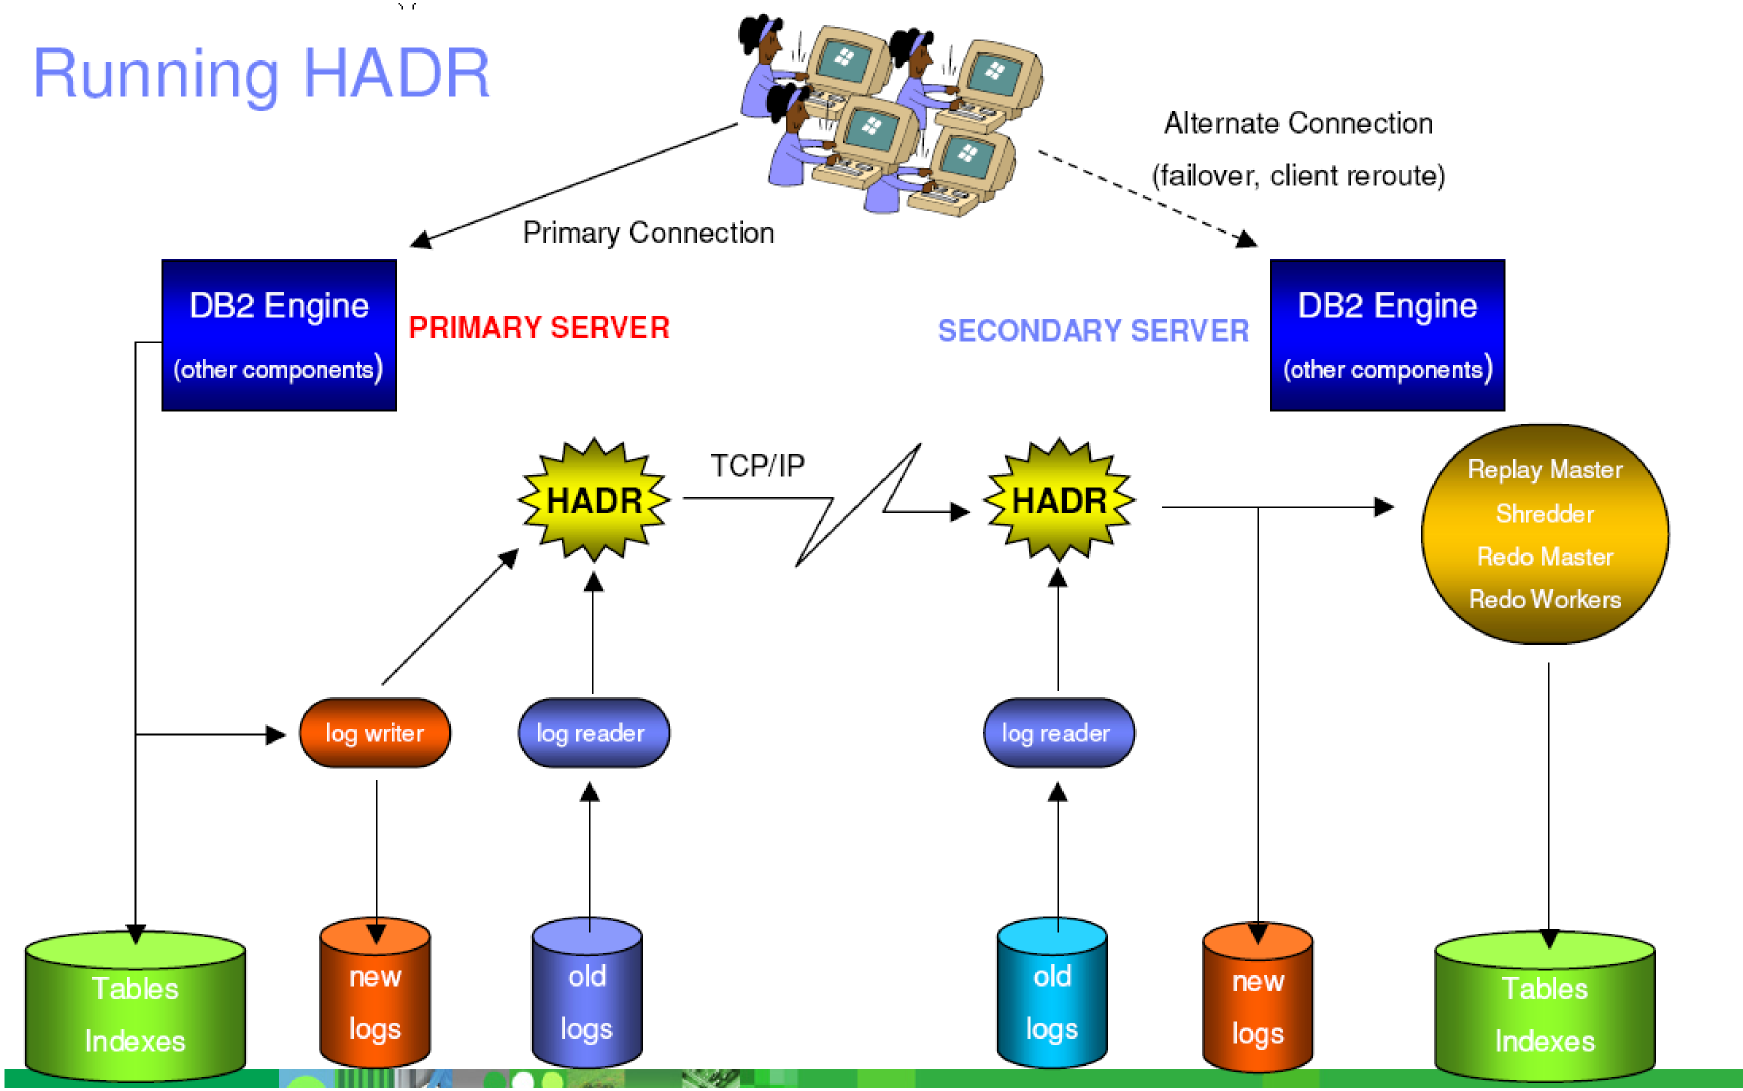

- Synchronous, Near Synchronous, Asynchronous

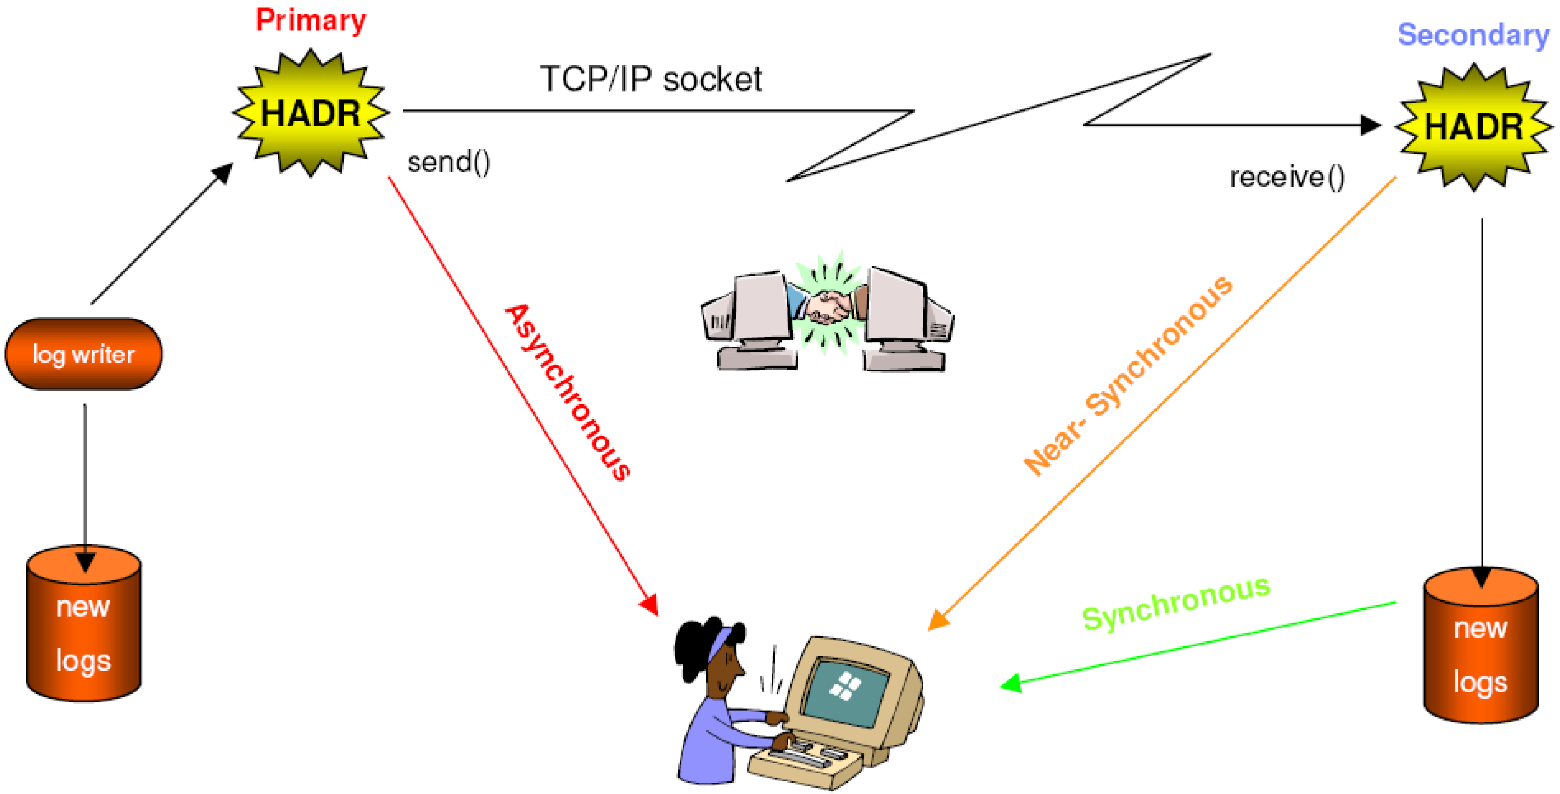

3.4 High Availability Concepts: DB/2 4/7

(b) Standby Database

- create Standby DB via Split Mirror or by Backup/Restore
- Log Shipping:
  - copy whole Log Files to standby Machine
    - from Archive Device
    - or directly via User Exit
  - secondary continuously rolls forward through Logs
  - when Primary fails:
    - copy over remaining Logs

3.4 High Availability Concepts: DB/2 5/7

- db2 rollforward database <name> to end of logs and stop
- Clients reconnect to Standby Database

\- Replication:

- use Data Propagator for Log Capturing at Primary
  - (Capture and Apply Subcomponents)
  - Standby DB available for Read Access
  - BiDi Replication possible
  - no need for physically identical Storage Partitions
  - Restriction to critical Tables possible

(c) MQ Replication

- Each message represents a transaction
- XML format option for publishing
- Highly parallel apply process
- Differentiated conflict detection and resolution
- Staged availability of heterogeneous support

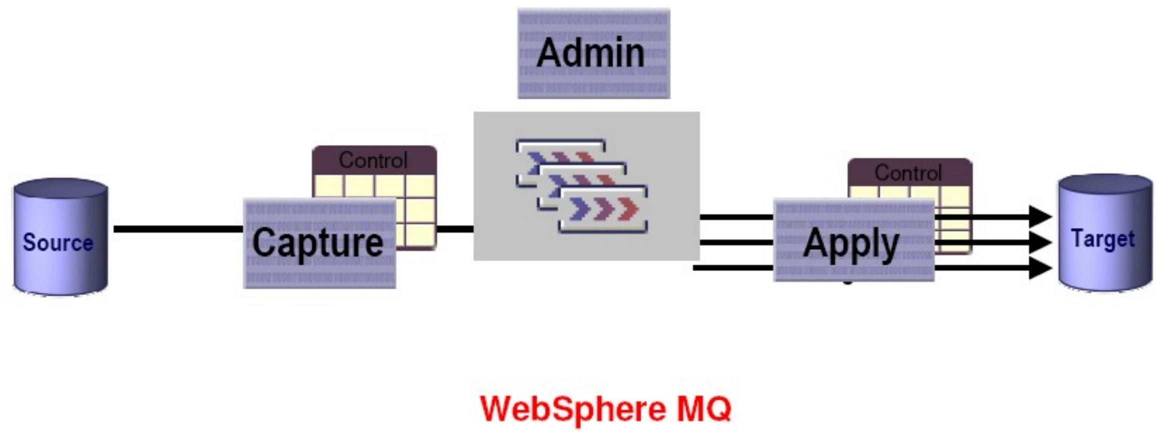

(d) Data Partitioning Feature

- not for High Availability but for Scalability
- Shared Nothing Architecture
- Agents may fail
- Hash Partitioning

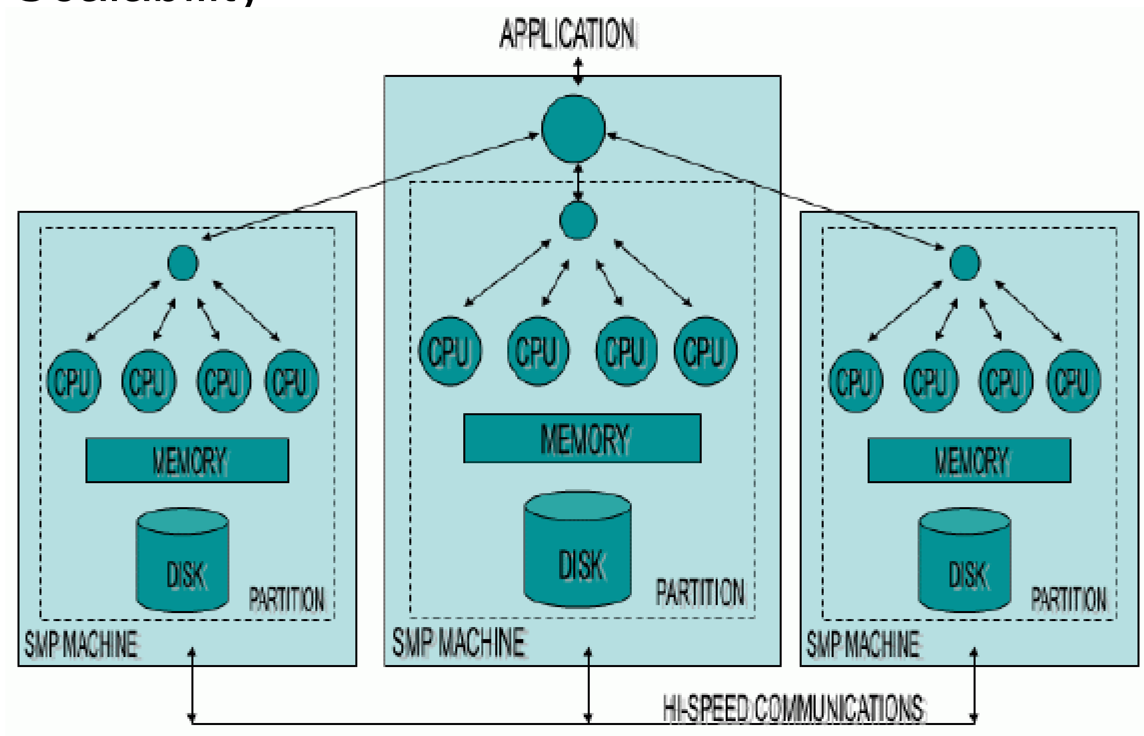

4. Vendor Comparison: Beyond Technology

4.1 Oracle Product Portfolio 1/3

Database	Oracle RDBMS
Middleware	Oracle Fusion Middleware - Application Server - Identity Management - Directory Services
Developer Tools	- Developer Suite ■ Developer ■ Discoverer ■ Designer - BPEL Designer (see note)

4.1 Oracle Product Portfolio 2/3

Enterprise Management	Administration Interface and Repository for managing all Database Instances (Grid Control)
Business Intelligence Project "Fusion" (all Products have to be merged)	- Oracle E-Business Suite - Peoplesoft Enterprise - Siebel CRM - JD Edwards

4.1 Oracle Product Portfolio 3/3

Collaboration	Collaboration Suite - Content Services - Real Time Collaboration - Messaging - Workspaces
Data Warehouse	- Oracle Warehouse Builder - Oracle OLAP

4.1 Sybase Product Portfolio 1/3

Database (OLTP)	Adaptive Server Enterprise (ASE)
Heterogeneous Data Services	- Sybase IQ (BI, DWH) - Replication Server - Dynamic Archive - Mirror Activator - Real Time Data Services

4.1 Sybase Product Portfolio 2/3

Information Anywhere	- SQL Anywhere - Adaptive Server Anywhere (ASA) - RFID Anywhere/Enterprise - Mobile Link - Afaria - Content Integrator - Enterprise Portal - Unwired Accelerator - Unwired Orchestrator

4.1 Sybase Product Portfolio 3/3

On Demand Mobility Services	- AvantGo
Development	- Workspace - Power Designer - Power Builder Family - EA Server

4.2 Platform Support / Market Shares 1/3

Operating System	Oracle	DB/2 UDB 8.2	Sybase ASE/RS 15	SQL Server 2005
Unix				
■ Solaris	x	x	x	-
■ AIX	x	x	x	-
■ HP-UX	x	x	x	-
■ Tru64	x	-	x	-
■ Linux	x	x	x	-
■ Irix (SGI)	<= 8.1.7	-	x	-

HVB Information Services
Member of UniCredit Group

4.2 Platform Support / Market Shares 2/3

Operating System	Oracle	DB/2 UDB 8.2	Sybase ASE/RS 15	SQL Server 2005
Open VMS	x	-	-	-
z/OS (OS 390) BS 2000	x <= 9.2	x -	- -	- -
OS 400	-	x	-	-
Microsoft Windows	x	x	x	x
Novell Netware	<= 8.1.7	-	-	-
Mac OS (>= X)	x	-	x	-

4.2 Platform Support / Market Shares (see note) 3/3

Vendor	Sales Volume 2005 (bn. USD)	Market Share 2005 (%)	Sales Volume 2004 (bn. USD)	Market Share 2004 (%)
Oracle	6.721	48.6	6.234	48.9
IBM	3.040	22.0	2.860	22.4
Microsoft	2.073	15.0	1.778	13.9
Teradata	0.441	3.2	0.412	3.2
Sybase	0.407	2.9	0.383	3.0
Other	1.135	8.2	1.090	8.5
Total Market	13.817	100.0	12.758	100.0

4.3 Cost Comparison

The numbers in this chapter are from

- Microsoft SQLServer 2005: Understanding Database Pricing (11/05, www.microsoft.com)

4.3 Cost Comparison: Pricing - Editions

Edition	Features	Price/CPU (1000 USD / CPU)
Free	- limited DB functionality - memory, CPU, DB size limits	0
Basic	- simplified Mgmt. Tools - basic Security - CPU, Memory Limits	0.5 – 5.0

4.3 Cost Comparison: Pricing - Editions

Edition	Features	Price/CPU (1000 USD / CPU)
Standard	- full DB functionality - Basic Mgmt. Tools - CPU limits	5.0 – 15.0
Enterprise	- HA - Scalability - High End Management Tools - Enterprise Security - no Limits	25.0 – 40.0

4.3 Cost Comparison: Pricing - Options

	SQL Server 2005	Oracle 10g	DB/2 8.2
EE	1.0	1.6	1.0
Mgmt. AddOn	0.0	0.5	0.8
OLAP AddOn	0.0	0.8	1.1
DataMining AddOn	0.0	0.8	2.4

4.3 Cost Comparison: SQL Server - Oracle

	Oracle10g		SQL Server 2005	
Standard		1.0		0.4
Enterprise	EE Mgmt. Pack Adv. Sec. BI	2.7 0.4 0.7 2.7	EE Mgmt. Tools Netw. Encr., SSO, PKI Integr., Rep. Analysis Serv.	1.7 0.0 0.0 0.0
HA	add. Inst. EE + above Opt.	6.5	add. Inst. EE	0.0

4.3 Cost Comparison: SQL Server – DB/2

	DB/2 8.2		SQL Server 2005	
Standard		1.0		0.8
Enterprise	EE Warehouse Manager OLAP Server	3.3 1.4 3.7	EE Integr.Services Reporting + Analysis Serv.	3.3 0.0 0.0
HA	add. Inst. EE + above Opt.	8.4	add. Inst.	0.0

4.3 Cost Comparison: TCO Saving SQL Serv. vs. Comp.

Saving = 100%- TCO(MSSQL/Comp.)

Cost factor	Departm. App. (50-100 Users)	Organisational App. (100+ Users)
Ongoing Operations	33%	45%
Design/Development	24%	13%
Software Licensing	67%	67%
Hardware	54%	54%
Maintenance	56%	56%
Training	54%	60%

4.4 Total Cost of Administration: SQL Server vs. Oracle

The numbers in this chapter are from

- Microsoft SQL Server and Oracle Databases: A Comparative Study on Total Cost of Administration (TCA) (05/06, Alinean, Inc., www.alinean.com)

4.4 Total Cost of Administration: SQL Server vs. Oracle

Figure	SQL Server	Oracle
average Number of DB per company	107	87
average Number of Users per Database	328	716
mission critical Databases (%) (see note)	66.1	63.8
OLTP Databases (%)	55.7	60.3
DSS Databases (%)	44.3	39.7
Databases per DBA	31.2	9.9
Users per DBA	6784	5567
annual TCA per Database (USD)	2847	10206
annual TCA per User (USD)	13.09	18.15

4.4 Total Cost of Administration: SQL Server vs. Oracle

DBA Training (Hours/Year):

- SQL Server: 52.2
- Oracle: 63.1

Database Administrator Experience Levels (%)

	SQL Server DBA	Oracle DBA
Junior (0-5)	38	25
Intermediate (5-8)	38	38
Senior (8+)	24	37

4.4 Total Cost of Administration: SQL Server vs. Oracle

Average Time / Week / Database (Hours) (see note)

Activity	SQL Server	Oracle
Design	1.1	2.2
Tuning	0.8	1.8
Problem	0.8	2.1
Security	0.4	1.2
Capacity	0.3	0.8
Backup	0.7	1.3
Reporting	0.5	1.6

Excursus 1: Oracle Basic Design Decisions 1/2

Under market pressure design decisions had to be done which up to now influence the product.

- Transactional concept: non blocking queries, no read locks (see Note 1)
- Architecture: Development on PDP-11 (see Note 2)
  - Two Task: separation of application process and DB engine
    - comm. with DB engine per IPC (Message Passing)
    - sharing of control information in Shared Global Area (SGA)
    - foundation for SQL*Net

Excursus 1: Oracle Basic Design Decisions 2/2

- 1985: IBM world market leader (DBMS)

  OLTP: IMS (no Ad Hoc, instead precompiled queries)

  DSS: DB/2

- Oracle: small budget => Ad Hoc Queries without precompilation

  Parsing

  Version 2: Rule Based Optimizer (heuristic, static rules)

Excursus 2: Early Tool Development 1/3

- Oracle 2: Appl. Programming required Knowledge in proc. Language
  - C, Fortran, Pascal, Cobol
  - expensive, difficult for On-Line User Interfaces
- Oracle had available:
  - UFI (User Friendly Interface, => SQL*Plus)
  - RPT (Report Production Utility)

Excursus 2: Early Tool Development 2/3

- Oracle wanted to create a procedural environment from scratch => was abandoned
  - without proper investigation of
    - user requirements
    - performance considerations
    - mathematical properties of the procedural language
- 1988: PL/SQL (came from public ADA + Syntax for embedding SQL)
- UFI => AFI (Advanced Friendly Interface)
  - Goal:
    - User can edit arbitrary Ad Hoc commands
    - command buffer exchange to system editor (see note)

Excursus 2: Early Tool Development 3/3

- 1985 – 1987: Oracle was used for small OLTP applications
- Oracle started Development of nonprocedural Tools for AdHoc Query Formulation
  - Easy*SQL, FastForms (Vers. 4.1) => did not survive
  - IAF (Interactive Application Facility)
  - breakthrough: SQL*Forms 2.1 (Oracle 5.1, 1987)

Excursus 3: Oracle Customer Relations 1/3

- 1980s: IT was Mainframe dominated, User Dissatisfaction
- Oracle Marketing:
  - if your IT Dept. cannot meet your Needs => buy some Mini-Computers (VAX) of your Choice + some Terminals
  - Oracle will run everywhere
  - Applications you can program with our Tools (there was no Standard Software in these Days)
  - if you are not able: Oracle can do it for you!
- begin of decentralization

Excursus 3: Oracle Customer Relations 2/3

- Oracle raised Expectations
- Problems:
  - Support was needed during Installation
  - Scalability problems
    - Version 5: maximal 40 (42?) simultaneous Sessions
    - Forms did not work with Transaction Monitors (Tuxedo)
    - analyzing Performance Problems not possible
    - Solution: bigger Machines, good for License Dept.
    - V6: 120-250 sess., V7: 1250 sess.

Excursus 3: Oracle Customer Relations 3/3

- Unix enters the Datacentre
- typically: dedicated Server per Application
  - transparent for Accounting and Provisioning
  - convenient for Users (nobody disturbs)
  - but: Users had to administrate their own Servers (Patches, Backups, …)
- growing Applications could no longer exist isolated
- beginning from Version 5 the Market was the driving Force concerning the functional Requirements of the RDBMS

Excursus 4: Oracle Platform Support 1/2

- V3-V6: VAX/VMS was development platform
  - process creation: huge CPU overhead
- growing number of supported platforms: Oracle sold "Porting Kits" to hardware vendors (no need to develop "Operating System Dependent Layer")
- large number of OS in data centers (=> market share!)
- many OS were abandoned
- Oracle ported in joint programs with most important hardware/OS vendors

Excursus 4: Oracle Platform Support 2/2

- 1987: Sequent + Oracle 5.1 : first SMP port (see note 1)
- 1990: nCube: MPP support
- Oracle opted also for the mainframe market (IBM dominated)
  - Version 5: VM/CMS + MVS
  - Oracle was not able to run its own business process on mainframe
  - mainframe business was a flop in the long run
- PC market: 1986 Oracle 4.1 "The Cube" (see note 2)
  - breakthrough with NT 3.51 and OS/2

Appendix A: The 12 Codd Rules (*) 1/7

(*) Codd, E.F. "Is Your DBMS Really Relational?" and "Does Your DBMS Run By the Rules?"

ComputerWorld, October 14 1985 and October 21 1985

1. Information Rule

- all data in tables (relations)
- tables consist of rows (tuples), no row ordering
- each row has the same columns
- column: scalar, no redefinition by the table definition, unique in table

2. Guaranteed Access Rule

Each datum has to be accesible without ambiguity (unique, primary keys!).

3. Systematic Treatment of NULL Values (see notes)

Appendix A: The 12 Codd Rules 2/7

4. Dynamic On-Line Catalog

- same access to the structure of the DBMS like data access
- "Dictionary"

-Oracle: Version 6: X$- and V$-Views yield SQL access to control and performance informations

Appendix A: The 12 Codd Rules 3/7

5. Comprehensive data Sublanguage

- SQL (many dialects)

  - declarative, non – procedural, no sequentiality, no loops, no branches
  - query result is also a table (relation)

    select xx.zero_sal_qty

    from (*select count(*)* zero_sal_qty *from emp where sal=0*) xx;

    xx .. Table alias, zero_sal_qty .. column name

    Oracle: View : Query in dictionary as table

Appendix A: The 12 Codd Rules 4/7

6. View Updating Rule

- Each data manipulation valid against a table has also to be valid against a view.

7. High – Level Insert, Update, Delete

- I,U,D has to be possible for any retrievable number of rows

  - update … set ..=.. where …;

8. Physical Data Independence

- SQL Programming possible without knowledge about storage of data.

  - DML: o.k.
  - DDL: not o.k. (mixture of log. and phys. specs) (see notes)

Appendix A: The 12 Codd Rules 5/7

9. Logical Data Independence

Data appearance should not change when table definitions change => views.

Very difficult in full generality (Normalization !).

10. Integrity Independence

- Definition of constraints forces data integrity.

- up to Oracle 7: unique index, NOT NULL constraint (enforced)

  - Application was responsible for data integrity.

- from Oracle 7: foreign key constraints, check constraints

  - Avoidance of user- or application errors.

Appendix A: The 12 Codd Rules 6/7

11. Distribution independence

- DBs may be distributed on different servers.
- Partly obsoleted by n-Tier Architectures and hardware developments.
- Oracle 5: Distributed Database Feature (query only)
- Oracle 7: 2PC
- Rule 2 may be violated when a table is located on different servers.

Appendix A: The 12 Codd Rules 7/7

12. Nonsubversion Rule

- No backdoor in the DB.
- All changes through SQL.
- Partly violated => SQL Loader, external tables

What is missing?

- transactional integrity (commit, rollback) (ACID)
- isolation levels
- privileges

Appendix B: Client / Server Architecture: Oracle 1/4

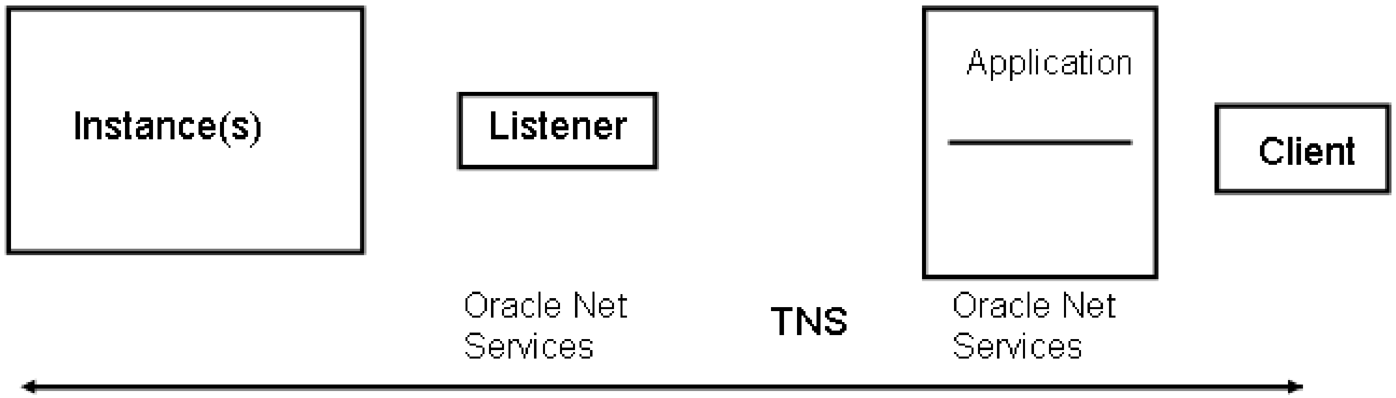

- TNS: Transparent Network Substrate
  - works with any standard network transport protocol
- Client may als be an Instance
- dedicated Server Arch. / Shared Server Arch. (see Note)
- Transparent Gateways for Connection to foreign RDBMS

Appendix B: Client / Server Architecture: Sybase 2/4

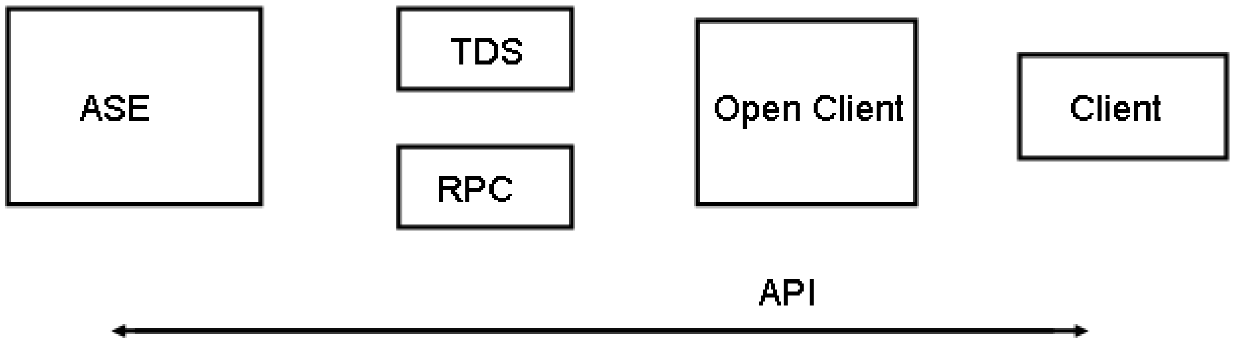

- ASE: no dedicated listener process
- comm. between ASE with RPC
- CIS: Component Integration Services (conn. to foreign RDBMS)
- Open Client: API for conn. to other servers

Appendix B: Client / Server Architecture: SQL Server 3/4

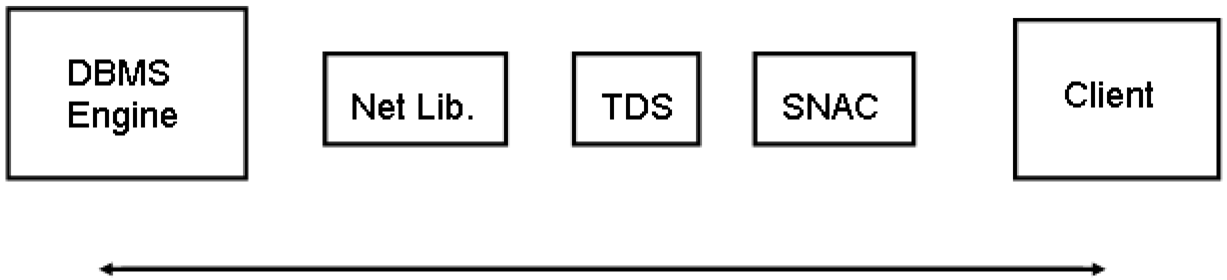

SNAC .. SQL Server Native Client (only 2005) (OLEDB+ODBC)

TDC .. Tabular Data Stream (data transfer protocol)

Appendix B: Client / Server Architecture: DB/2 4/4

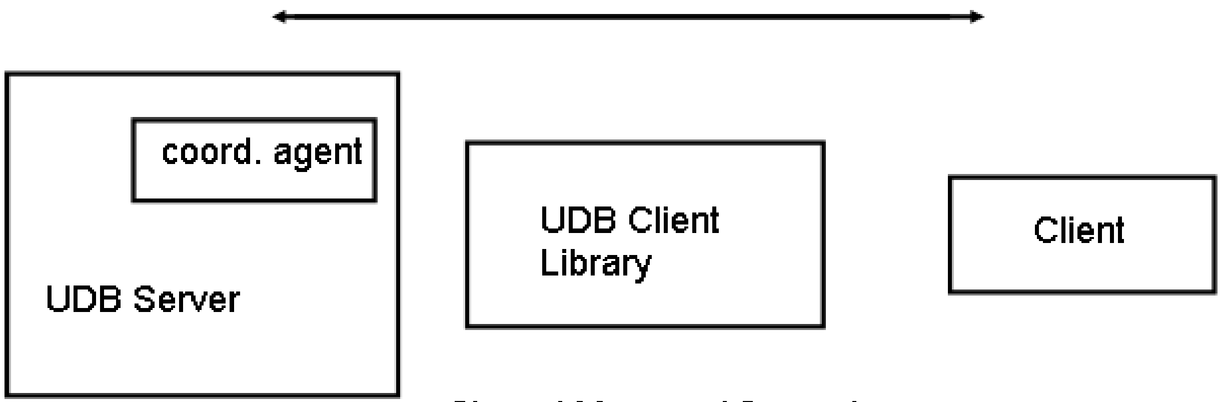

- Shared Memory / Semaphores
- TCP / IP
- Named Pipes
- NetBios
- SNA
- IPX/SPX

App. C Data Encryption / Network Security: Oracle 1/2

- DBMS_CRYPTO Package for selectively encrypting sensitive Data
    - DES
    - AES
    - SHA-1 (see note)
- listener is password protected
- Virtual Private Database: provides database integrated access control

  (access by different user groups to the same data: only relevant data are visible)
- Oracle Label Security: preconfigured and implemented VPD
- Network Traffic: ASO

App. C Data Encryption / Network Security: Oracle 2/2

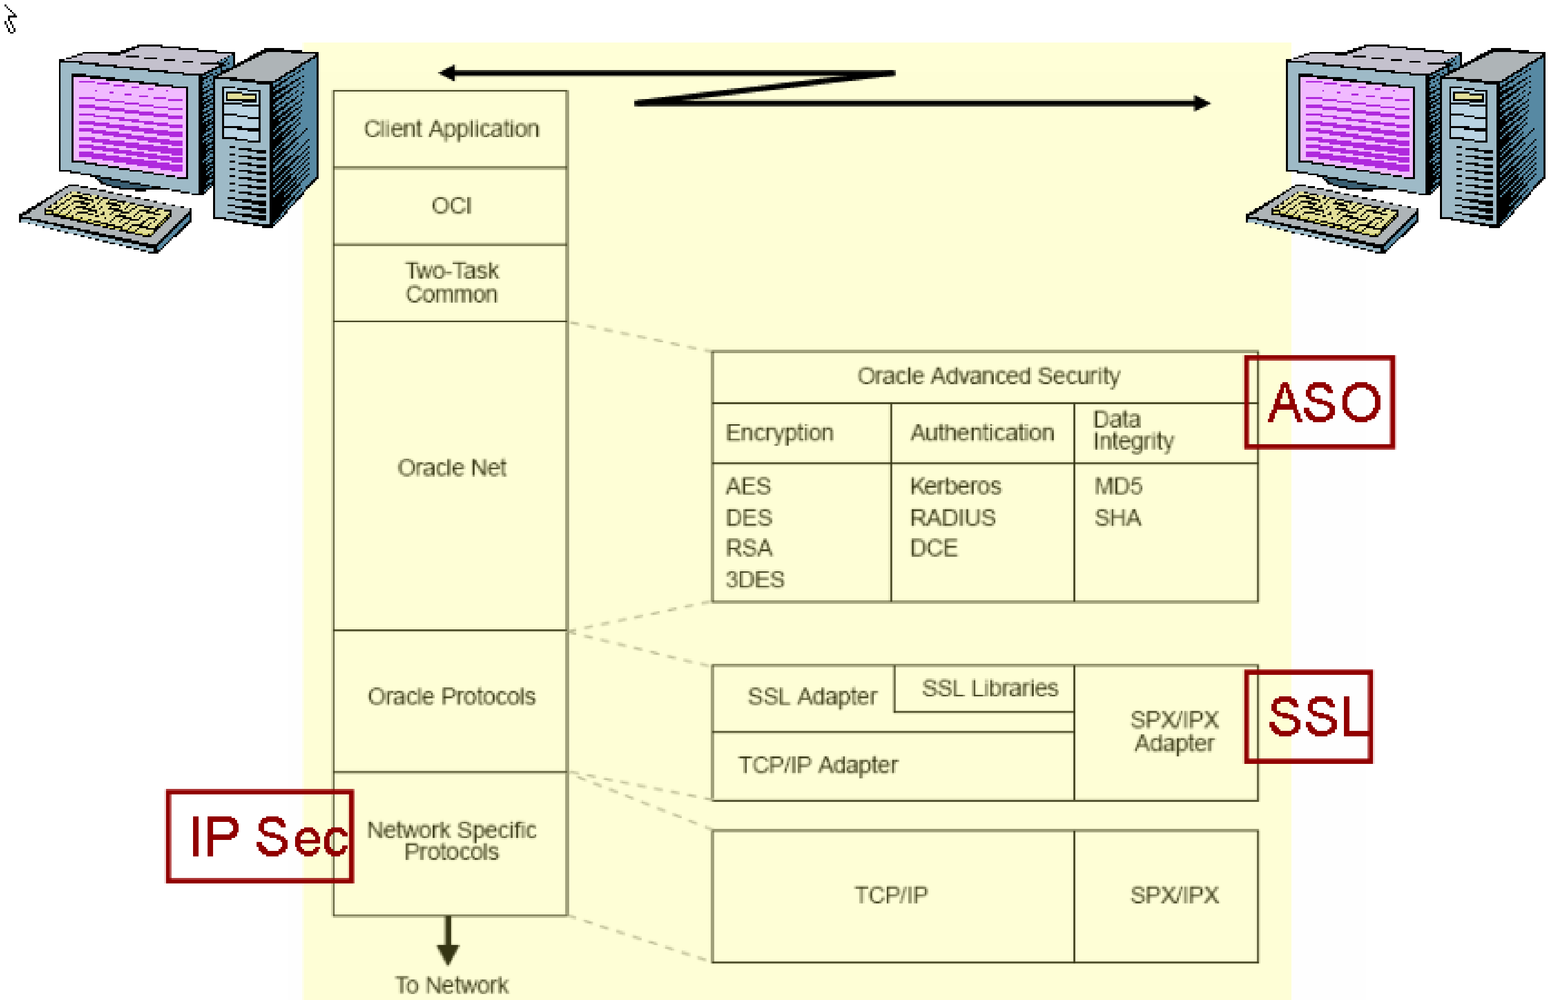

App. C Data Encryption / Network Security: Sybase 1/1

- Database Dump/Load may be password protected

- PKI with SSL is available (SSL Plus Library API from Certicom Corp.)

- ASE 15 Encryption Option:

  - no external keys
  - permission based access control
  - no application modification

App. C Data Encr. / Network Security: SQL – Server 1/1

Hierarchical Encryption and Key Management Infrastructure

Data Protection API (created d. setup)

SQL Server	Instance	Service Master Key
SQL Server	Databases	Database Master Key

(further drilldown possible)

network traffic: SSL SQL native Client Library (40 bit or 128 bit)

App. C Data Encryption / Network Security: DB/2 1/1

- on Mainframe: Data Encryption for IMS and DB/2 uses Crypto Hardware
  (DES, AES) (Table Level)
- UDB: SQL built in Functions (column level)
  Application Modification needed

- Network Security: relies on OS measures for encrypting Network Traffic
  (ssh Port Forwarding)

App. D: High Availability Concepts: Sybase 1/2

normal Operation

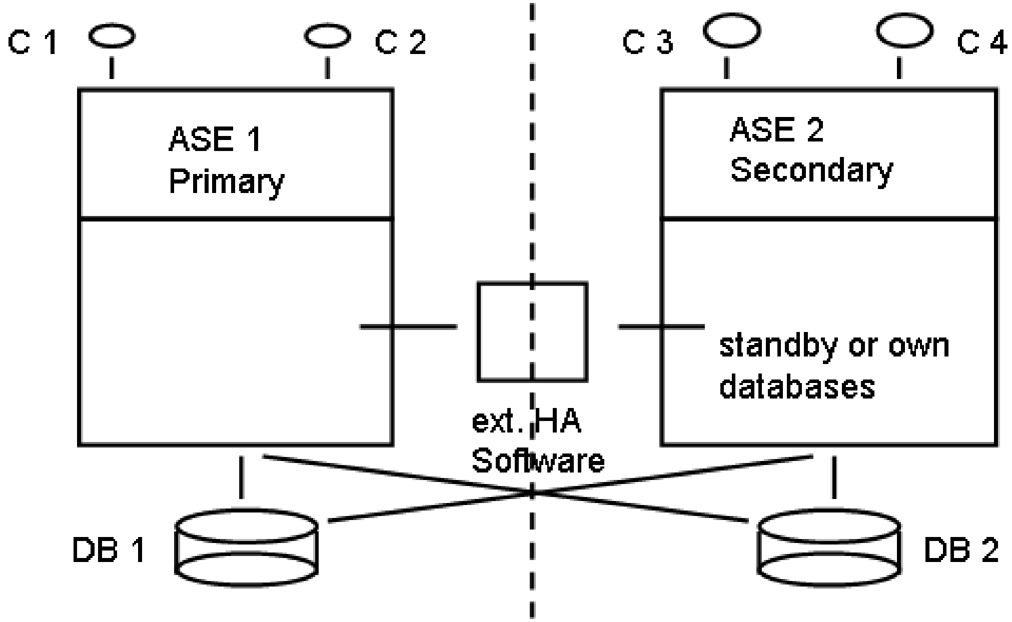

emergency Operation

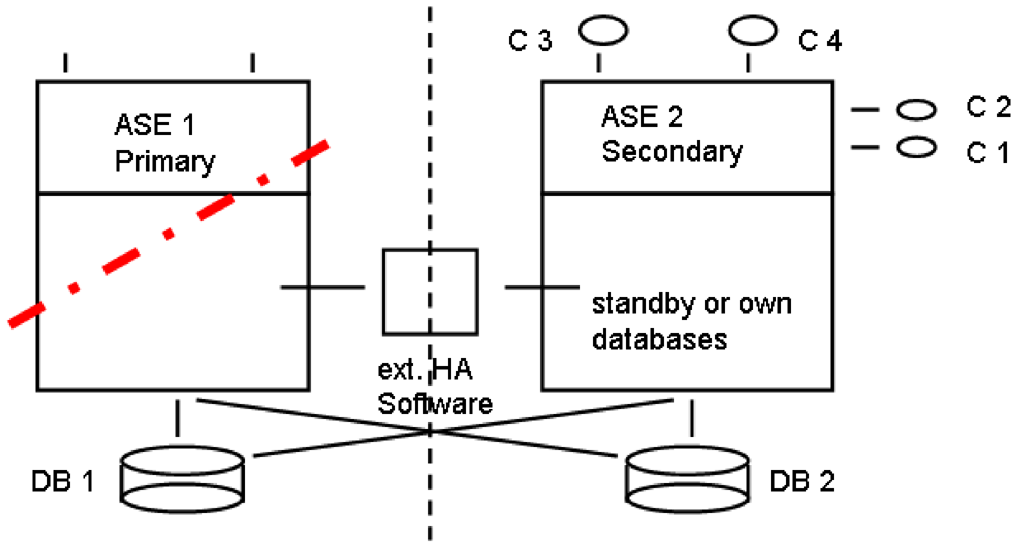

History and Comparison of Relational Database Management Systems

TechnoCircle HVB Information Services

Contents

HVB Information Services
Member of UniCredit Group

1. History of Relational Database Management Systems

2. Architecture Overview of important RDBMS

2.1 A Compilation of Generic Terms

2.2 Timeline of a Transaction

2.3 RDBMS Architecture Comparison

- Oracle
- DB/2
- Sybase
- SQL Server
- MySQL

3. Technology Comparison: Oracle, DB/2, Sybase,SQLServer

3.1 Server Security Model

3.2 Backup/Recovery and Logging

3.3 Data Consistency vs. Concurrency

3.4 High Availability Concepts

4. Vendor Comparison: Beyond Technology

4.1 Vendor Product Portfolios

- Oracle and Sybase

4.2 Platform Support / Market Shares

4.3 Cost Comparison

- SQL Server 2005
- Oracle 10g
- DB/2 8.2

4.4 Total Cost of Administration

- SQL Server vs. Oracle

1.1 Origin and Foundation 1/11

Q.: Why is the database industry dominated by US companies like

- IBM
- Oracle Corp
- Sybase, Inc.
- Teradata Corp. (NCR Corp.)
- Microsoft?

A.: There was in the beginning a strong competition between industry and academic research.

1.1 Origin and Foundation 2/11

- **19th century**: Us Gov. needed reports from large datasets
- **1890**: Herman *Hollerith*: first automatic information processing equipment => punchcard machine => census 1890 / 1900 were processed
- **1911**: IBM founded (*Hollerith* et.al.)
- control of industry production, tax calculation not possible without automatic data processing
- **1935**: Social Security Act: records of 26 mio individuals, special equipment by IBM
- Bureau of Census: bought UNIVAC I
- **1959**: 200 computers in the Pentagon
- **1960s**: remove hardware constraints from programmers, term: "database"

UNIVAC I: first digital computer on commercial market

emerged from a project of the UPenn: Electronix Discrete Variable Automatic Computer (EDVAC)

US dominance in the punch card industry laid the foundation for the dominance in electronic computing in general

database: the information stored on a computer system can structured and manipulated regardless from the specific hardware details

1.1 Origin and Foundation 3/11

- Standardization efforts
  - computing => commercial market
  - techniques for: data access, data quality, security, control
  - **1960**: US DoD: Conference on Data system Languages (Codasyl)

    => Cobol (Common Business Oriented Language)
  - tapes => hard disks (serial => random)
  - **1961**: Charles *Bachmann* (GE): IDS (Integrated Data Store System)
  - **1960**s: Database Task Force Group (Codasyl): *Navigation* with computer languages in databases

HVB Information Services Member of UniCredit Group

1.1 Origin and Foundation 4/11

- **1968**: IBM: IMS (partly developed during the Apollo project)
- **1971**: formal standard: Codasyl approach to database management
  - hierarchical database model
  - network database model
  - => navigational databases (*Bachmann* 1973: The programmer as navigator)
- adoption in mainframe market outside IBM (Eckert-Mauchley Comp. Corp., Honeywell, Inc., Siemens AG, DEC, Prime Comp. Corp.)

IMS: Information Management System:

2006: 100 million transactions in one day on a single sysplex (customer)

3000 days without outage (customer)

95% of Fortune 1000 companies use IMS

200 million end users

15 million GB of data

50 billion transaction / day

Hierarchical and network database model: navigational databases

•data in tree structure

•parent child records (1:N, M:N)

•loops, pointers

•different record types, subrecords, subfields

•User had to navigate within the data sets

1.1 Origin and Foundation 5/11

- *Edgar F. Codd* (IBM) was dissatisfied ("taking the old line view, that the burden of finding information should be placed on users …")
- landmark paper: *"A relational Model of Data for large shared Data Banks" (see note)*
  - Indepence of Data from the Hardware- and Storage Implementation
  - automatic navigation to the data set (high level nonprocedural language for data access (Record => Set)
  - pointers => keys (primary, secondary)
  - theoretical proposal, no practical design or implementation

E. F. Codd: A Relational Model of Data for Large Shared Data Banks. Commun. ACM 13(6): 377-387 (1970)

1.1 Origin and Foundation 6/11

- impact was not fully recognized (even by *Codd*)
- IBM: IMS was preferred (sole strategic) product (invest, infrastructure, success, profit)
- relational model had to survive in and outside IBM
- *Codd* initiated public debate with *Bachmann* (Codasyl)
- 1970s: two projects for development of relational products:
    - Army, Air Force, Navy, NSF, UC Berkeley funded: *Ingres*
    - IBM funded: *System R* (1973)
        - begin of competition!

1.1 Origin and Foundation 7/11

System R: (see notes)

- high level, non navigational data independent user interface
- 1974/75 prototype
- 1978/79 fully functional, multi user version => SQL

Ingres (Interactive Graphics and Retrieval System)

- handling of geographic data
- 1974: Prototype, own query language: QUEL
- several revisions, testing through user community => feedback
- rapid spread because of DEC hardware at universities

System/R Features

•SQL (Chamberlin, Boyce)

•CBO

•Query Compiler

•Ad Hoc Query Formulation (in contrast to navigational databases)

•Online Data Definition

1.1 Origin and Foundation 8/11

Diffusion and Commercialization (Ingres)

- Source free available (1000 copies in circulation)
- *Stonebraker* founded Ingres Corp. (1994 bought by CA)
- *Robert Epstein* (Chief Programmer) => Britton-Lee, Inc. => Sybase, Inc.

  (ideas from Ingres, Public Agencies as first customers)
- *Held/Youseffi* => Tandem Comp., Inc. (predecessor of Nonstop SQL)
- *Held*: up to 1998 Sen. Vice Pres. Engineering, Oracle
- *Hawthorn/Ubell* => Britton-Lee => cofounder of Illustra Inform. Technol.,

  => *Informix*
- …

1.1 Origin and Foundation 9/11

Epstein: "What came from Ingres was the experience of having built a prototype … to say what parts need to be done differently."

- hard competition between Ingres and System R
- System R: 15 people
- Ingres 1973-1979: 30 people (never more than six)

- SQL persisted (would have been possible years earlier)

1.1 Origin and Foundation 10/11

- *Larry Ellison* (IBM): read publications of the System R group => Oracle
  - sold SQL compatible product before IBM
- 1980:
  - IBM developed SQL/DS => mainframe market
  - second generation of relational systems
  - *Codd*: ACM Turing Award
- 1982: End of Ingres Project
- 1985: continues in Postgres (UC Berkeley), object relational

1.1 Origin and Foundation 11/11

Lessons learned

- strong governmental influence at dev. of rel. systems (tasks like census)
- funding speeds up commercialization (competition System R <=> Ingres)
- free publication of research results => technology transfer in community
- scientists from industry laboratories founded their own enterprises

HVB Information Services Member of UniCredit Group

1.2 History of Oracle (*) 1/16

1975-1979 1st Rel.: Oracle 2	- Market Demand: RDBMS - Larry Ellison (Marketing) , Bob Miner, Ed Oates (Code): Software Development Laboratories (SDL) => Implementation of a RDBMS => Oracle 2 (see notes) - less than one year development time
1979 1981 VAX Factor 1983	SDL => RSI (Relational Software, Inc.) market pressure: fast avail., cheap product, - PDP-11 (widely used VAX platform), subqueries, joins, no transact. - Oracle 2.3 : complete RDBMS -RSI => Oracle Corp.

- Oracle Insights: Tales of the Oak Table, Springer, Apress, 2004
- Version 1 would not sell
- Oracle was the Codename of a CIA Project (Ellison, Miner)

1.2 History of Oracle 2/16

1983 Oracle 3 Portability	- Rewrite in C (White Smith C) - many basic design decisions (influence up to date) (see excurs 1) - not very stable !
1985	- hard competition in database industry, begin of tool development (see excursus 2) - Oracle Marketing (see excursus 3) - many systems only remakes of old hierarchical systems - Codd: 12 Rules for RDBMS (Appendix A)

1.2 History of Oracle 3/16

1984 Version 4	- rewrite in K&R C - Portability on almost any platform (flex. FS and IPC facility) (Note) (see Excursus 4) - Views, Named Accounts - Read Consistency
1985/86 Vers. 5.1	- encrypted passwords - subqueries - precompiler (Cobol, Fortran IV, PL/1) ■ SQL embedding in 3GL programs

Hardware had to provide atomic "Test+Set"-Operation (no integrity violations allowed)

Oracle Kernel has to check, if memory block can be used exclusively by 1 process (check if used and occupy in one atomic operation, which cannot be interrupted)

1.2 History of Oracle 4/16

	■bad quality until Version 8 ■OCI was the choice (Library Calls for exec SQL in 3GL prg.) - SQL*Net: ■C/S Arch. ■system, platform independence ■DB Server also Client (Database Link) ■Query Only Distrib. Database Support (see note) ■Gateways for other DB engines

- Version 7:
  - 2PC => Distributed Updating
  - synchronous and asynchronous Replication
- Version 8: Advanced Queueing Option

1.2 History of Oracle 5/16

	■ growing spread of PCs => mismatch ■Server: DB + Appl., PC: only Terminal + Interface ■Two Task Arch. IPC => TCP/IP was possible (see note) ■ better networks => more effic. implementation of SQL*Net ■n-Tier architectures were possible - first tries with clusters
1987	- standardized applications were needed - Founding of Applications Division at Oracle (Business Mgmt. Software)

HVB Information Services Member of UniCredit Group

at first the Error was made: bring the application on the PC (10ms (10 Mbit Ethernet), 0.5s (X.25) Roundtriptime)

Two Task:

•separation of database and application on different servers

•implicite assunption of instantaneous connection between Server and Client

•architecture could not get rapid spread because of network performance

1.2 History of Oracle 6/16

late 1980s Problems: Perf./ Scaling	- Oracle: good for rapid dev . of OLTP Applications - no scaling, few sessions - no precompiled queries, Parsing Overhead - RBO: good for OLTP on well normalized schemas bad for DSS - growing number of users aggravated performance problems
1988 - 1991 Version 6	- true row level locking (additional charge) (see notes) - Online Backups - Shared Pool (minimization of parsing overhead)

Matthias Mann / HIS14DB, 28.06.2006
Seite 23

HVB Information Services Member of UniCredit Group

Row Level Locking:
no entries in lock pool
lock management in blocks => good for scaling

Competitors: Block-Level and Table Level Locking with Lock escalation

1.2 History of Oracle 7/16

	- change from block level logging to field level logging - keep rollback segments in the database ■ recovery of undo structures via redologs (see notes) ■ Read Consistency - huge memory consumption when many clients => MTS (see notes) - no Upgrade Path from Oracle 5 to Oracle 6

Rollback Segments:

RBS are also database blocks as the rest of the database => changes to the redologs

therefore rollback of uncommitted changes during crash recovery possible

memory problems when many client sessions were active => MTS (100 sessions => 1GB memory !)

(at the same time: Sybase at US fin. services: 16 kB per session)

MTS:

instable at the beginning (thread blocking)

on VMS: because of overhead during process creation (dedicated server, many clients) MTS was good

MTS did not yield the desired memory saving effect

1.2 History of Oracle 8/16

1992 Version 7	- Oracle had to rescue non-tech educated users from their errors - data integrity: declarative constraints (up to now only unique keys and NULL-value checking) - PL/SQL storage in database (procedural language + SQL) - triggers: implement data rules (insert into a => update b) - stored procedures: implement process rules - =>radical changes in application design - emergence of standard software (database independent design) - role concept

HVB Information Services
Member of UniCredit Group

1.2 History of Oracle 9/16

	- Cost Based Optimizer: not perfect - Rule Based Optimizer since Version 6 no longer improved
1993	- Business Applications in Client/Server Mode
1994	- Port to Windows
	- Distributed Transactions (2PC)
1995	- Parallel Query => Data WareHouse
1996	- porting of all Tools to Windows
	- graphical EM: not usable (see note)

even Oracle staff used internally 3rd party management tools

1.2 History of Oracle 10/16

1997 Version 8	- LOBs (storage of unstructured data) - Partitions (ease of administration) - Materialized Views (Benchmark Opportunism) (see note) - Function Based Indexes (error tolerance) - Bitmapped Indexes: Performance Boost for DWH - Oracle Discoverer: Ad Hoc End User Query Tool for BI - Support for object oriented Developing, new datatypes - completely revised development suite (native Java): ■ Jdeveloper - Oracle Application Server 4.0 (middle tier)

View: storage of a query

Materialized View: View execution and storage of results

1.2 History of Oracle 11/16

1998 Version 8i	- native Java Runtime Environment - SQLJ (SQL in Java Code) - interMedia: managing of Multimedia Content - Repository for Designer (Modeling and Application Generation Tool) - move away from C/S Paradigma: Oracle Applications run in browser - Port of RDBMS and Application Server to Linux - RAW Iron Initiative: Flop (Pack. of Server+cust. OS+Oracle Env.) - XML Support

1.2 History of Oracle 12/16

1999	- Oracle WebDB (manage DB driven Websites) => Oracle Portal
	- Applications 11 => Applications 11i
2000	- Internet Filesystem (iFS) (see note)
	- Oracle AS9i Portal Technology (Content Mgmt, Portlets, Deployment Options)
	- OAS9i Wireless
	- Webcache Technology (Part of OAS)
	- Preview Cache Fusion (RAC)

iFS:

Drag and Drop of Windows Files of arbitrary Format in a directory structure => move transparent into the DB

Versioning

Check In, Check Out

Fulltext Search

APIs for FS (FTP, Mail, NFS, …)

1.2 History of Oracle 13/16

2001 Version 9i	- DEC could not enter the mainframe market - research in ■ SMP- architectures (cache coherency) ■ cluster architectures (Shared Disks) also with Oracle Shared Instances (Instance:DB = N:1) (see note) - Real Application Cluster (Oracle Cluster Technology)

Oracle Cluster

Oracle 5: "Clustered Oracle"

Oracle 6.2 - Oracle 8: Oracle Parallel Server (OPS), scaling problems, problems with synchronizing memory contents of the nodes

Oracle 9i: RAC: breakthrough in performance and scaling

1.2 History of Oracle 14/16

2001	- native XML Support - Jdeveloper: J2EE- and XML-Support - Table Compression (Block Level) (Index Compr. in Vers. 5 and 8) - Flashback Query - OEM becomes usable (scalable architecture, repository)
2004 Oracle 10g	-development was always customer driven, but now: -enormous improvements in manageability and self-tuning of the database (see note) - Automatic Workload Repository (AWR) - Automatic Database Diagnostic Monitor (ADDM) - Automatic Space Managament (ASM)

weakness of Oracle was ever:

- DBA had to invest time and skill (see later)

1.2 History of Oracle 15/16

	- Selftuning must be licensed! - Selftuning is valuable for small businesses without dedicated DBAs - Oracle can afford not to satisfy small customers at moderate prices and let big customers pay extra charges - Product Quality: ■ new features often half hearted and error prone (see note) ■ initial release of each version is unusable in the real world ■ this is in 10g different

Errors in Oracle:

Version 6: Oracle could not run ist own business on Oracle => tkprof, hidden init.ora - Parameters

1.2 History of Oracle - Summary 16/16

1979	Oracle 2	first commercially available RDBMS
1983	Oracle 3	Portability
1984	Oracle 4	Transactional Integrity, Read Consistency
1986	Oracle 5	Client/Server
1988	Oracle 6	Scalability
1992	Oracle 7	Data Integrity
1997	Oracle 8	Partitioning
2001	Oracle 9	Cluster technology comes of age
2004	Oracle 10	Manageability, Selftuning

1.3 History of Sybase 1/13

1984 1986	- Robert Epstein (UCB), Mark Hoffmann (Univ. of Arizona) => Robert Epstein Consulting => Sybase (System+Database) - Sun: major Partner - SQL Server V1.0 (Beta) and DB Library on SunOS: ■ Client/Server Database for OLTP (Appl. Logic in Server) ■ pioneer with Stored Procedures, Triggers
1987	- SQL Server (Beta) on VMS - 100 employees, 150 customers

1.3 History of Sybase 2/13

1988	- second major Partner: Microsoft - SQL Server for OS/2 (joint dev. of Sybase, Microsoft, Aston-Tate) - Version 3.0: Stored Procedures, Triggers, CBO ■ first market Version - Sybase: Unix/VMS, Microsoft: Windows, OS/2
1989	- Open Client / Open Server: API for Comm. between Clients and different Data Sources
1990	- Sybase SQL Server 4.0: text, image datatypes - Vision of "Total Solution C/S Company" : aquis. of SQL Solutions - Integration of IBM MVS mainframes in "C/S World"

HVB Information Services Member of UniCredit Group

1.3 History of Sybase 3/13

1991	- Version 4.2: international language support - enters Chinese Market - Powersoft introduces Powerbuilder 1.0
1992	- Version 4.8: SMP support (engines), union, new datatypes - Version 4.9: multibyte Character Sets - Version 4.9.1: extraordinary stable - Version 4.9.2: log-based Replication

1.3 History of Sybase 4/13

1993 System X	-Backup Server, Monitor Server, - RI Constraints, Cursors, Identity Columns, system roles - thresholds, Auditing, encrypt. Password Storing, sybsystemprocs DB - CT Library replaces DB Library (C/S API) ■ but: severe Installation- and Stabilityproblems - Replication Server (Version 10): Open Replication Technology allows replication regardless of Database - Vertical Industry Solutions Group ■ Healthcare, Financial Services, Oil & Gas, Telecom

HVB Information Services Member of UniCredit Group

1.3 History of Sybase 5/13

1994	- Microsoft quits Partnership (reseller for OS/2 and Win NT) ■ gets copy of Sourcecode Version 4.2 Sybase SQL Server ■ Microsoft SQL Server now an independent Product - both comp.: SQL Server and Transact SQL (T-SQL) (Trademark still today owned by Sybase)
1995	- Merger with Powersoft (leading supplier of developm. tools for 4GL Applications) - with Powersoft came Watson: small RDBMS and Enterprise Synchronization Solution on sev. OS for mobile and embedded products ■ SQL Anywhere Studio => Adaptive Server Anwhere ■ undisputed mobile DB market leader

1.3 History of Sybase 6/13

	- Version 11.0: ■ Improvement in Quality ■ Table Partitioning ■ Memory Partitioning (named Caches) ■ ANSI Transaction Isol. Level 0 (Dirty Reads) ■ sp_sysmon, online database, Server config File
1996	- dominant provider in mobile DBMS Market ■ ASA: market share 22% (2000: 68%) - Strategy: Focus on OLTP, DataWarehouses, Internet Applications

1.3 History of Sybase 7/13

1997	- Version 11.5: ■ Sybase SQL Server renamed to Adaptive Server Enterprise ■ case expression ■ CIS (Proxy tables) ■ XP Server, Historical Server ■ dbcc checkstorage, user defined Roles ■ Query Parallelism, Resource Limits ■ log. Process Manager - Jaguar: 1999 EA-Server (Web Appl. Server for >30 DB)

1.3 History of Sybase 8/13

1998	- Version 11.9 "EARL" (Eagerly Awaited Row Level Locking) ■ DOL on row level / page level (see note) ■ lock timeouts , lock table statement ■ standby_access database, reorg command ■ ANSI Isol. Level 2 ■ license monitor ■ Optimizer Statistics Redesign - 11.0.3.3 : Port to Linux

Before:

Table Locking => argument was: concurrency problems should be treated within the application => not real life

1.3 History of Sybase 9/13

1999	- ASA for Handhelds - Aqu. Home Financial networks, Comb. with Financ. Server Unit ■ => own subsidiary: Finacial Fusion (leader in fin. Services) ■ 2000: 68% Wall Street comp., 95% Fortune 100 with Sybase ■ market share in Financial Sector: 58% - Aqu. of Data Warehose Network => Sybase Industry Warehose Studio - ASE Version 12.0: ■ exec immediate, identity_gap, license manager ■ abstract query plans, alter table, quiesce database ■ optional licensable features: Java in ASE, Adv. Sec., DTM, HA

HVB Information Services Member of UniCredit Group

1.3 History of Sybase 10/13

2001	- ASE 12.5 ■ union views, login trigger, dynamic server reconfiguration ■ larger page sizes, XQL-Parser ■ ddlgen ■ new licensable Options: ■ LDAP, Filesystem access with Proxy Tables ■ builtin Jave Enterprise Beans Server - Customer Base: 90% of Investment Banks, 60% dep. banks, NYSE

1.3 History of Sybase 11/13

2002	- Version 12.5.0.1: ■ ASE Replicator (lightweight Data Replication) ■ sybmigrate ■ free Developer Edition - Version 12.5.0.2: Port on MacOS
2003	- Version 12.5.0.3: ■ MDA monitoring Tables ■ multiple Temp DBs ■ statistics sampling - Version 12.5.1: job scheduler, dynamic data caches

1.3 History of Sybase 12/13

	Version 12.5.3: cross platform transfer with dump/load
2005 ASE 15	- revised partitioning support -computed columns, function indexes - Adv. System Metrics, auto update Statistics - native conn. to MQSeries - completely rewritten CBO - Unicode Data Support - native Encryption in the Database - selective load

1.3 History of Sybase – Summary 13/13

1984	Foundation
1998	SQL Server: 1sc C/S Database for distr. Appl. and mobile Dev. Mgmt.
1990	Open Replication Technology – Database independent
1991	early activities in China: today market leader
1995	IT Infrastructure Standard in Wall Street => Dominance Fin.Service Market
2003	Sybase IQ: worldwide largest DWH at Nielson Media
2004	Sybase iAnywhere: Marketleader in Mobile Devices

1.4 History of SQL Server – Milestones 1/2

1987	Sybase launches SQL Server for Unix
1988	Microsoft, Sybase and Aston-Tate port SQL Server to OS/2
1989	SQL Server 1.0 for OS/2
1990	SQL Server 1.1 supports Win 3.0 Clients Aston – Tate quits Development of SQWL Server
1991	Microsoft and IBM quit joint Deveopment of OS/2
1992	Microsoft SQL Server 4.2 for 16-Bit OS/2 1.3 is released
1992	Port of SQL Server to Win NT (MS and Sybase)

1.4 History of SQL Server – Milestones 2/2

1993	Win NT 3.1 is released
1993	Microsoft and Sybase release SQL Server 4.2 for Windows NT
1994	Quit of joint Development of SQL Server by sybase and Microsoft. Sybase develops Unix variant, Microsoft Windows variant separately.
1998	SQL Server 7.0
2000	SQL Server 2000
2005	SQL Server 2005

1.5 Overview DB/2 History: Foundations 1/2

System/R Project	- IBM Implementation of the relational Model
Project Aris	- Row Level Locking
Rstar Project	- Rel. Model for Distributed Environments
Starburst Project	- Optimization Strategies in the Relational Model
Garlic Project	- Data Management in diverse Systems

1.5 Overview DB/2 History: Milestones 2/2

1980	Database Integration in System/38 (first Implementation of System / R)
1982	SQL/DS (Mainframe VM, VSE)
1983	DB/2 (Database 2 on VMS)
1987	Database Manager in OS/2 (first on distr. Systems)
1988	SQL/400 for AS/400
1993	DB/2 for AIX
1994	DB/2 for HP-UX, Solaris
1995	DB/2 for Windows
1999	DB/2 for Linux

1.6 Landmarks and Cross Relations 1/3

1970	Codd: concept RDBMS			
1976	System/R: RDBMS prototype			
1977		Larry Ellison: Oracle		
1979		Oracle Corp. 1st comm. Database (SQL)		
1983	DB/2 for MVS			

1.6 Landmarks and Cross Relations 2/3

1988		Oracle 6	Sybase – Microsoft Partnership	
1989				MS SQL 1.0
1993	DB/2 for AIX		End of Partnership	
1995	DB/2 for Windows		ASE 11.0 Merger Sybase Powersoft	MS SQL 6.0
1996			ASE 11.5	MS SQL 6.5
1998		Oracle 8	ASE 11.9.2	MS SQL 7.0

HVB Information Services Member of UniCredit Group

1.6 Landmarks and Cross Relations 3/3

2000	DB/2 V7	Oracle 8i	ASE 12.0	MS SQL 2000
2001		Oracle 9i		
2003			Sybase IQ	
2004		Oracle 10g		
2005			ASE 15	MS SQL 2005
			Replication Server 15	

2. Architecture Overview of important RDBMS

2.1 A Compilation of Generic Terms

2.2 Timeline of a Transaction

2.3 RDBMS Architecture Comparison

- Oracle
- DB/2
- Sybase
- SQL Server
- MySQL

2.1 Generic Terms: Logical Components 1/10

Generic	*DB/2*	*Informix*	*MySQL*	*Oracle*	*Sybase*	*MSSQL*
Background Processes and Memory Structures						
Instance	Instance	Instance	Server	Instance	Server	Server
Persistent Objects on Storage (Figure A)						
Database	Database	Database	Database	Database	Database	Database
Relation (Grouping of related information, Figure B)						
Table	Table	Table	Table	Table	Table	Table
special purpose table for quicker lookups						
Index	Index	Index	Index	Index	Index	Index

Gen. Terms: 2/10

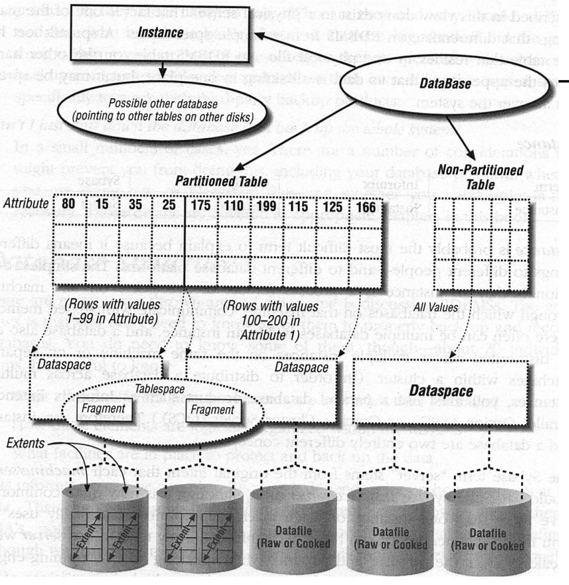

Figure A

Graphical Representation

of a Database

W. Curtis Preston

Unix Backup and Recovery

O'Reilly

2.1 Generic Terms - Logical Components 3/10

Generic	*DB/2*	*Informix*	*MySQL*	*Oracle*	*Sybase*	*MSSQL*
Data which do not fit in a "normal" table						
BLOB space		**(smart) BLOBspace**		BLOB, CLOB	image data type	varchar (max) varbinary (max)
any type of table						
Object	Object	Object	Object	Object	Object	Object
Collection of related Attributes (Figure B)						
Row = Tuple	Row	Row	Row	Row	Row	Row

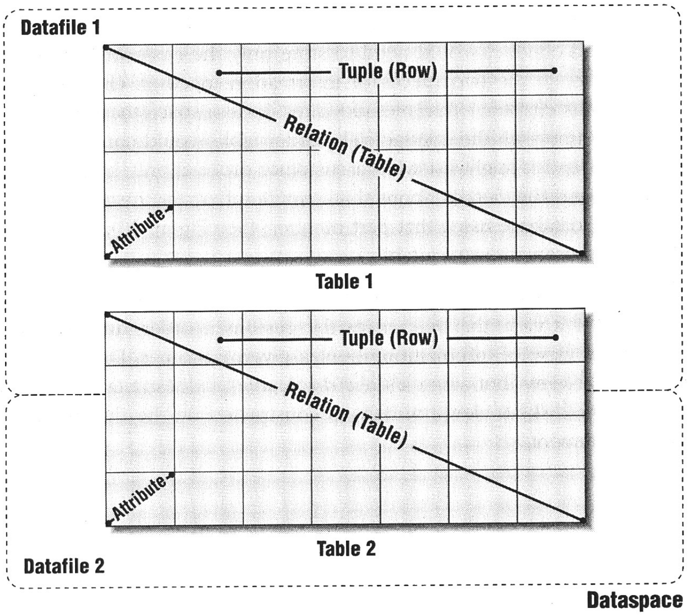

Gen.Terms 4/10

Figure B

Table Layout

HVB Information Services

Member of UniCredit Group

W. Curtis Preston

Unix Backup and Recovery

O'Reilly

2.1 Generic Terms: Logical Components 5/10

Generic	*DB/2*	*Informix*	*MySQL*	*Oracle*	*Sybase*	*MSSQL*
Basic Element of Data within a Table (Figure B)						
Attribute	Attribute	Attribute	Attribute	Attribute	Attribute	Attribute
Activity, which changes one or more Attributes						
Transaction	Transaction	Transaction	Transaction	Transaction	Transaction	Transaction
Point in Time where everything is on Disk						
Checkpoint	Checkpoint	Checkpoint	Checkpoint	Checkpoint	Checkpoint	Checkpoint

2.1 Generic Terms: Physical Components 6/10

Generic	*DB/2*	*Informix*	*MySQL (Inno)*	*Oracle*	*Sybase*	*MSSQL*
Basic Building Block						
Page/Block	Page	Page		Block	Page	Page
Structure for storing data (raw or cooked files)						
Datafile	Container	Chunk	Datafile	Datafile	Device	Datafile/ Filegroup
logically contiguous number of pages (1)						
Extent	Extent	Extent (2)	N/A	Extent	Extent	Extent (3)

(1) may or may not physically contiguous

(2) also physically contiguous

(3) 8 physically contiguous pages

2.1 Generic Terms: Physical Components 7/10

Generic	*DB/2*	*Informix*	*MySQL*	*Oracle*	*Sybase*	*MSSQL*
collection of logically contiguous extents (1)						
Fragment	N/A	Tblspace	N/A	N/A	Allocation Unit (2)	Allocation Unit
Space occupied by a single Table (Figure C)						
Tablespace	Object	Tblspace	**Tablespace**	Segment	Disk fragment	N/A
structures where tables go in (4, Figure C)						
Dataspace	**Tablespace (3)**	**Dbspace**		**Tablespace**	**Segment**	**N/A**

(1) may or may not be physically contiguous

(2) smallest possible storage unit which can be added to a database (32 extents)

(3) database managed / system managed / DMS raw

(4) may consist of several datafiles, may consist of severasl tablespaces

Terms: 8/10

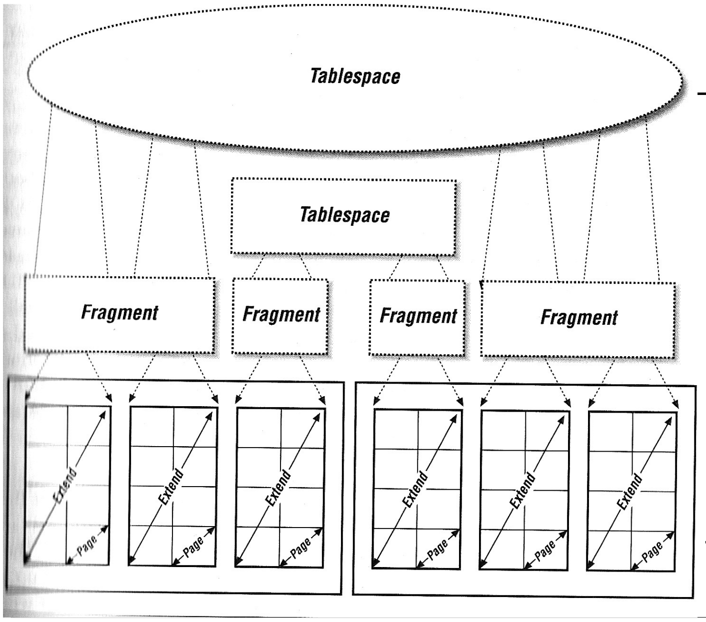

Figure C

Tablespace Layout

2.1 Generic Terms: Physical Components 9/10

Generic	*DB/2*	*Informix*	*MySQL*	*Oracle*	*Sybase*	*MSSQL*
structure for distributing a table across multiple dataspaces						
Partition	Partition	Fragment	N/A	Partition	Partition	Partition
Inventory of all database parts						
Master Database	control files	sysmaster onconfig file rootdb	Information Schema (1)	control file	master database	master database
storage place of the "before image" of a transaction						
Rollback Log	Logfiles	Physical Log	Undo (2) Segments	Undo Segments	Transaction Log	Transaction Log

(1) MyIsam Files

(2) InnoDB TS

2.1 Generic Terms: Physical Components 10/10

Generic	*DB/2*	*Informix*	*MySQL*	*Oracle*	*Sybase*	*MSSQL*
log place for transaction and page changes						
Transaction Log	Logfiles	Logical Log	Binary Log	Redolog	**Transaction Log**	**Transaction Log**

2.2 General Transaction Properties - ACID -

Gliederung durch Klicken hinzufügen

Atomicity	either all of a transaction happens or nothing
Consistency	Database is brought from one consistent state to another
Isolation	The effects of a transaction may not be visible to other transactions until the transaction is committed.
Durability	After commit the transaction is permanent.

2.2 General Transaction Properties - Timeline -

Gliederung dur

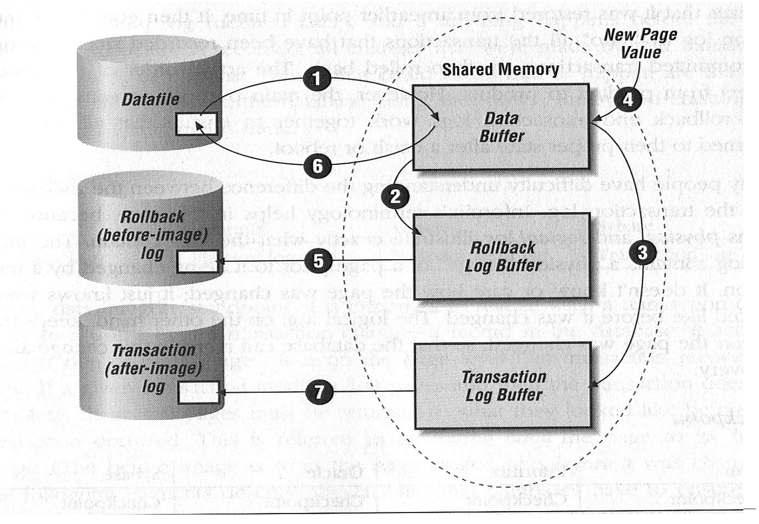

What happens when a user enters a transaction?

3. Page to be modified is loaded in the data buffer.
4. The "before image" of the page is copied to the Rollback Log Buffer.
5. Writing of the "after image" of the page in the transaction log buffer.
6. changes are applied in the data buffer

Up to now nothing is on disk. In the event of a system crash the transaction will not persist.

5. rollback log buffer is flushed to the rollback log (safety of the before image)
6. flush changed data page to disk (datafile)
7. write the after image to the transaction log (for redo purpose)

Now the transaction will be committed in the transaction log.

2.3 RDBMS Architecture Comparison - Oracle -

Gliederung

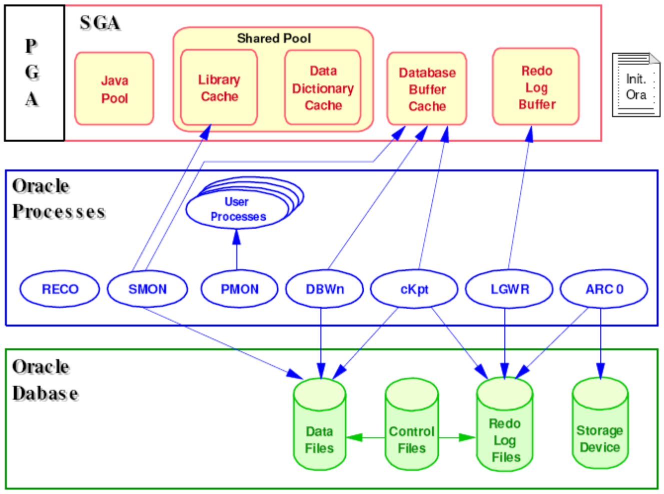

Matthias Mann / HIS14DB, 28.06.2006
Seite 67

Process Model:

Unix. multiple processes

Windows: multiple Threads

Processes:

PMON: Process Monitor (cleanup)

SMON: System Monitor (crash recovery)

DBWR: Database Writer (blocks -> disk)

LGWR: Logwriter (redolog buffer -> redolog)

CKPT: Checkpointing, updates Block Headers

ARCH: copies Redologs to archive logs

2.3 RDBMS Architecture Comparison - Oracle -

Gli

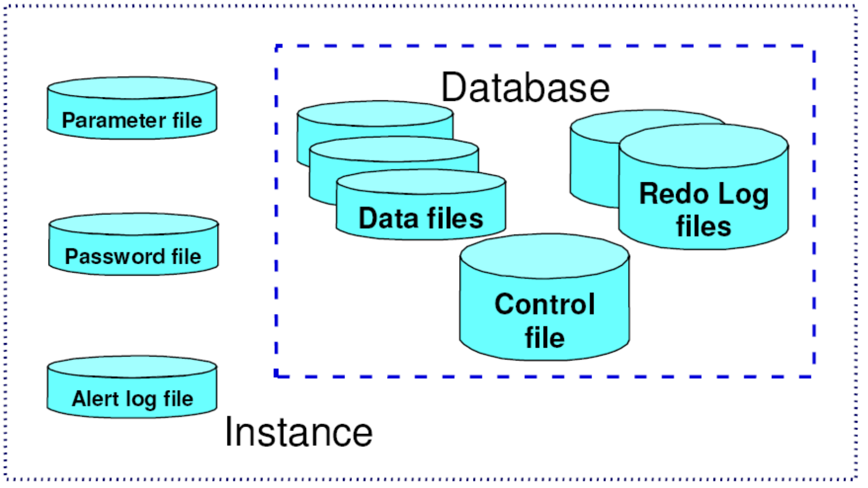

Datafiles: Datafiles and Tempfiles

2.3 RDBMS Architecture Comparison - DB2/UDB -

Gliederung (

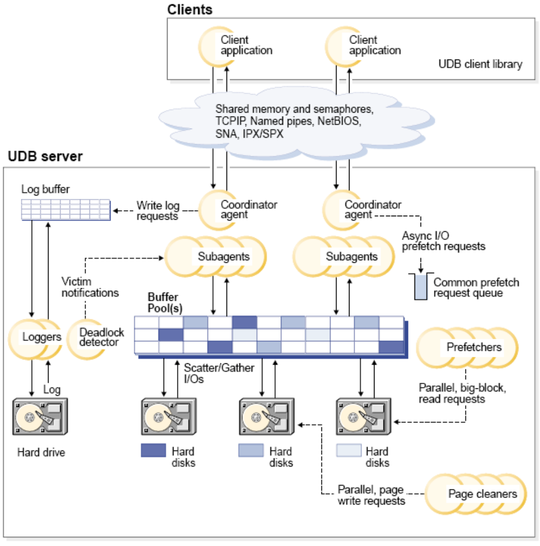

db2wdog: watchdog, cleanup

db2sysc: system controller (startup/Shutdown)

db2pclnr: page cleaner (dirty pages -> disk)

db2pfchr: Prefetcher (async data -> buffer pools)

db2loggr: Log writer (log buffer -> files)

db2dlock: deadlock detector

db2agent: coordination agent

db2agntp: Subagent (works on partition of plan)

2.3 RDBMS Architecture Comparison - DB2/UDB -

Gliederun

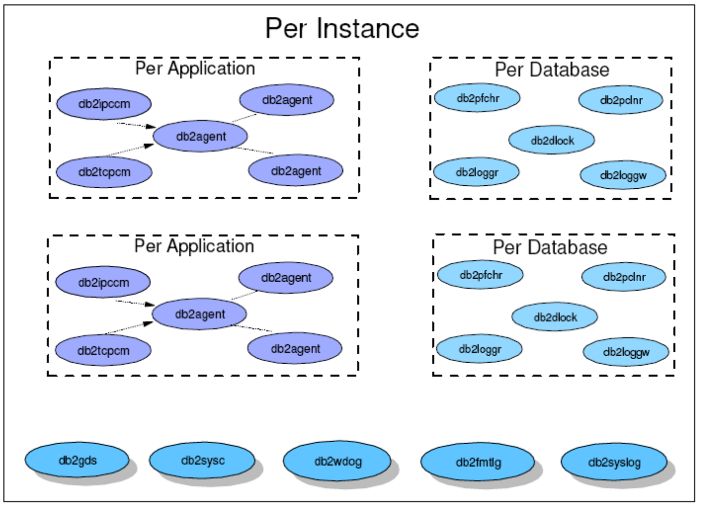

HVB Information Services Member of UniCredit Group

2.3 RDBMS Architecture Comparison - DB2/UDB -

Gliederung durch K

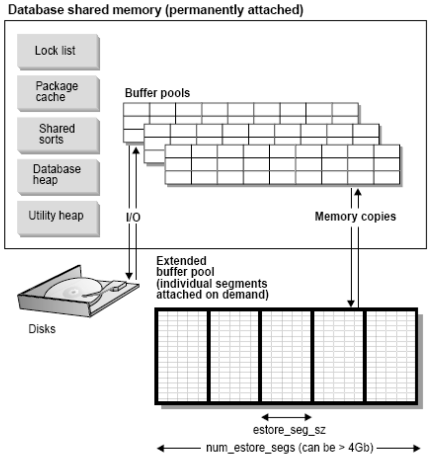

HVB Information Services Member of UniCredit Group

2.3 RDBMS Architecture Comparison - DB2/UDB -

Gliederung durc

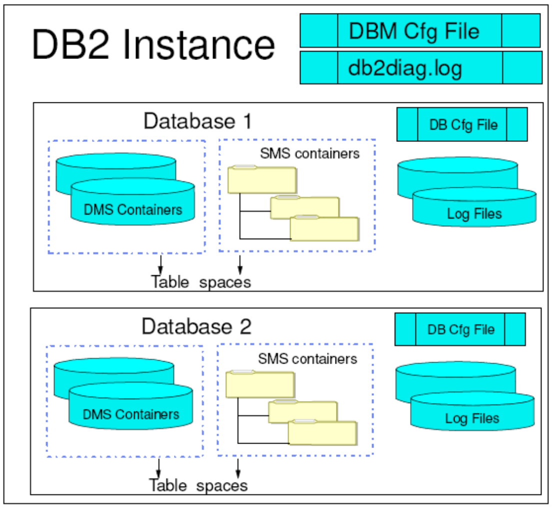

HVB Information Services
Member of UniCredit Group

Categories of Tablespaces:
SMS – system managed
DMS – Database managed

Types of tablespaces:
-regular
-- temporary
-- long DMS

Space Allocation Methods:
-Directory – SMS
-File, Device - DMS

2.3 RDBMS Architecture Comparison - Sybase ASE -

Glie

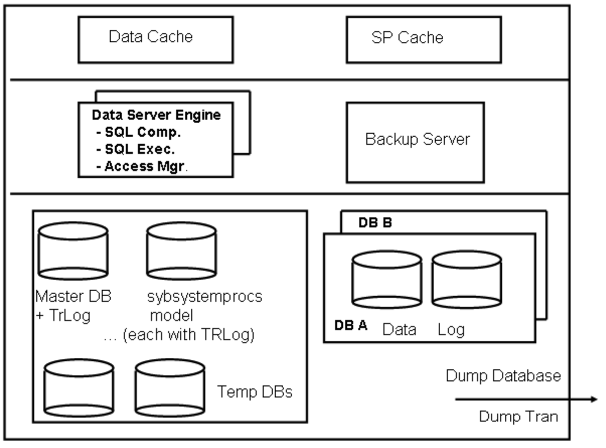

HVB Information Services
Member of UniCredit Group

2.3 RDBMS Architecture Comparison - SQL Server -

Gliederur

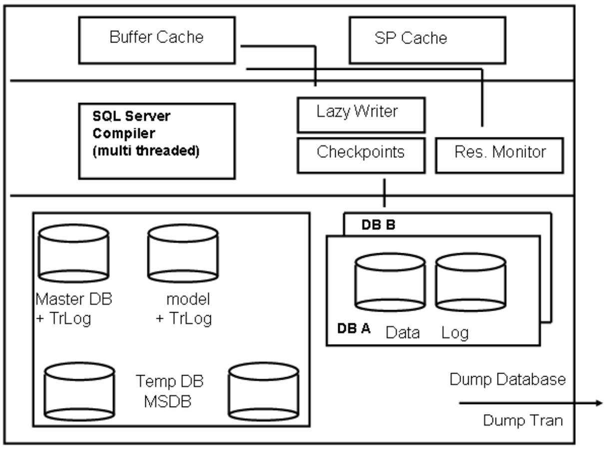

HVB Information Services Member of UniCredit Group

2.3 RDBMS Architecture Comparison - MySQL -

G

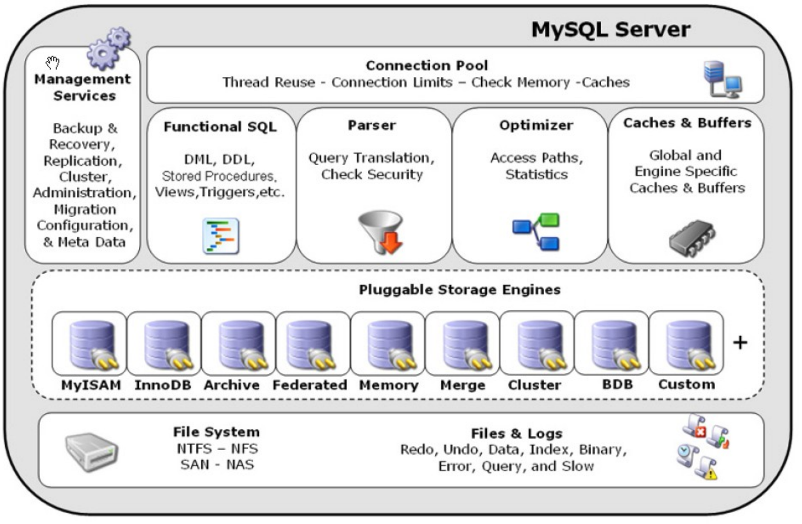

HVB Information Services Member of UniCredit Group

myISAM – File for Dictionary always (Information Schema)

InnoDB:

Version <= 4.0 one datafile

Version >= 4.1 one datafile or one datafile per table

2.3 RDBMS Architecture Comparison - MySQL -

The MySQL Pluggable Storage Engines

- MyISAM
  - Table Locking
  - no Transactions
- InnoDB
  - ACID compatible
  - row level locking
  - Tablespaces

2.3 RDBMS Architecture Comparison - MySQL -

- default Isolation Level: Read Committed
- multi version concurrency control (readers do not block writers and vice versa)

- Archive
  - compressed data -> ILM

- Federated
  - links to tables on another server

ILM: Information Lifecycle Management

3. Technology Comparison: Oracle, DB/2,Sybase,SQLServer

3.1 Server Security Model

3.2 Backup/Recovery and Logging

3.3 Data Consistency vs. Concurrency

3.4 High Availability Concepts

3.1 Server Security Model 1/5

Gliederung durch Klicken hinzufügen

Feature	Definition	Oracle	Sybase	SQL Server	DB/2
Schema	- collection of logical structures of data - classification of objects (namespace)	rudimentary	-	new in 2005	yes
- Login Account - DB User	- right to access the database server - security domain and ■ Quotas ■ Privileges ■ Resource Limits	N / A User	Login DB User	Login DB User	

HVB Information Services
Member of UniCredit Group

3.1 Server Security Model 2/5

Gliederung durch Klicken hinzufügen

Feature	Definition	Oracle	Sybase	SQL Server	DB/2
- Groups	- collection of similar users	N/A	up to 11.5	via Windows Groups	
-Privileges (Authorization)	right to run a particular SQL statement or to access a database object	system pr. object pr.	system pr. object pr.	statement lev. pr. object pr.	database pr. object pr.
-Roles	named (protected) group of related privileges	yes	yes	yes	no

HVB Information Services Member of UniCredit Group

3.1 Server Security Model 3/5

Authentication: Identity Verification

Oracle	- Operating System (externally) - Network (globally) ■ 3rd Party (DCE, Kerberos) ■ PKI: SSL, OCI, Wallets, LDAP ■ Remote: Radius - Database: (password) - Multitier: (with OCI)

DCE: Distributed Computing Environment (De Facto Standard in the commercial Unix Environment)

DCE provides a complete Distributed Computing Environment infrastructure. It provides security services to protect and control access to data, name services that make it easy to find distributed resources, and a highly scalable model for organizing widely scattered users, services, and data. DCE runs on all major computing platforms and is designed to support distributed applications in heterogeneous hardware and software environments.

www.opengroup.org

3.1 Server Security Model 4/5

Authentication: Identity Verification

Sybase	- ASE (password) - external: ■ Kerberos (Network) ■ LDAP ■ PAM (OS)
SQL Server	- Windows Authentication (Integrated Security, NTLM or Kerberos) - SQL Server Authentication

PAM: Pluggable Authentication Modules

NTLM: Windows NT Lan Manager

Windows NT Challenge/Response (NTLM) is the authentication protocol used on networks that include systems running the Windows NT operating system and on stand-alone systems. NTLM stands for Windows NT LAN Manager, a name chosen to distinguish this more advanced challenge/response-based protocol from its weaker predecessor LAN Manager (LM).

3.1 Server Security Model 5/5

Authentication: Identity Verification

DB/2	- Server: comparison with OS - Server_ Encrypt - Client: comparison to valid credentials on client node - Kerberos - Krb_Server_Encrypt - add. Security Plugins may be installed (LDAP, PKI, …)

HVB Information Services
Member of UniCredit Group

3.2 Backup/Recovery and Logging: Oracle 1/6

Logical Backup

- import/export (proprietary format), impdp/expdp (XML)
- SQL Loader for large quantities of data

Physical Backup

- noarchivelog Mode: redologs will be overwritten in a circular fashion
  - no PITR possible
  - protection from instance failure
  - most recent changes available only for crash or instance recovery
- archivelog mode: redologs will be duplicated to an archivelog
  - complete PITR from instance or media failure possible
  - all changes are permanently saved in archivelogs

3.2 Backup/Recovery and Logging: Oracle 2/6

Glied

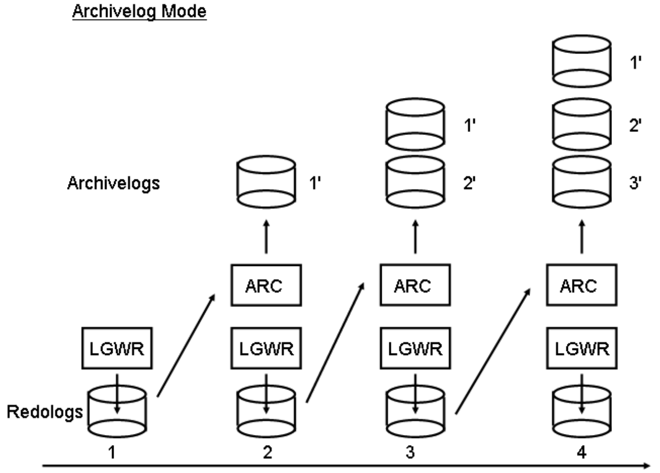

HVB Information Services
Member of UniCredit Group

3.2 Backup/Recovery and Logging: Oracle 3/6

- Offline Backup: copy of DB files while Database ist down
- Online Backup:
  - User Managed B&R:
    - begin .. end backup + OS commands
    - recover Statement
  - RMAN + (opt. Recovery Catalog) + API for Integration with 3rd – Party Media Manager

3.2 Backup/Recovery and Logging: Oracle 4/6

Gliede

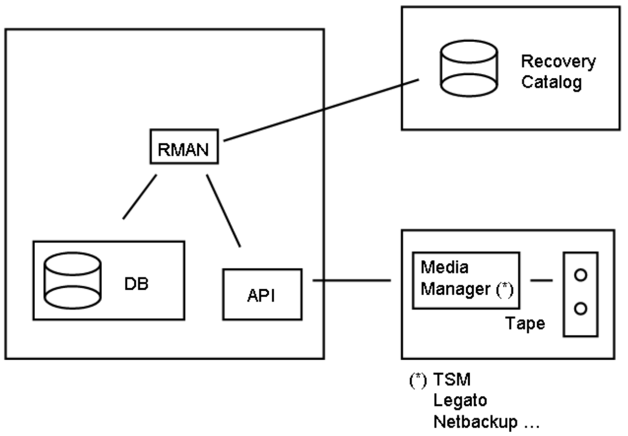

HVB Information Services Member of UniCredit Group

3.2 Backup/Recovery and Logging: Oracle 5/6

Gliederung durch Klicken hinzufügen

Feature	RMAN	User managed	Export
closed DB Backup	x	x	-
open DB Backup	x (no begin .. end Backup)	x (begin .. end Backup)	x (Undo required for consistency)
incremental Backup	x	-	-
corrupt Block Detection	x	-	x
automatic Backup	x	-	x (full, user, table) - Level

3.2 Backup/Recovery and Logging: Oracle 6/6

Gliederung durch Klicken hinzufügen

Feature	RMAN	User managed	Export
Backup catalog	x	-	-
Backup to Media Manager	x (API)	x	x
Backup init-File, Password – File	x	x	-
OS independent Language	x	-	x

3.2 Backup/Recovery and Logging: Sybase 1/3

Gliederur

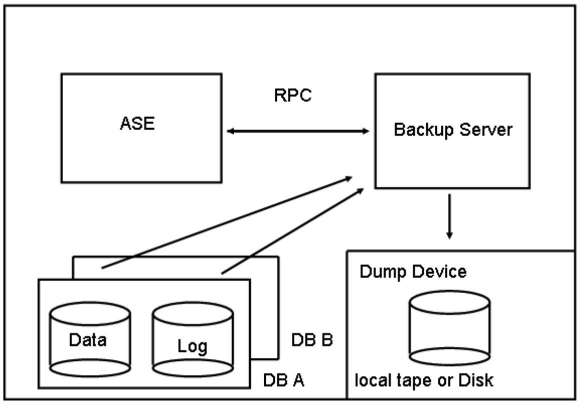

3.2 Backup/Recovery and Logging: Sybase 2/3

Physical Backup

- same Logging Model as SQL Server (per database one Transaction Log)
- dump database (only whole Database)
- dump transaction (thereafter Log will be truncated)
- threshold procedures have to be implemented
- database options may be set:
  - abort transaction on log full
  - truncate log on checkpoint (removal of committed transactions)
- Media Manager Integration through SQL*Backtrack (BMC) or Veritas Netbackup (Veritas)

3.2 Backup/Recovery and Logging: Sybase 3/3

Logical Backups

- bcp (for data, only at object level)
- ddlgen (for structures)

3.2 Backup/Recovery and Logging: SQL Server 1/3

Physical Backup

- auto growing tran log
  may be defined otherwise
  stored threshold procedure
- full backup (backup database)
  (local tape or disk)
- full differential backup
  (backup database with
  differential) based on full backup

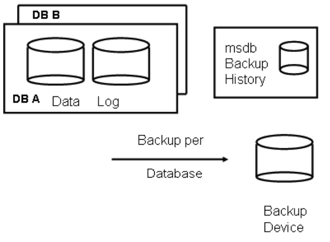

HVB Information Services
Member of UniCredit Group

3.2 Backup/Recovery and Logging: SQL Server 2/3

- partial Backup: Backup of write intensive Portions of a Database
- partial differential backup
- File backup, Filegroup Backup: backup of a Database at the File Level (faster Recovery)
- transaction Log backup: default autogrow enabled (backup log) (after Backup it will be truncated)
- API for split Mirror hot Backup available (with suspension of Database Activity)
- no integration of media management software

3.2 Backup/Recovery and Logging: SQL Server 3/3

Logical Backup

- bcp
- DTS (Data Transformation Service)
- new in 2005: SSIS (SQL Server Integration Service) – an ETL Tool

3.2 Backup/Recovery and Logging: DB/2 1/4

Logical Backup

- import/export (sql select in various formats)
  - Formats:
    - IXF (Integration Exchange Format) (see note)
    - WSF (Worksheet File Format) (IBM spreadsheet format)
    - DEL Delimited ASCII (seq. file with row and column del.)
    - ASC Undelimited ASCII (seq. file with row del.)
- load for large quantities of data

IXF: for data exchange between RDBMS

independent of exporting/importing DB

unbroken sequence of variable length records

H Header

T Table

C Column

D Rows

3.2 Backup/Recovery and Logging: DB/2 2/4

Logging

- circular Logging
  - Default
  - only logs with current transaction data (active logs) are kept
  - only full and Offline Backups ("Version Recovery")
  - protects from Instance Failure
- retained Logging
  - PITR possible
  - protects from Instance and Media Failure
  - Log Archiving via Userexit (call of external Program: diskcopy, TSM, tape, any XBSA compliant utility)

Version Recovery is the restoration of a previous version of the database using an image that was created during a backup operation.

XBSA X/Open Backup Services API

3.2 Backup/Recovery and Logging: DB/2 3/4

Backup per Database

Log Files

- Transaction Recording

Recovery History File

- Summary of Backup Information to determine Recovery Options

Tablespace Change History File

- determine Log Files for Recovery of a Tablespace

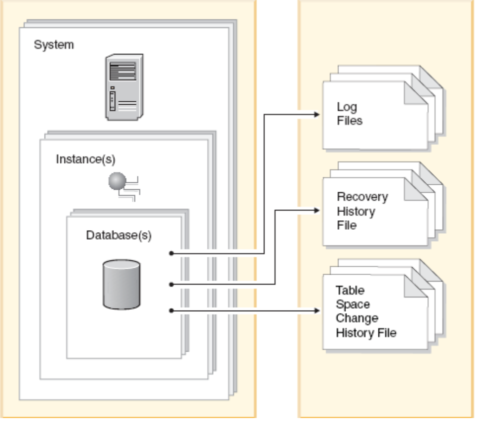

HVB Information Services
Member of UniCredit Group

3.2 Backup/Recovery and Logging: DB/2 4/4

Physical Backup

- retained (archive) Logging:
  - db2 backup <database> (Program)
  - db2backup (API)
  - incremental Backups (cumulative, delta) possible
  - split Mirror Backups
- circular Logging:
  - OS Backup only offline (Integration with TSM possible)

3.3 Data Consistency vs. Concurrency 1/27

Fundamental problem of Multi User Databases

How to ensure Data Consistency and maintain Concurrency at the same Time?

RDBMS handle this differently:

- Oracle: Multi Versioning
- Non Oracle (Sybase, DB/2, SQL Server): Read locking

3.3 Data Consistency vs. Concurrency 2/27

Multi Versioning

- read consistent queries: consistent results with respect to a given point in time
- non – blocking queries: queries are never blocked by writers of data

- Oracle is able to simultaneously materialize multiple versions of the data.

Read Consistent Queries:

any resultset would be current with respect to one of two points in time:

-The point at which a cursor was opened (default in Read Committed Isolation mode).

-The point at which the transaction to which the query belongs began (default in Read Only and Serializable isolation mode).

3.3 Data Consistency vs. Concurrency 3/27

Demo Example

- cursor does not fetch data
- delete command places data in undo segment

```
SQL> create table t
  2  as
  3  select * from all_users;

Table created.

SQL> variable x refcursor
SQL> begin
  2  open :x for select * from t;
  3  end;
  4  /

PL/SQL procedure successfully completed.

SQL> delete from t;

8 rows deleted.

SQL> commit;

Commit complete.

SQL> print x

USERNAME                          USER_ID CREATED
------------------------------ ---------- --------
SYSMAN                                 31 26.06.06
PERFSTAT                               25 26.06.06
DIP                                    19 26.06.06
TSMSYS                                 21 26.06.06
DBSNMP                                 24 26.06.06
SYS                                     0 26.06.06
SYSTEM                                  5 26.06.06
OUTLN                                  11 26.06.06

8 rows selected.

SQL>
```

3.3 Data Consistency vs. Concurrency 4/27

Transaction Isolation Levels

Purpose of a Transaction: Take DB from one consistent State to another.

In a Multiuser DB different Transactions are carried out simultaneously.

The Sensitivity of a Transaction to Changes made by others is determined via *Transaction Isolation Levels*.

Isolation Levels are defined in Terms of Phenomena.

3.3 Data Consistency vs. Concurrency 5/27

Gliederung durch Klicken hinzufügen

Phenomenon	Explanation
Dirty Read	read of uncommitted data no data integrity
Nonrepeatable Read	read of the same row at times t1,t2 may yield different results
Phantom Read	query at times t1,t2 (t2 > t1) may yield additional rows at t2

3.3 Data Consistency vs. Concurrency 6/27

Transaction Isolation Levels (x: permitted, -: not permitted)

Isolation Level	Dirty Read	Nonrepeatable Read	Phantom Read
L0: Read Uncommitted	x	x	x
L1: Read Committed	-	x	x
L2: Repeatable Read	-	-	x
L3: Serializable	-	-	-

3.3 Data Consistency vs. Concurrency 7/27

Locking and Default Isolation Levels (Non-Oracle)

- all locks are stored in memory
- locking schemes: row (page), table
- Concurrency and Overhead grow from table level locking to row level locking.

3.3 Data Consistency vs. Concurrency 8/27

Sybase	- Allpages Locking (data + index) - Datapages Locking - Datarows Locking - Lock Escalation may occur (page => table, row => table) - shared (for read) and exclusive locks - Default Isolation Level: L1
DB/2	- row Level Locking - Lock Escalation may occur. - Default Isolation Level: L1

3.3 Data Consistency vs. Concurrency 9/27

SQL Server	- Row Level Locking - Lock Escalation may occur (Row => Table) - Default Isolation Level : L1 - new in 2005: Snapshot Isolation (attempt to mimic Oracle Behaviour, uses tempdb)

3.3 Data Consistency vs. Concurrency 10/27

Locking and Isolation Levels (Oracle)

- Locks are stored in Data Blocks (no lock escalation).

- Locks only when data are modified.

- Oracle *always* uses non-blocking reads (Multi Versioning).

- *Writes do not block Reads and vice versa.*

- L1 and L3 are supported (L1 default)

- Intention of SQL Standard:

  - L0: to mimic non blocking reads
  - L1: does not give consistent results
  - L2: to guarantee read consistent results

3.3 Data Consistency vs. Concurrency 11/27

However, in Oracle:

- L1: because all queries are read consistent, it is equivalent to L2

- L0: The "spirit" of L0 is in Oracle always achieved (all reads are non – blocking)

- Additionally L4 (Read Only) Level is supported (L3 + no data modification).

  - Transaction sees only changes, that were committed when transaction began, but IUD are not allowed.

3.3 Data Consistency vs. Concurrency 12/27

Example: Take a look at a Transaction in all Isolation Levels and compare Oracle- and Non-Oracle Behaviour.

Row	Table Accounts		
	Account (PK)	Balance (t1)	Balance (t8)
1	123	500.00	100.00
2	456	240.25	240.25
3	781	0.00	0.00
4	987	100.00	500.00
Sum		840.25	840.25

3.3 Data Consistency vs. Concurrency 13/27

2 simultaneous Sessions:

- Session 1: select sum(balance) from account;
- Session 2: transfer 400.00 from Account 123 to Account 798.

Caption:

S: Session	O: Oracle like Database
R: Row	NO: Non – Oracle like Database
L: Level	
t: time	

3.3 Data Consistency vs. Concurrency 14/27

Read Uncommitted (L0) (Oracle does not support L0, so L1 is used.)

Time	S1 (NO) : L0	S1 (O) : L1	Transfer
T1	R1=500.00, Sum=500,00	R1=500.00 Sum=500.00	
T2	R2=240.25 Sum=740.25	R2=240.25 Sum=740.25	
T3			Update R1, excl. Lock on R1 set R1=100.00
T4	R3=0.00 Sum=740.25	R3=0.00 Sum=740.25	

3.3 Data Consistency vs. Concurrency 15/27

Time	S1 (NO) : L0	S1 (O) : L1	Transfer
T5			Update R4, excl. Lock on R4 set R4=500.00
T6	R4=500.00 Sum=1240.25	R4=100.00 Sum=840.25 (see Note)	
T7			Commit
T8	S=1240.25 Wrong!	Sum=840.25	

HVB Information Services Member of UniCredit Group

T6(O): Value as it existed, when query began (T5>T1), is retrieved from Undo.

3.3 Data Consistency vs. Concurrency 16/27

Read Committed (L1)

Time	S1 (NO) : L1	S1 (O) : L1	Transfer
T1	R1=500.00 Sum=500.00	R1=500.00 Sum=500.00	
T2	R2=240.25 Sum=740.25	R2=240.25 Sum=740.25	
T3			Update R1; excl. Lock Set R1=100.00
T4	R3=0.00 Sum=740.25	R3=0.00 Sum=740.25	

3.3 Data Consistency vs. Concurrency 17/27

Time	S1 (NO) : L1	S1 (O) : L1	Transfer
T5			Update R4; excl. Lock R4=500.00
T6	try to read R4 (is locked) S1 blocks, Query stops	R4=100.00 (from Undo) Sum=840.25	
T7			Commit;
T8	R4=500.00 S=1240.25 Wrong!	-	
The same Isolation Levels yield different answers!			

3.3 Data Consistency vs. Concurrency 18/27

Repeatable Read (L2) (not supported/needed in Oracle, uses L1)

Time	S1 (NO) : L2	S1 (O) : L1	Transfer NO	Transfer O
T1	R1=500.00 Sum=500.00 Lock R1	R1=500.00 Sum=500.00		
T2	R2=240.25 Sum=740.25 Lock R2	R2=240.25 Sum=740.25		
T3			try to update R1 (is locked) Update suspended	Update R1; excl. Lock Set R1=100.00
T4	R3=0.00 Sum=740.25 Lock R3	R3=0.00 Sum=740.25		

3.3 Data Consistency vs. Concurrency 19/27

Time	S1 (NO) : L2	S1 (O) : L1	Transfer NO	Transfer O
T5				Update R4; excl. Lock Set R4=500.00
T6	R4=100.00 Sum=840.25 Lock R4	R4=100.00 (from Undo) Sum=840.25		
T7				Commit;
T8	Commit; release all Locks			

3.3 Data Consistency vs. Concurrency 20/27

Time	S1 (NO) : L2	S1 (O) : L1	Transfer NO	Transfer O
T9			Update R1 Set R1=100.00 excl. Lock	
T10			Update R4 Set R4=500.00 excl. Lock	
T11			Commit;	

3.3 Data Consistency vs. Concurrency 21/27

Remarks

- Correct Answer also in NO.
- Cost: Transactions had to be executed sequentially!
- Readers of Data will block Writers of Data!
- Spurious Errors still possible: see next Page.

3.3 Data Consistency vs. Concurrency 22/27

Gliederung durch Klicken hinzufügen

Time	S1 (NO) L2	Transfer 50.00 from R4 to R1
T1	R1=500.00 Sum=500.00 Lock R1	
T2	R2=240.25 Sum=740.25 Lock R2	
T3		Update R4=50.00; excl. Lock
T4	R3=0.00 Sum=740.25	
T5		try to Update R1 (is locked) Update suspended
T6	try to read R4 (is locked) Query suspended	
Classical Deadlock One Transaction will be chosen as victim, killed and rolled back.		

HVB Information Services
Member of UniCredit Group

3.3 Data Consistency vs. Concurrency 23/27

Porting Issues

In DB with shared Read Locks and L2: Lost Updates cannot happen.

But:

In DB without shared Read Locks: Lost Update Problem!

3.3 Data Consistency vs. Concurrency 24/27

Lost Update Problem (Database without Shared Read Locks)

T1: S1: read R1

T2: S2: read R1

T3: S1: Update R1; Commit;

T4: S2: Update R1; Commit:

⇒ Result from T3 is lost!

How to avoid: Manual Locking (optimistic or pessimistic)

optimistic Locking: defer locking as late as possible (Application Design!)

pessimistic Locking: select for update (lock for write, not for read)

3.3 Data Consistency vs. Concurrency 25/27

Serializable (L3)

Each Transaction appears to be the only Transaction in the Database at a given Time.

Oracle:

- L3 extends the Read Consistency from Statement to Transaction Level (using Undo).

- Price: ORA-8177: can't serialize access for this transaction

  (when updating data which has changed since the begin of the transaction)

Non Oracle:

- locking may yield Deadlocks, Blocking, …

3.3 Data Consistency vs. Concurrency 26/27

Read Only (L4) (Oracle Only)

- as L3 but no modifications allowed => no ORA-8177
- in other systems: L2
- because in Oracle without Locking: ORA-1555 might happen.

ORA-1555: Snapshot too old.

3.3 Data Consistency vs. Concurrency 27/27

Gliederung durch Klicken hinzufügen

Isolation Level	Implementation	Writes Block Reads	Reads Block Writes	Deadlock Sensitive Reads	Incorrect Query Results	Lost Updates	Lock Escalation or Limits
L0	NO	n	n	n	y	y	y
L1	NO	y	n	n	y	y	y
	O	n	n	n	n	n (*)	n
L2	NO	y	y	y	n	n	y
L3	NO	y	y	y	n	n	y
	O	n	n	n	n	n	n

(*) with select for update

3.4 High Availability Concepts: Oracle 1/14

- Real Application Cluster
- Flashback
- Data Guard
- Replication (Streams)

3.4 High Availability Concepts: Oracle 2/14

(a) Real Application Cluster (RAC) and Oracle ClusterWare (OCW)

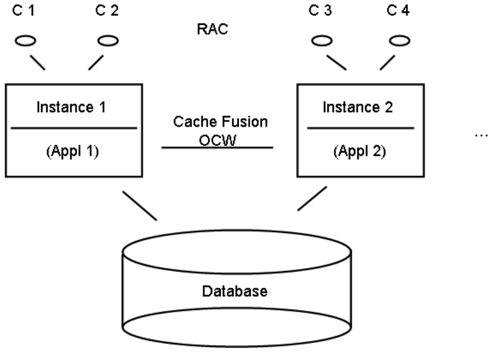

High Availability and horicontal Scaling

3.4 High Availability Concepts: Oracle 3/14

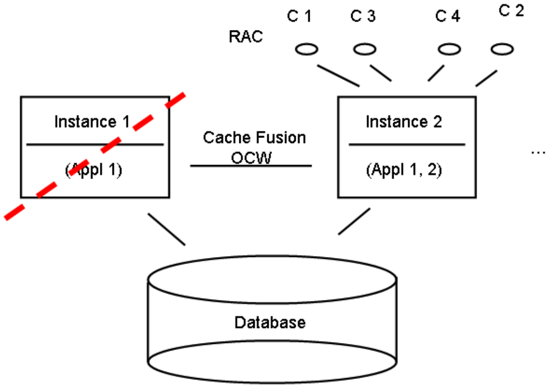

HVB Information Services
Member of UniCredit Group

3.4 High Availability Concepts: Oracle 4/14

Oracle Clusterware

- comes from Tru-Cluster (DEC)
- OCW + ASM:
  - Ressource Monitoring
  - Ressource Management
  - Cluster Control
  - Failover APIs
  - Cluster Membership
  - LVM

OCW: Oracle Clusterware

ASM: Automatic Storage Management

3.4 High Availability Concepts: Oracle 5/14

OCW:

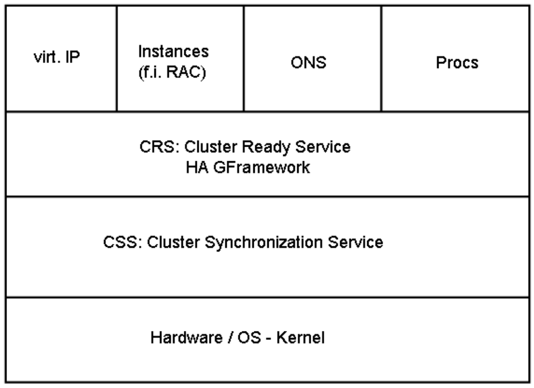

ONS: Oracle Notification Server

3.4 High Availability Concepts: Oracle 6/14

additional features of OCW:

- Cluster for any Application
- Loadbalancing, Workloadmanagement
- Fast Application Notification (FAN) for subscribed Clients
- Transparent Application Failover (TAF)

3.4 High Availability Concepts: Oracle 7/14

(b) Flashback – Protection from logical Errors

Feature	Purpose
Flashback Versions Query	to reproduce historical Operations on Database Objects
Flashback Transaction Query	generate Undo-SQL for historical Operations
Flashback Table	set Database Objects back in Time
Flashback Drop	restore Database Objects after dropping them
Flashback Database	set whole Database back in Time

3.4 High Availability Concepts: Oracle 8/14

(c) DataGuard (Standby Database)

Type of Standby Database	Purpose /Protection from
Physical Standby Database	- User Errors, logical Corruptions - Disaster Recovery - Fast Site and Database Failover - Planned Switchover for Maintenance - Divergence for Reporting and Testing (Flashback) - taking Backups / Load ! - Read Only Capability

3.4 High Availability Concepts: Oracle 9/14

Type of Standby Database	Purpose /Protection from
Logical Standby Database	- Disaster Recovery - Read/Write Capability for normal Operation - additional Objects may be created for Reporting - Rolling Database Upgrades

Maximum Availability Architecture: RAC + DataGuard

3.4 High Availability Concepts: Oracle 10/14

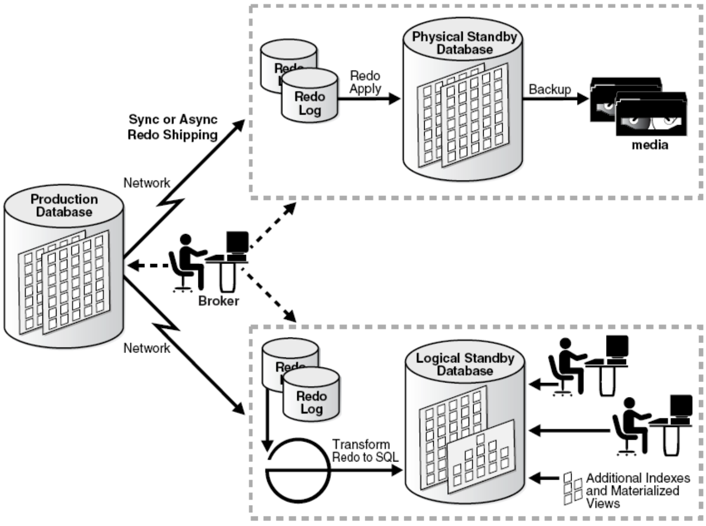

3.4 High Availability Concepts: Oracle 11/14

(d) Replication: Oracle Streams

- highest Degree of Flexibility
- highly customizable
- heterogeneous Platform Support
- high Configuration and Administration Flexibility

3.4 High Availability Concepts: Oracle 12/14

1. capture Messages at a Source Database
2. stage Messages in Queues
3. propagate Messages to another Queue
4. consume Messages
5. apply Changes in Destination Database

3.4 High Availability Concepts: Oracle 13/14

Gliederung durc

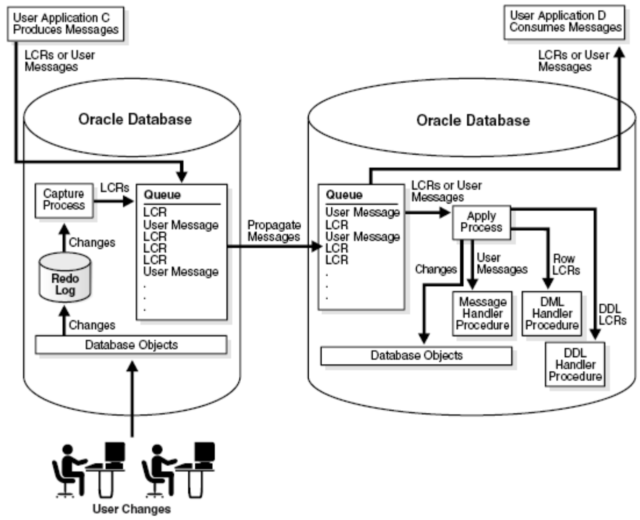

When you use Streams, replication of a DML or DDL change typically includes three steps:

A capture process or an application creates one or more logical change records (LCRs) and enqueues them into a queue. An LCR is a message with a specific format that describes a database change. A capture process reformats changes captured from the redo log into LCRs, and applications can construct LCRs. If the change was a data manipulation language (DML) operation, then each LCR encapsulates a row change resulting from the DML operation to a shared table at the source database. If the change was a data definition language (DDL) operation, then an LCR encapsulates the DDL change that was made to a shared database object at a source database.

A propagation propagates the staged LCR to another queue, which usually resides in a database that is separate from the database where the LCR was captured. An LCR can be propagated to a number of queues before it arrives at a destination database.

At a destination database, an apply process consumes the change by applying the LCR to the shared database object. An apply process can dequeue the LCR and apply it directly, or an apply process can dequeue the LCR and send it to an apply handler. In a Streams replication environment, an apply handler performs customized processing of the LCR and then applies the LCR to the shared database object.

Step 1 and Step 3 are required, but Step 2 is optional because, in some cases, an application can enqueue an LCR directly into a queue at a destination database. In addition, in a heterogeneous replication environment in which an Oracle database shares information with a non-Oracle database, an apply process can apply changes directly to a non-

3.4 High Availability Concepts: Oracle 14/14

Gliederung durch

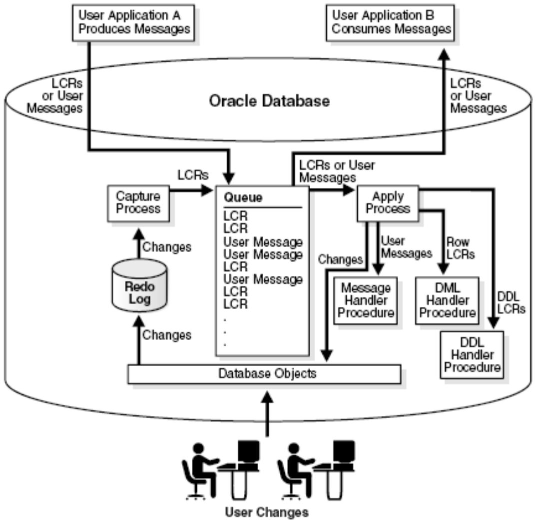

HVB Information Services Member of UniCredit Group

3.4 High Availability Concepts: Sybase 1/9

(a) Dump/Load (Poor Man's Solution)

Primary	Secondary
Dump DB	
	Load DB
Dump Transaction Log	
	Load Transaction Log
Dump Transaction Log	
	Load Transaction Log
Fail	

3.4 High Availability Concepts: Sybase 2/9

(a) Dump/Load (Poor Man's Solution)

Primary	Secondary
Dump Transaction Log (if possible)	
	Load Transaction Log (if available)
	Online DB
	Point Clients here (manually or automatically)

• Drawback: secondary unavailable prior Online DB

HVB Information Services
Member of UniCredit Group

3.4 High Availability Concepts: Sybase 3/9

(b) Dump/Load (Poor Man's Solution) with Standby Access

Primary	Secondary
Dump DB	
	Load DB
Dump Transaction Log with standby_access	
	Load Transaction Log Online DB for standby_access
Dump Transaction Log with standby_access	

3.4 High Availability Concepts: Sybase 4/9

(b) Dump/Load (Poor Man's Solution) with Standby Access

Primary	Secondary
	Load Transaction Log Online DB for standby_access
Fail Dump Transaction Log (if possible)	
	Load Transaction Log (if available)
	Online DB

3.4 High Availability Concepts: Sybase 5/9

(b) Dump/Load (Poor Man's Solution) with Standby Access

Primary	Secondary
	Point Clients here (automatically or manually)

Advantage: Secondary available for non change Queries during normal Operation

3.4 High Availability Concepts: Sybase 6/9

(c) Database Device Mirroring: Protection against Disk Failures

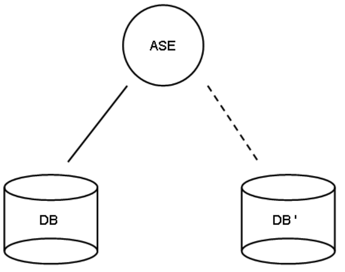

3.4 High Availability Concepts: Sybase 7/9

(d) Replication Server (or ASE Replicator)

- Publish/Subscribe Model
- Warm Standby Model (Database Granularity, also DDL Replication)
- Latency Period
- automatic Failover
- automatic Client Failover (includes Application Change)

3.4 High Availability Concepts: Sybase 8/9

(e) Open Switch

- no application change resp. client redirect

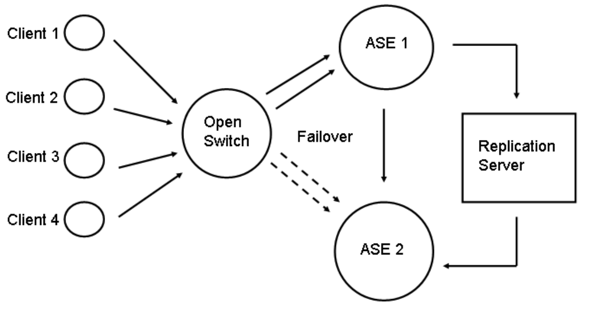

3.4 High Availability Concepts: Sybase 9/9

(f) HA Companion Server

- depends on Cluster Technology
  - SunCluster
  - HACMP
  - MC/ServiceGuard
  - VeritasCluster
- ASE 2 in Standby Mode or with own Databases
- Client Redirect automatic through network interfaces

3.4 High Availability Concepts: SQL Server 1/4

(a) Database Mirroring

- per Database
- possible in conjunction with Microsoft Cluster Service
- mirror Database available for limited reporting Needs
- optional "Witness" for automatic Failover
- Client redirection is accomplished through local Caching of the Server Names

3.4 High Availability Concepts: SQL Server 2/4

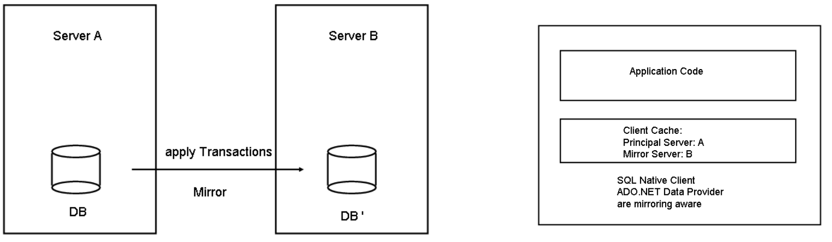

3.4 High Availability Concepts: SQL Server 3/4

(b) Database Snapshots

- read only Databases as exact copy of a User Database in a specific Point of Time
  - at Page Level
  - Copy – on – Write (only modified Blocks will be copied)
  - for Reporting
  - for Recovery Purposes

3.4 High Availability Concepts: SQL Server 4/4

(c) Replication

- copy Databases and Database Objects from one Database to another and synchronize
- transactional
- also as multi-master variant (peer – to – peer)
- Replication with non SQL Server Databases is supported

(d) Log Shipping

- same Principle as Oracle Standby Database

3.4 High Availability Concepts: DB/2 1/7

(a) HADR (High Availability Disaster Recovery)

- takeover command for role switch
- automatic client reroute (client has both connect information)

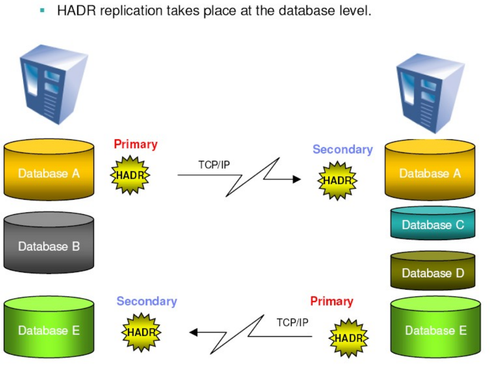

3.4 High Availability Concepts: DB/2 2/7

Running HADR

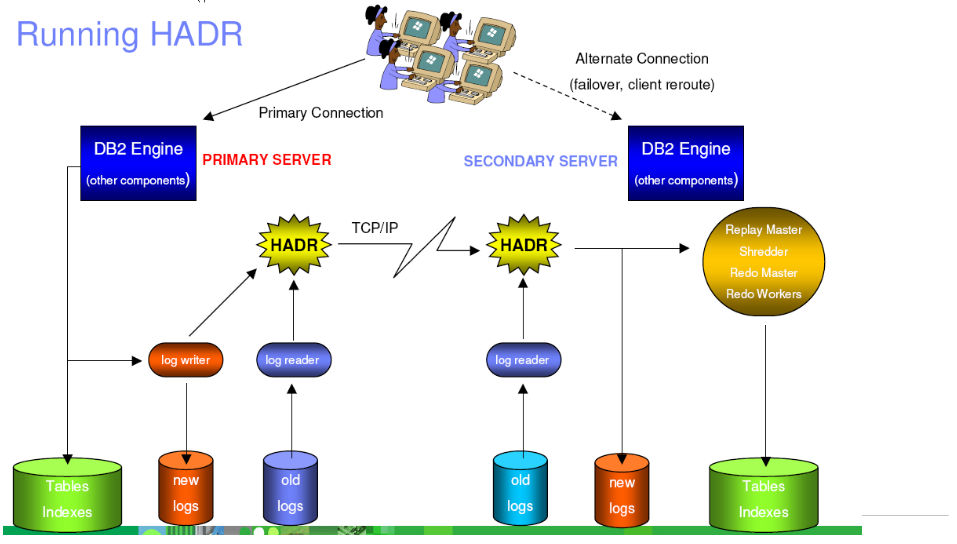

3.4 High Availability Concepts: DB/2 3/7

- Synchronous, Near Synchronous, Asynchronous

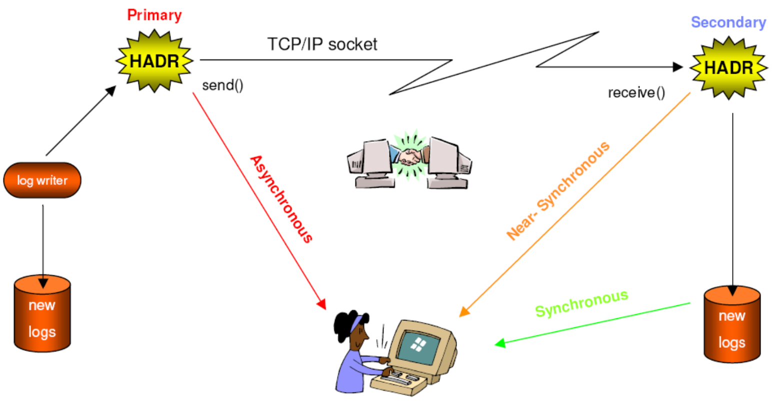

Matthias Mann / HIS14DB, 28.06.2006
Seite 156

HVB Information Services
Member of UniCredit Group

Synchronization Modes

synchronuous:

commit succeeds, if log data is on disk on primary and secondary

near synchronuous:

commit succeeds, if log data is on disk at primary and has been received by secondary

asynchronuous:

commit succeeds, when log data is on disk at primary and has been sent to secondary

3.4 High Availability Concepts: DB/2 4/7

(b) Standby Database

- create Standby DB via Split Mirror or by Backup/Restore

- Log Shipping:
  - copy whole Log Files to standby Machine
    - from Archive Device
    - or directly via User Exit
  - secondary continuously rolls forward through Logs
  - when Primary fails:
    - copy over remaining Logs

3.4 High Availability Concepts: DB/2 5/7

- db2 rollforward database <name> to end of logs and stop
- Clients reconnect to Standby Database

- Replication:

- use Data Propagator for Log Capturing at Primary
  - (Capture and Apply Subcomponents)
  - Standby DB available for Read Access
  - BiDi Replication possible
  - no need for physically identical Storage Partitions
  - Restriction to critical Tables possible

3.4 High Availability Concepts: DB/2 6/7

(c) MQ Replication

- Each message represents a transaction
- XML format option for publishing
- Highly parallel apply process
- Differentiated conflict detection and resolution
- Staged availability of heterogeneous support

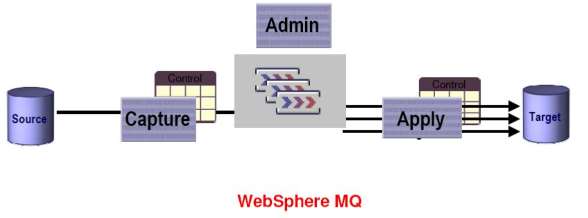

HVB Information Services Member of UniCredit Group

3.4 High Availability Concepts: DB/2 7/7

(d) Data Partitioning Feature

- not for High Availability but for Scalability
- Shared Nothing Architecture
- Agents may fail
- Hash Partitioning

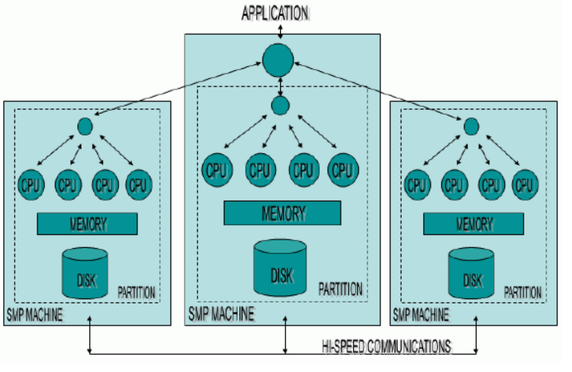

HVB Information Services Member of UniCredit Group

4. Vendor Comparison: Beyond Technology

4.1 Oracle Product Portfolio 1/3

Gliederung durch Klicken hinzufügen

Database	Oracle RDBMS
Middleware	Oracle Fusion Middleware - Application Server - Identity Management - Directory Services
Developer Tools	- Developer Suite ■ Developer ■ Discoverer ■ Designer - BPEL Designer (see note)

Business Process Engineering Language

4.1 Oracle Product Portfolio 2/3

Gliederung durch Klicken hinzufügen

Enterprise Management	Administration Interface and Repository for managing all Database Instances (Grid Control)
Business Intelligence Project "Fusion" (all Products have to be merged)	- Oracle E-Business Suite - Peoplesoft Enterprise - Siebel CRM - JD Edwards

4.1 Oracle Product Portfolio 3/3

Collaboration Gliederung durch Klicken hinzufügen	Collaboration Suite - Content Services - Real Time Collaboration - Messaging - Workspaces
Data Warehouse	- Oracle Warehouse Builder - Oracle OLAP

HVB Information Services
Member of UniCredit Group

4.1 Sybase Product Portfolio 1/3

Gliederung durch Klicken hinzufügen

Database (OLTP)	Adaptive Server Enterprise (ASE)
Heterogeneous Data Services	- Sybase IQ (BI, DWH) - Replication Server - Dynamic Archive - Mirror Activator - Real Time Data Services

4.1 Sybase Product Portfolio 2/3

Information Anywhere Gliederung durch Klicken hinzufügen	- SQL Anywhere - Adaptive Server Anywhere (ASA) - RFID Anywhere/Enterprise - Mobile Link - Afaria - Content Integrator - Enterprise Portal - Unwired Accelerator - Unwired Orchestrator

4.1 Sybase Product Portfolio 3/3

Gliederung durch Klicken hinzufügen

On Demand Mobility Services	- AvantGo
Development	- Workspace - Power Designer - Power Builder Family - EA Server

4.2 Platform Support / Market Shares 1/3

Gliederung durch Klicken hinzufügen

Operating System	Oracle	DB/2 UDB 8.2	Sybase ASE/RS 15	SQL Server 2005
Unix				
■ Solaris	x	x	x	-
■ AIX	x	x	x	-
■ HP-UX	x	x	x	-
■ Tru64	x	-	x	-
■ Linux	x	x	x	-
■ Irix (SGI)	<= 8.1.7	-	x	-

HVB Information Services Member of UniCredit Group

only Operating Systems, no Hardware Platforms

4.2 Platform Support / Market Shares 2/3

Operating System	Oracle	DB/2 UDB 8.2	Sybase ASE/RS 15	SQL Server 2005
Open VMS	x	-	-	-
z/OS (OS 390) BS 2000	x <= 9.2	x -	- -	- -
OS 400	-	x	-	-
Microsoft Windows	x	x	x	x
Novell Netware	<= 8.1.7	-	-	-
Mac OS (>= X)	x	-	x	-

Gliederung durch Klicken hinzufügen

HVB Information Services Member of UniCredit Group

only Operating Systems, no Hardware Platforms

4.2 Platform Support / Market Shares (see note) 3/3

Gliederung durch Klicken hinzufügen

Vendor	Sales Volume 2005 (bn. USD)	Market Share 2005 (%)	Sales Volume 2004 (bn. USD)	Market Share 2004 (%)
Oracle	6.721	48.6	6.234	48.9
IBM	3.040	22.0	2.860	22.4
Microsoft	2.073	15.0	1.778	13.9
Teradata	0.441	3.2	0.412	3.2
Sybase	0.407	2.9	0.383	3.0
Other	1.135	8.2	1.090	8.5
Total Market	13.817	100.0	12.758	100.0

Source: Gartner Group: Dataquest, 2006

Sales Volume:

•License Fees

•Updates

•Support

•technical Consulting

4.3 Cost Comparison

The numbers in this chapter are from

- Microsoft SQLServer 2005: Understanding Database Pricing (11/05, www.microsoft.com)

4.3 Cost Comparison: Pricing - Editions

Gliederung durch Klicken hinzufügen

Edition	Features	Price/CPU (1000 USD / CPU)
Free	- limited DB functionality - memory, CPU, DB size limits	0
Basic	- simplified Mgmt. Tools - basic Security - CPU, Memory Limits	0.5 – 5.0

4.3 Cost Comparison: Pricing - Editions

Edition	Features	Price/CPU (1000 USD / CPU)
Standard	- full DB functionality - Basic Mgmt. Tools - CPU limits	5.0 – 15.0
Enterprise	- HA - Scalability - High End Management Tools - Enterprise Security - no Limits	25.0 – 40.0

Gliederung durch Klicken hinzufügen

4.3 Cost Comparison: Pricing - Options

Gliederung durch Klicken hinzufügen

	SQL Server 2005	Oracle 10g	DB/2 8.2
EE	1.0	1.6	1.0
Mgmt. AddOn	0.0	0.5	0.8
OLAP AddOn	0.0	0.8	1.1
DataMining AddOn	0.0	0.8	2.4

4.3 Cost Comparison: SQL Server - Oracle

Gliederung durch Klicken hinzufügen

	Oracle 10g		SQL Server 2005	
Standard		1.0		0.4
Enterprise	EE	2.7	EE	1.7
	Mgmt. Pack	0.4	Mgmt. Tools	0.0
	Adv. Sec.	0.7	Netw. Encr., SSO, PKI	0.0
	BI	2.7	Integr., Rep. Analysis Serv.	0.0
HA	add. Inst. EE + above Opt.	6.5	add. Inst. EE	0.0

4.3 Cost Comparison: SQL Server – DB/2

Gliederung durch Klicken hinzufügen

	DB/2 8.2		SQL Server 2005	
Standard		1.0		0.8
Enterprise	EE Warehouse Manager OLAP Server	3.3 1.4 3.7	EE Integr.Services Reporting + Analysis Serv.	3.3 0.0 0.0
HA	add. Inst. EE + above Opt.	8.4	add. Inst.	0.0

HVB Information Services

Member of UniCredit Group

4.3 Cost Comparison: TCO Saving SQL Serv. vs. Comp.

Saving = 100%- TCO(MSSQL/Comp.)

Cost factor	Departm. App. (50-100 Users)	Organisational App. (100+ Users)
Ongoing Operations	33%	45%
Design/Development	24%	13%
Software Licensing	67%	67%
Hardware	54%	54%
Maintenance	56%	56%
Training	54%	60%

Dtudy conducted by NerveWire
(http://www.microsoft.com/sql/evaluation/compare/tco.asp)

Period considered: 3 years

MS SQL 2005

DB/2 8.2

Oracle 10g

4.4 Total Cost of Administration: SQL Server vs. Oracle

The numbers in this chapter are from

- Microsoft SQL Server and Oracle Databases: A Comparative Study on Total Cost of Administration (TCA) (05/06, Alinean, Inc., www.alinean.com)

4.4 Total Cost of Administration: SQL Server vs. Oracle

Gliederung durch Klicken hinzufügen

Figure	SQL Server	Oracle
average Number of DB per company	107	87
average Number of Users per Database	328	716
mission critical Databases (%) (see note)	66.1	63.8
OLTP Databases (%)	55.7	60.3
DSS Databases (%)	44.3	39.7
Databases per DBA	31.2	9.9
Users per DBA	6784	5567
annual TCA per Database (USD)	2847	10206
annual TCA per User (USD)	13.09	18.15

percentage of oracle database resp sql server databases (classification of the corresponding application)

4.4 Total Cost of Administration: SQL Server vs. Oracle

DBA Training (Hours/Year):

- SQL Server: 52.2
- Oracle: 63.1

Database Administrator Experience Levels (%)

	SQL Server DBA	Oracle DBA
Junior (0-5)	38	25
Intermediate (5-8)	38	38
Senior (8+)	24	37

4.4 Total Cost of Administration: SQL Server vs. Oracle

Average Time / Week / Database (Hours) (see note)

Activity	SQL Server	Oracle
Design	1.1	2.2
Tuning	0.8	1.8
Problem	0.8	2.1
Security	0.4	1.2
Capacity	0.3	0.8
Backup	0.7	1.3
Reporting	0.5	1.6

total time of all dbas within an organization allocated to various administration tasks / number of databases (sql or oracle)

Excursus 1: Oracle Basic Design Decisions 1/2

Under market pressure design decisions had to be done which up to now influence the product.

- Transactional concept: non blocking queries, no read locks (see Note 1)
- Architecture: Development on PDP-11 (see Note 2)
  - Two Task: separation of application process and DB engine
    - comm. with DB engine per IPC (Message Passing)
    - sharing of control information in Shared Global Area (SGA)
    - foundation for SQL*Net

1. Transactional concept: - bevor image file (predecessor of undo) for queries and rollback
2. PDP-11: 16 bit address space, => execs > 32kB needed overlays => performance overhead => Baby Oracle should run as executable < 32 kB
   1. PDP–11: good on IPC

3. SGA: Shared Global Area => (Version 4 on PC -> single user DB) => System Global Area

5. same story: Undo – segments => rollback segments (Version 6) => undo segments (Version 9)
6. same story: SCN (System Change Number => System Commit Number)
7. same story: MTS is not multi threaded

Excursus 1: Oracle Basic Design Decisions 2/2

- 1985: IBM world market leader (DBMS)
  - OLTP: IMS (no Ad Hoc, instead precompiled queries)
  - DSS: DB/2
- Oracle: small budget => Ad Hoc Queries without precompilation

  Parsing

  Version 2: Rule Based Optimizer (heuristic, static rules)

Excursus 2: Early Tool Development 1/3

- Oracle 2: Appl. Programming required Knowledge in proc. Language
  - C, Fortran, Pascal, Cobol
  - expensive, difficult for On-Line User Interfaces
- Oracle had available:
  - UFI (User Friendly Interface, => SQL*Plus)
  - RPT (Report Production Utility)

Excursus 2: Early Tool Development 2/3

- Oracle wanted to create a procedural environment from scratch => was abandoned
  - without proper investigation of
    - user requirements
    - performance considerations
    - mathematical properties of the procedural language
- 1988: PL/SQL (came from public ADA + Syntax for embedding SQL)
- UFI => AFI (Advanced Friendly Interface)
  - Goal:
    - User can edit arbitrary Ad Hoc commands
    - command buffer exchange to system editor (see note)

Development Platform: VAX/VMS: system editor: EDT = from here comes the well known afiedt.buf file

Excursus 2: Early Tool Development 3/3

- 1985 – 1987: Oracle was used for small OLTP applications
- Oracle started Development of nonprocedural Tools for AdHoc Query Formulation
  - Easy*SQL, FastForms (Vers. 4.1) => did not survive
  - IAF (Interactive Application Facility)
  - breakthrough: SQL*Forms 2.1 (Oracle 5.1, 1987)

Excursus 3: Oracle Customer Relations 1/3

- 1980s: IT was Mainframe dominated, User Dissatisfaction
- Oracle Marketing:
  - if your IT Dept. cannot meet your Needs => buy some Mini-Computers (VAX) of your Choice + some Terminals
  - Oracle will run everywhere
  - Applications you can program with our Tools (there was no Standard Software in these Days)
  - if you are not able: Oracle can do it for you!
- begin of decentralization

Excursus 3: Oracle Customer Relations 2/3

- Oracle raised Expectations
- Problems:
  - Support was needed during Installation
  - Scalability problems
    - Version 5: maximal 40 (42?) simultaneous Sessions
    - Forms did not work with Transaction Monitors (Tuxedo)
    - analyzing Performance Problems not possible
    - Solution: bigger Machines, good for License Dept.
    - V6: 120-250 sess., V7: 1250 sess.

40 sessions: Douglas Adams (42!) used MicroVAX2 + Oracle 5 as word proceessor

Excursus 3: Oracle Customer Relations 3/3

- Unix enters the Datacentre
- typically: dedicated Server per Application
    - transparent for Accounting and Provisioning
    - convenient for Users (nobody disturbs)
    - but: Users had to administrate their own Servers (Patches, Backups, …)
- growing Applications could no longer exist isolated
- beginning from Version 5 the Market was the driving Force concerning the functional Requirements of the RDBMS

Excursus 4: Oracle Platform Support 1/2

- V3-V6: VAX/VMS was development platform
  - process creation: huge CPU overhead
- growing number of supported platforms: Oracle sold "Porting Kits" to hardware vendors (no need to develop "Operating System Dependent Layer")
- large number of OS in data centers (=> market share!)
- many OS were abandoned
- Oracle ported in joint programs with most important hardware/OS vendors

Excursus 4: Oracle Platform Support 2/2

- 1987: Sequent + Oracle 5.1 : first SMP port (see note 1)
- 1990: nCube: MPP support
- Oracle opted also for the mainframe market (IBM dominated)
  - Version 5: VM/CMS + MVS
  - Oracle was not able to run its own business process on mainframe
  - mainframe business was a flop in the long run
- PC market: 1986 Oracle 4.1 "The Cube" (see note 2)
  - breakthrough with NT 3.51 and OS/2

HVB Information Services Member of UniCredit Group

1. both companies were owned by Oracle
2. box with diskettes and manual: Oracle 4.1 + UFI + RPT + FastForms needed >640kB memory

Appendix A: The 12 Codd Rules (*) 1/7

(*) Codd, E.F. "Is Your DBMS Really Relational?" and "Does Your DBMS Run By the Rules?"

ComputerWorld, October 14 1985 and October 21 1985

1. Information Rule

- all data in tables (relations)
- tables consist of rows (tuples), no row ordering
- each row has the same columns
- column: scalar, no redefinition by the table definition, unique in table

2. Guaranteed Access Rule

Each datum has to be accesible without ambiguity (unique, primary keys!).

3. Systematic Treatment of NULL Values (see notes)

3. early Oracle: NULL="empty string" (tragical error)

Appendix A: The 12 Codd Rules 2/7

4. Dynamic On-Line Catalog

- same access to the structure of the DBMS like data access
- "Dictionary"

-Oracle: Version 6: X$- and V$-Views yield SQL access to control and performance informations

Appendix A: The 12 Codd Rules 3/7

5. Comprehensive data Sublanguage

- SQL (many dialects)
  - declarative, non – procedural, no sequentiality, no loops, no branches
  - query result is also a table (relation)

    select xx.zero_sal_qty

    from (*select count(*)* zero_sal_qty *from emp where sal=0*) xx;

    xx .. Table alias, zero_sal_qty .. column name

    Oracle: View : Query in dictionary as table

Appendix A: The 12 Codd Rules 4/7

6. View Updating Rule

- Each data manipulation valid against a table has also to be valid against a view.

7. High – Level Insert, Update, Delete

- I,U,D has to be possible for any retrievable number of rows
  - update … set ..=.. where …;

8. Physical Data Independence

- SQL Programming possible without knowledge about storage of data.
  - DML: o.k.
  - DDL: not o.k. (mixture of log. and phys. specs) (see notes)

OFA: Best Practices

phys. location of datafiles can be changed without recreating the database

Oracle 10g: ASM implements this rule

Appendix A: The 12 Codd Rules 5/7

9. Logical Data Independence

Data appearance should not change when table definitions change => views.

Very difficult in full generality (Normalization !).

10. Integrity Independence

- Definition of constraints forces data integrity.

- up to Oracle 7: unique index, NOT NULL constraint (enforced)

  - Application was responsible for data integrity.

- from Oracle 7: foreign key constraints, check constraints

  - Avoidance of user- or application errors.

Appendix A: The 12 Codd Rules 6/7

11. Distribution independence

- DBs may be distributed on different servers.
- Partly obsoleted by n-Tier Architectures and hardware developments.
- Oracle 5: Distributed Database Feature (query only)
- Oracle 7: 2PC
- Rule 2 may be violated when a table is located on different servers.

Appendix A: The 12 Codd Rules 7/7

12. Nonsubversion Rule

- No backdoor in the DB.
- All changes through SQL.
- Partly violated => SQL Loader, external tables

What is missing?

- transactional integrity (commit, rollback) (ACID)
- isolation levels
- privileges

Appendix B: Client / Server Architecture: Oracle 1/4

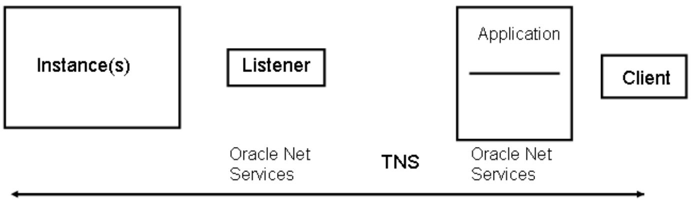

- TNS: Transparent Network Substrate
  - works with any standard network transport protocol
- Client may als be an Instance
- dedicated Server Arch. / Shared Server Arch. (see Note)
- Transparent Gateways for Connection to foreign RDBMS

dedicate Server architecture: one server process per client connection

shared server architecture: connection pooling

Oracle Net Services: support diffferebt naming sources (flat file, ldap, easy naming)

Appendix B: Client / Server Architecture: Sybase 2/4

Gliede

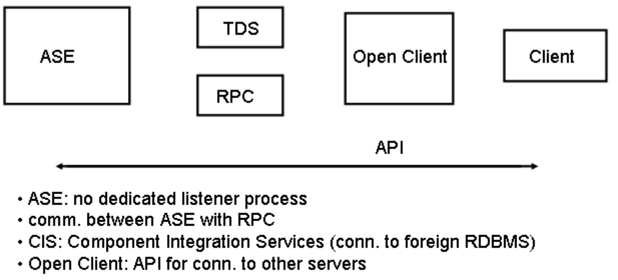

HVB Information Services
Member of UniCredit Group

Appendix B: Client / Server Architecture: SQL Server 3/4

Gliede

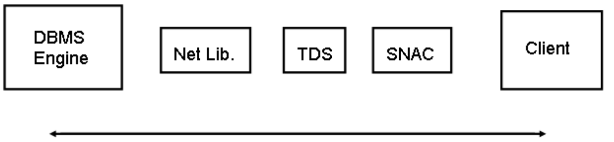

SNAC .. SQL Server Native Client (only 2005) (OLEDB+ODBC)

TDC .. Tabular Data Stream (data transfer protocol)

Appendix B: Client / Server Architecture: DB/2 4/4

Gliederung

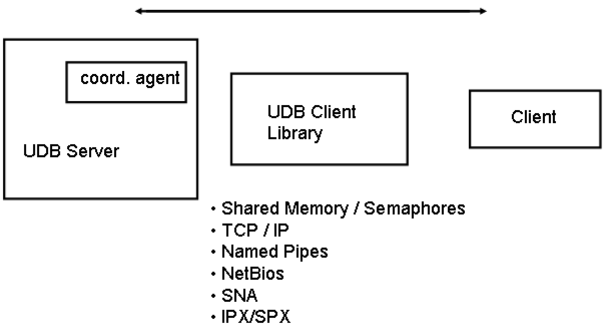

App. C Data Encryption / Network Security: Oracle 1/2

- DBMS_CRYPTO Package for selectively encrypting sensitive Data
  - DES
  - AES
  - SHA-1 (see note)
- listener is password protected
- Virtual Private Database: provides database integrated access control

  (access by different user groups to the same data: only relevant data are visible)
- Oracle Label Security: preconfigured and implemented VPD
- Network Traffic: ASO

SHA: Secure Hash Algorithm

MD5: Message Digest

AES: Advanced Encryption Standard

DES: Data Encryption Standard

App. C Data Encryption / Network Security: Oracle 2/2

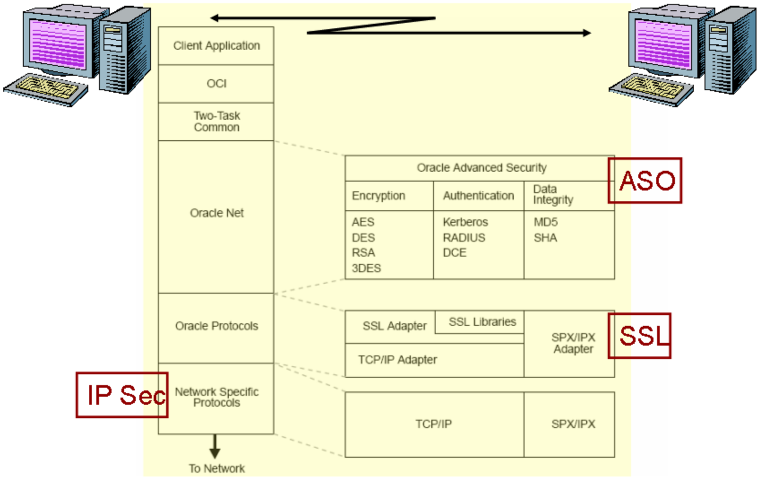

App. C Data Encryption / Network Security: Sybase 1/1

- Database Dump/Load may be password protected
- PKI with SSL is available (SSL Plus Library API from Certicom Corp.)
- ASE 15 Encryption Option:
  - no external keys
  - permission based access control
  - no application modification

App. C Data Encr. / Network Security: SQL – Server 1/1

Gliederu

Hierarchical Encryption and Key Management Infrastructure

Data Protection API (created d. setup)

SQL Server	Instance	Service Master Key
SQL Server	Databases	Database Master Key

(further drilldown possible)

network traffic: SSL SQL native Client Library (40 bit or 128 bit)

HVB Information Services Member of UniCredit Group

App. C Data Encryption / Network Security: DB/2 1/1

- on Mainframe: Data Encryption for IMS and DB/2 uses Crypto Hardware (DES, AES) (Table Level)
- UDB: SQL built in Functions (column level)
  Application Modification needed

- Network Security: relies on OS measures for encrypting Network Traffic (ssh Port Forwarding)

App. D: High Availability Concepts: Sybase 1/2

normal Operation

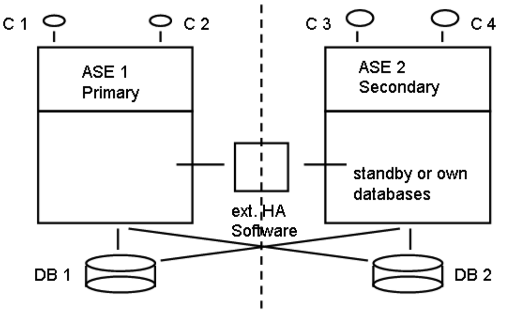

HVB Information Services
Member of UniCredit Group

App. D High Availability Concepts: Sybase 2/2

emergency Operation

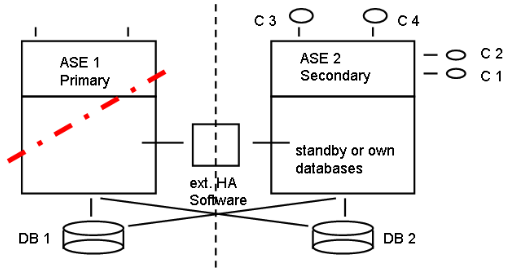

Titel durch Klicken hinzufügen

Gliederung durch Klicken hinzufügen