System Architecture
for
In-Memory Database

Anil Vasudeva
President & Chief Analyst
imex@imexresearch.com
408-268-0800

IMEX RESEARCH.COM

IT Industry Dynamics - Roadmap

IMEX RESEARCH.COM

IT Industry Roadmap
Source: IMEX Research

Analytics – BI

Predictive Analytics - Unstructured Data
From Dashboards Visualization to Prediction Engines using Big Data.

Automation/SDDC

Automatically Maintains Application SLAs
(Self-Configuration, Self-Healing©IMEX, Self-Acctg. Charges etc.)

Cloudization

On-Premises > Private Clouds > Public Clouds
DC to Cloud-Aware Infrast. & Apps. Cascade migration to SPs/Public Clouds.

Virtualization

Pools Resources. Provisions, Optimizes, Monitors
Shuffles Resources to optimize Delivery of various Business Services

Integration/Consolidation

Integrate Physical Infrast./Blades to meet CAPSIMS ©IMEX
Cost, Availability, Performance, Scalability, Inter-operability, Manageability & Security

Standardization

Standard IT Infrastructure- Volume Economics HW/Syst SW
(Servers, Storage, Networking Devices, System Software (OS, MW & Data Mgmt. SW)

IT Infrastructure: DataCenters & Cloud

IMEX RESEARCH.COM

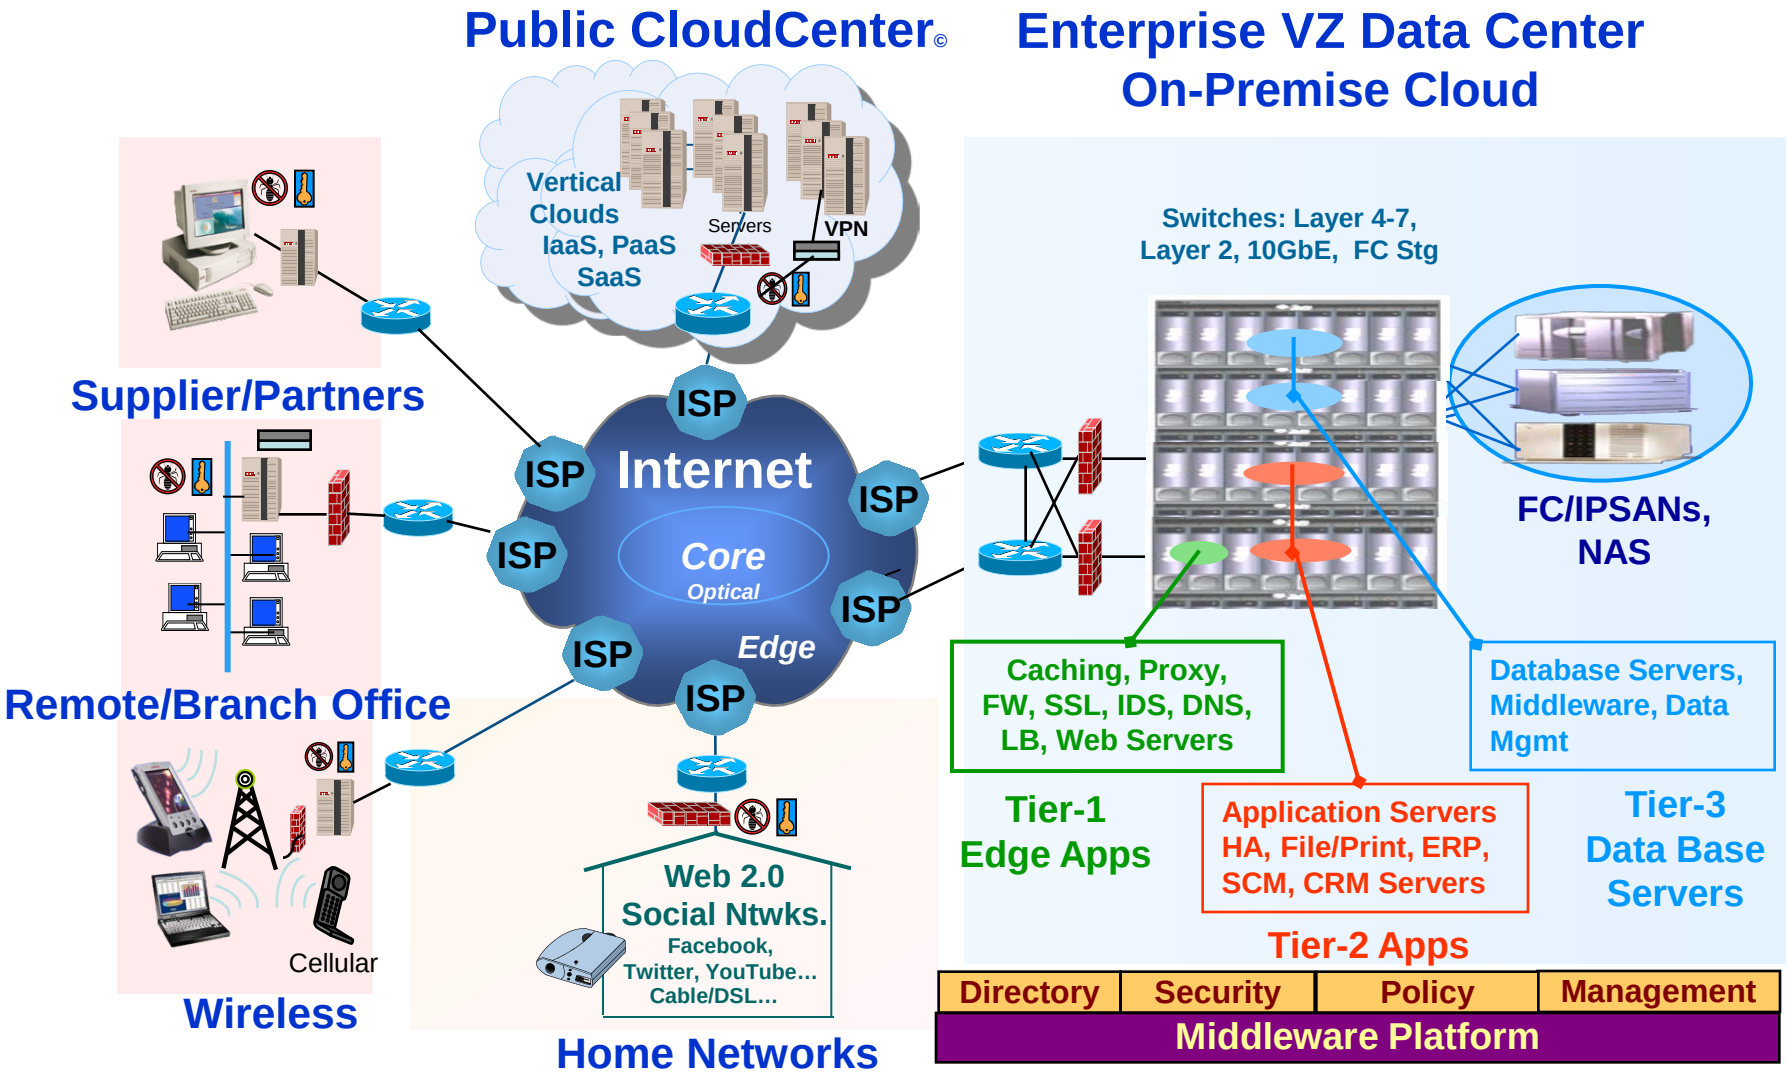

Source:: IMEX Research - Cloud Infrastructure Report ©2009-11

3

IT Industry Dynamics - Roadmap

IMEX RESEARCH.COM

IT Industry Roadmap
Source: IMEX Research

Analytics – BI

Predictive Analytics - Unstructured Data
From Dashboards Visualization to Prediction Engines using Big Data.

Automation/SDDC

Automatically Maintains Application SLAs
(Self-Configuration, Self-Healing©IMEX, Self-Acctg. Charges etc.)

Cloudization

On-Premises > Private Clouds > Public Clouds
DC to Cloud-Aware Infrast. & Apps. Cascade migration to SPs/Public Clouds.

Virtualization

Pools Resources. Provisions, Optimizes, Monitors
Shuffles Resources to optimize Delivery of various Business Services

Integration/Consolidation

Integrate Physical Infrast./Blades to meet CAPSIMS ©IMEX
Cost, Availability, Performance, Scalability, Inter-operability, Manageability & Security

Standardization

Standard IT Infrastructure- Volume Economics HW/Syst SW
(Servers, Storage, Networking Devices, System Software (OS, MW & Data Mgmt. SW)

Workloads: Mapped on Infrastructure Metrics

IMEX RESEARCH.COM

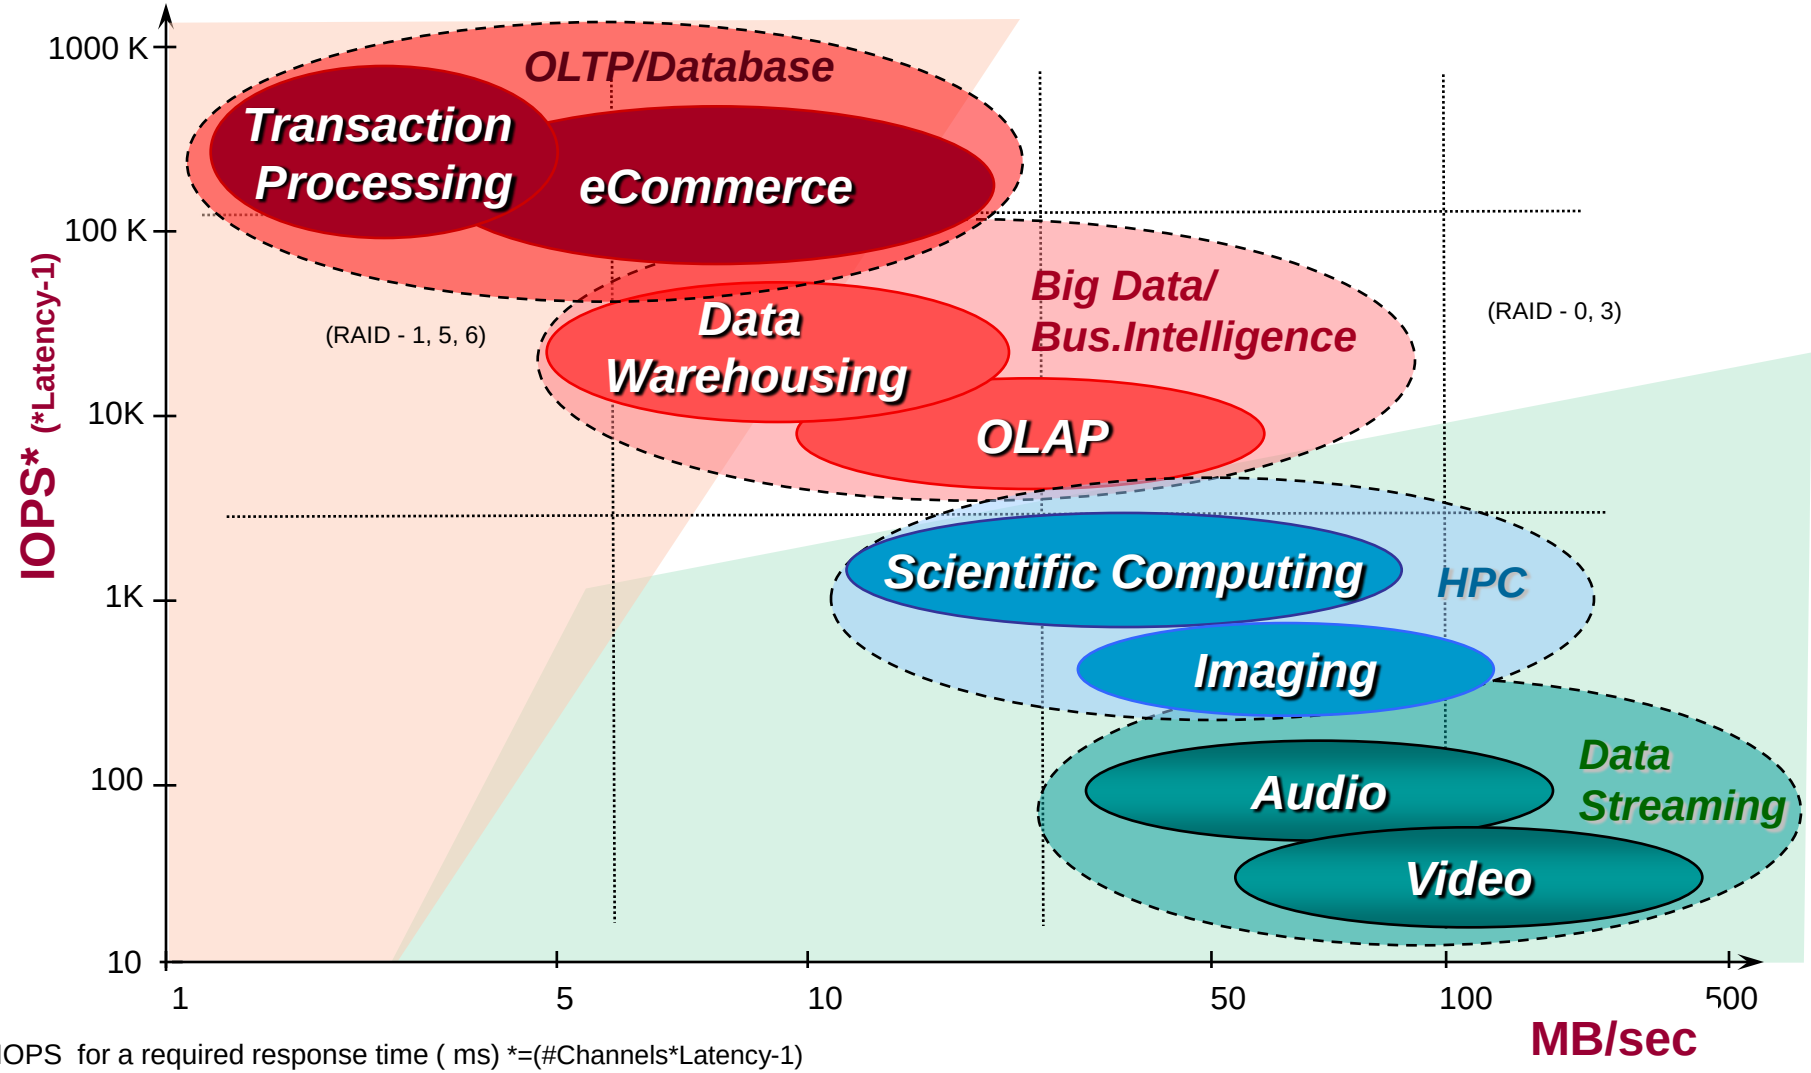

*IOPS for a required response time (ms) *=(#Channels*Latency-1)

Workloads need Infrastructure Optimized for Cost, Availability, Performance …

Workloads: I/O Characteristics

IMEX RESEARCH.COM

Storage performance, management and costs are big issues in running Databases

- **Data Warehousing Workloads are I/O intensive**
  - Predominantly read based with low hit ratios on buffer pools
  - High concurrent sequential and random read levels
    - ✓ Sequential Reads requires high I/O Bandwidth (MB/sec)
    - ✓ Random Reads require high IOPS
  - Write rates driven by life cycle management and sort operations
- **OLTP Workloads are strongly random I/O intensive**
  - Random I/O is more dominant
    - ✓ Read/write ratios of 80/20 are most common but can be 50/50
    - ✓ Difficult to build out test systems with sufficient I/O characteristics
- **Batch Workloads (Hadoop) are more write intensive**
  - Sequential Writes requires high I/O Bandwidth (MB/sec)
- **Backup & Recovery times are critical for these workloads**
  - Backup operations drive high level of sequential IO
  - Recovery operation drives high levels of random I/O

Source: IMEX Research SSD Industry Report ©2011

Driver : Need Real Time Analytics

IMEX RESEARCH.COM

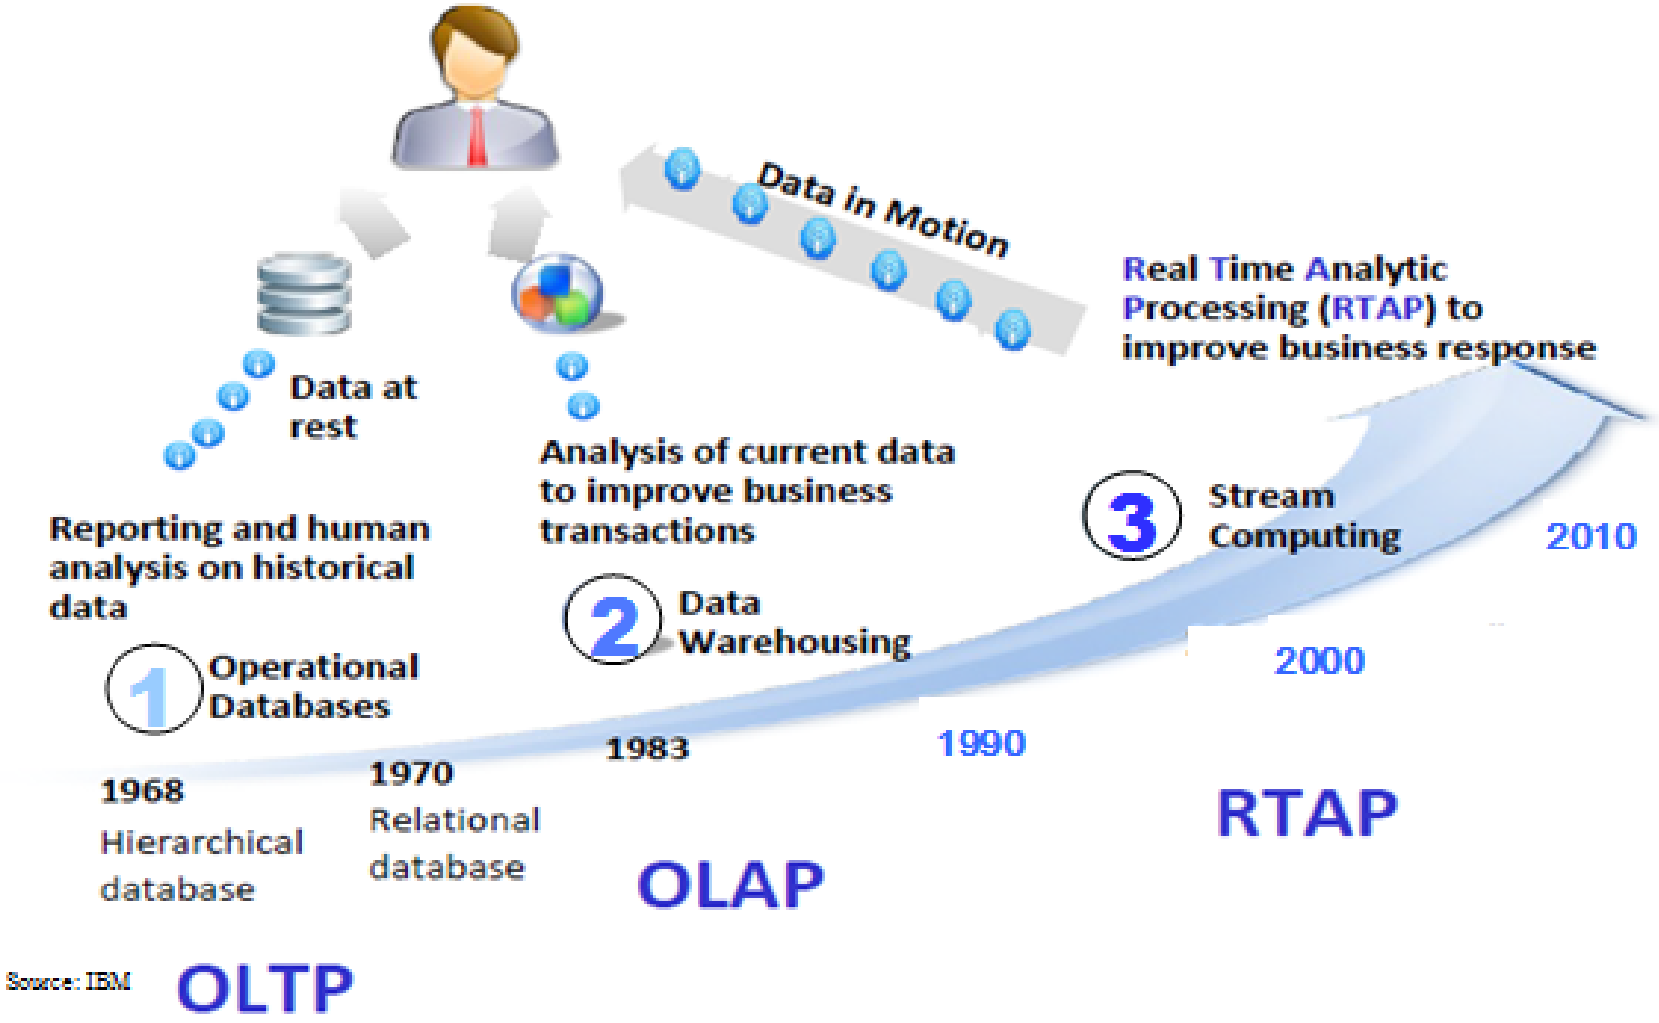

Issue: Server to Storage I/O Gap

IMEX RESEARCH.COM

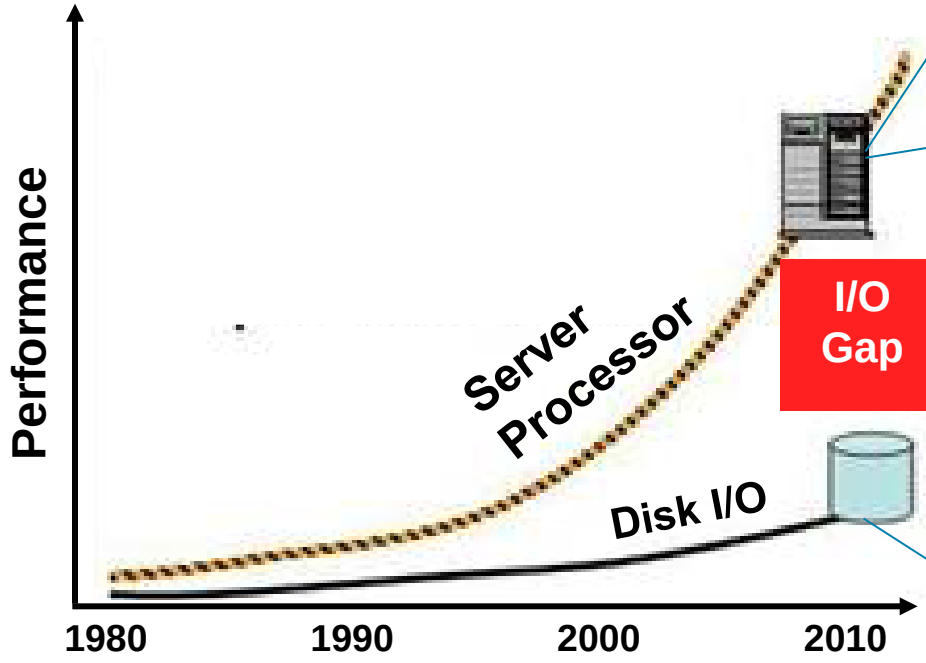

Memory - MultiSlots

Connect to Storage as Direct Attached Storage (Internal Storage)

I/O Gap

For Each Disk Operation,
Millions of CPU Opns. or Thousands of Memory Opns. can be accomplished

PCIe used for HBAs to connect to (External Shared Storage) via Storage Switches/Fabric as SAN/NAS on HDDs front ended by DRAM Cache creating a gap of 100,000x latency gap

L1 cache reference	0.5 ns
Branch mispredict	5 ns
L2 cache reference	7 ns
Mutex lock/unlock	25 ns
Main memory reference	100 ns
Compress 1K bytes with Zippy	3,000 ns
Send 2K bytes over 1 Gbps network	20,000 ns
Read 1 MB sequentially from memory	250,000 ns
Round trip within same datacenter	500,000 ns
Disk seek	10,000,000 ns
Read 1 MB sequentially from disk	20,000,000 ns
Send packet CA->Netherlands->CA	150,000,000 ns

A 7.2K/15k rpm HDD can do 100/140 IOPS

Solution: SSDs Improving DB Query Responses

IMEX RESEARCH.COM

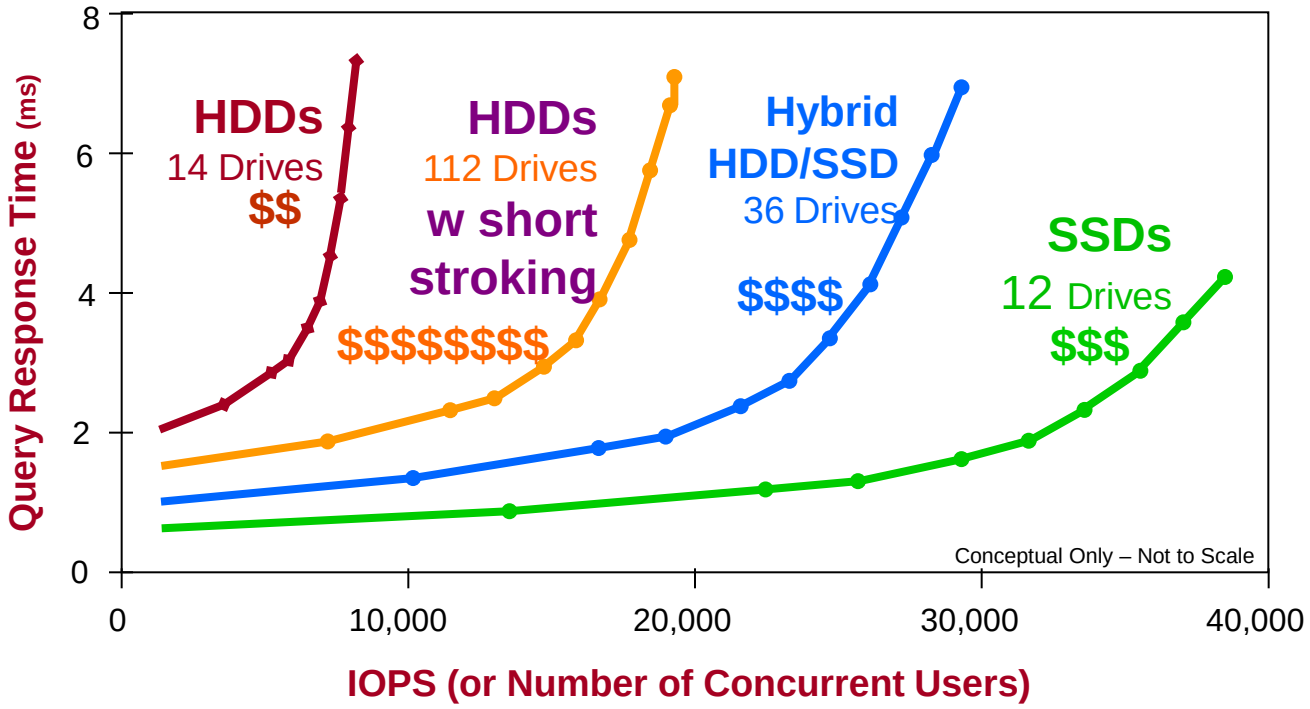

For a targeted query response time in DB & OLTP applications, many more concurrent users can be added cost-effectively when using SSDs or SSD + HDDs storage vs. adding more HDDs or short-stroking HDDs

Source: IMEX Research SSD Industry Report ©2011

Industry Trends: Impact on Infrastructure

IMEX RESEARCH.COM

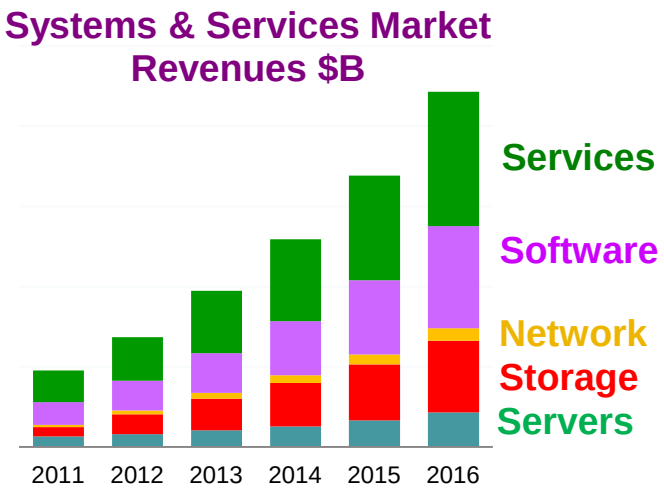

Industry Dynamics	Impact on Infrastructure
Serviceable PCIe SSDs	Same Connector - PCIe+SAS/SATA
Lower Cost DRAM	NVMe based Flash Memory
BYOD / Boot Storms	Client Images on Servers
Big Data/RealTm Analytics	PCIe Servers based SSD
Server/Stg. Price/Perf.	New Storage Class Memory
Cloud Computing	New Protocols / REST, HTTP..
Multicores	Flash for Concurrent Multitasking
Power Efficiency	Green Memory
Virtualization	New Infrastrecture for Multi-VMs
Scale Out Clustering	Distributed Memory Architecture
Dense Blades	Fast, Low Power Memory
64 bit Computing	Larger Size Memories

Workloads need Infrastructure - Optimized for Cost, Availability, Performance …

Solution: SSDs Filling Price/Perf Gap

IMEX RESEARCH.COM

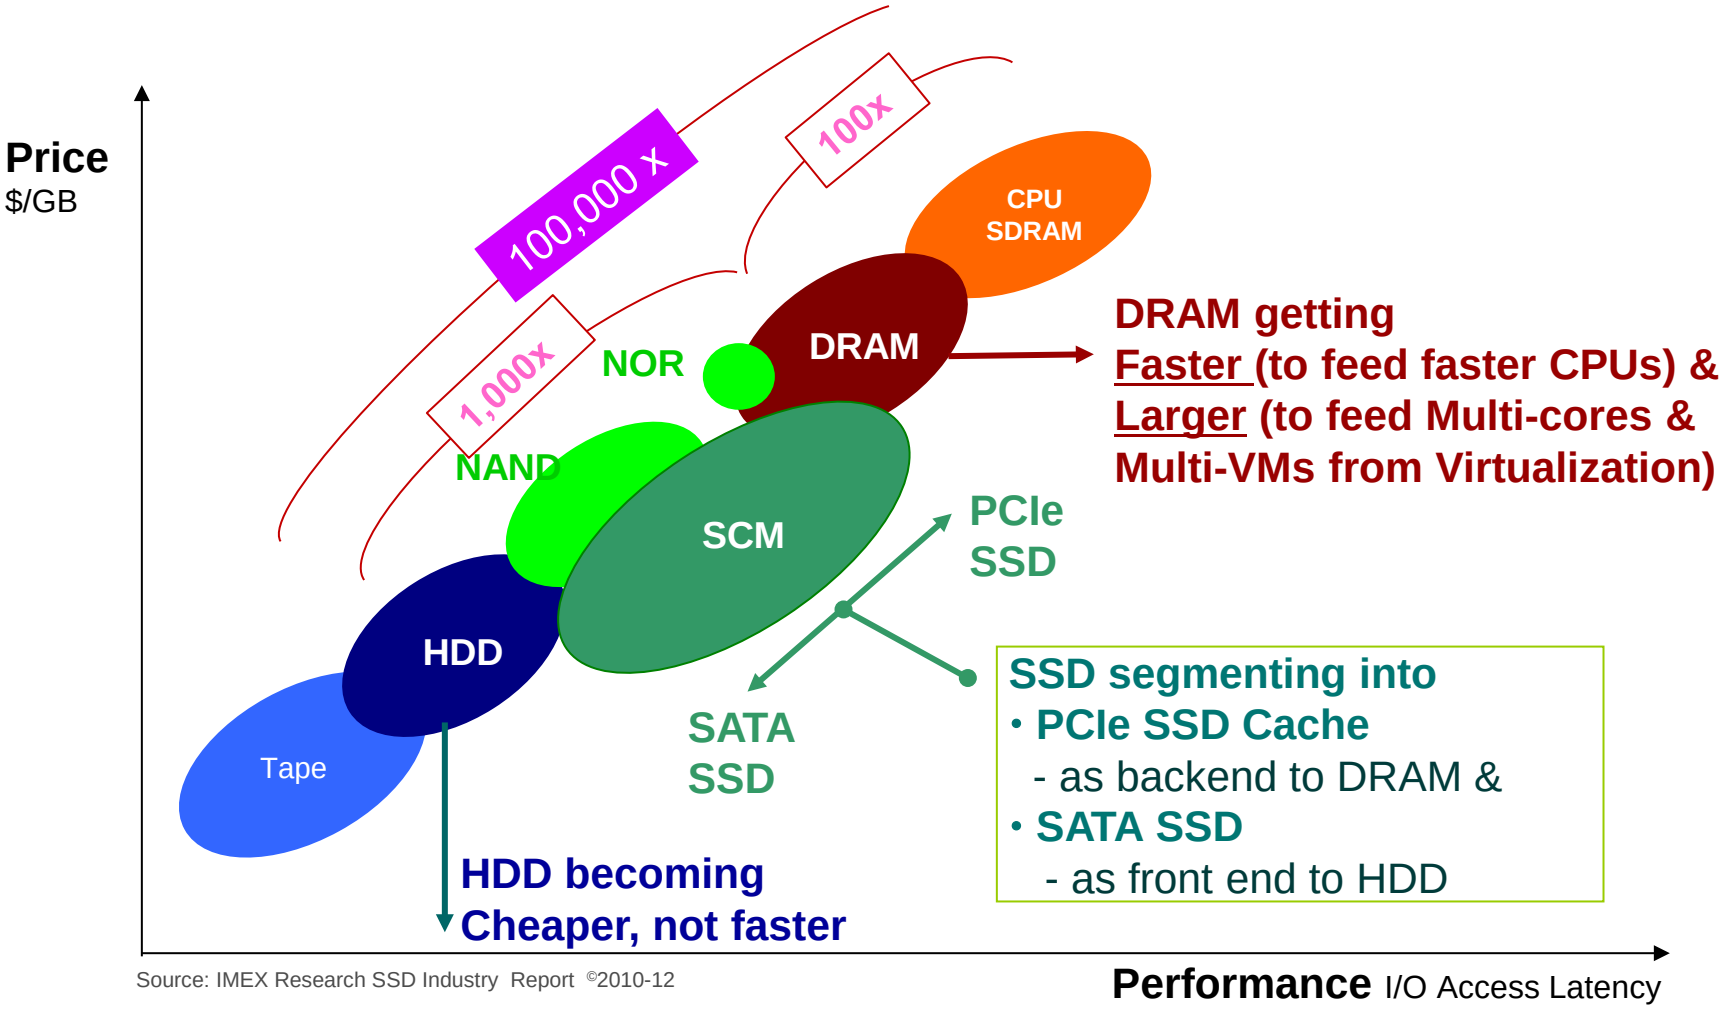

Best Opportunity to fill the gap is for storage to be close to Server CPU.

11

Innovations Roadmap: HW Technologies

IMEX RESEARCH.COM

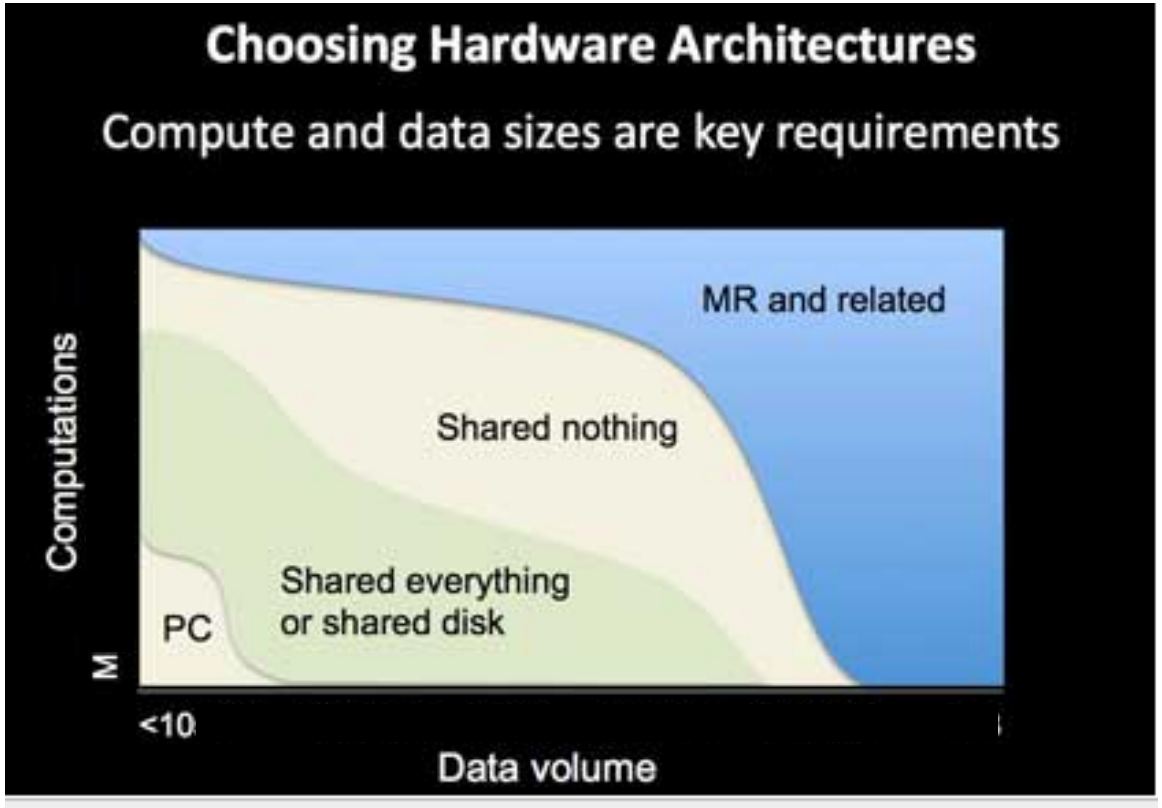

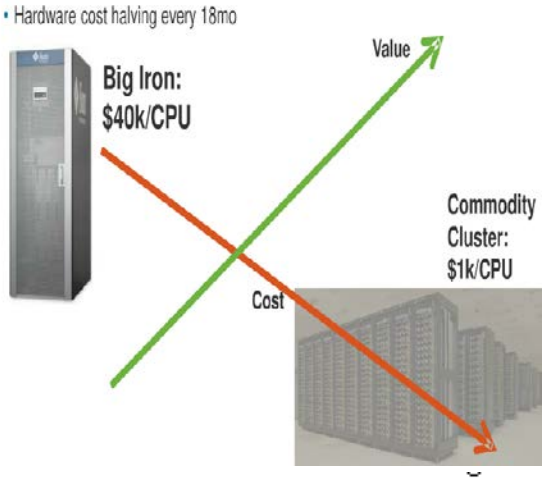

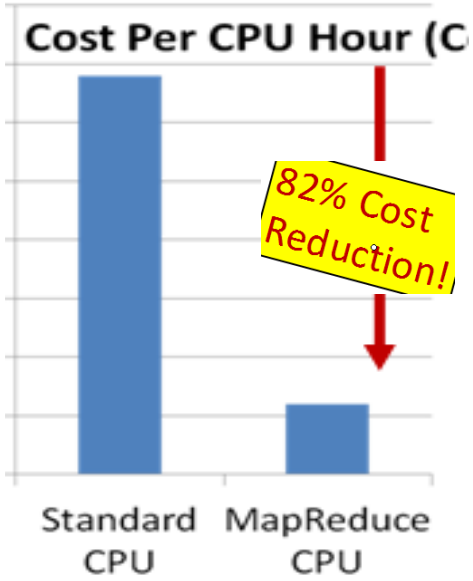

Innovations Roadmap – DB SW Technologies

IMEX RESEARCH.COM

OLTP Database Innovation Progress

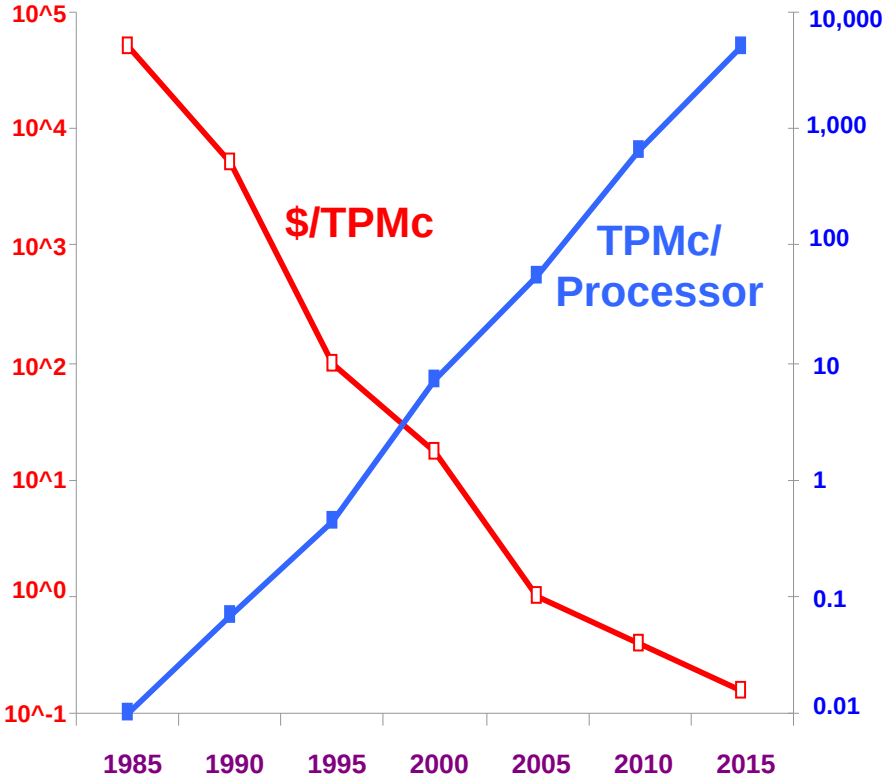

EDW/Big Data Database Innovation Progress

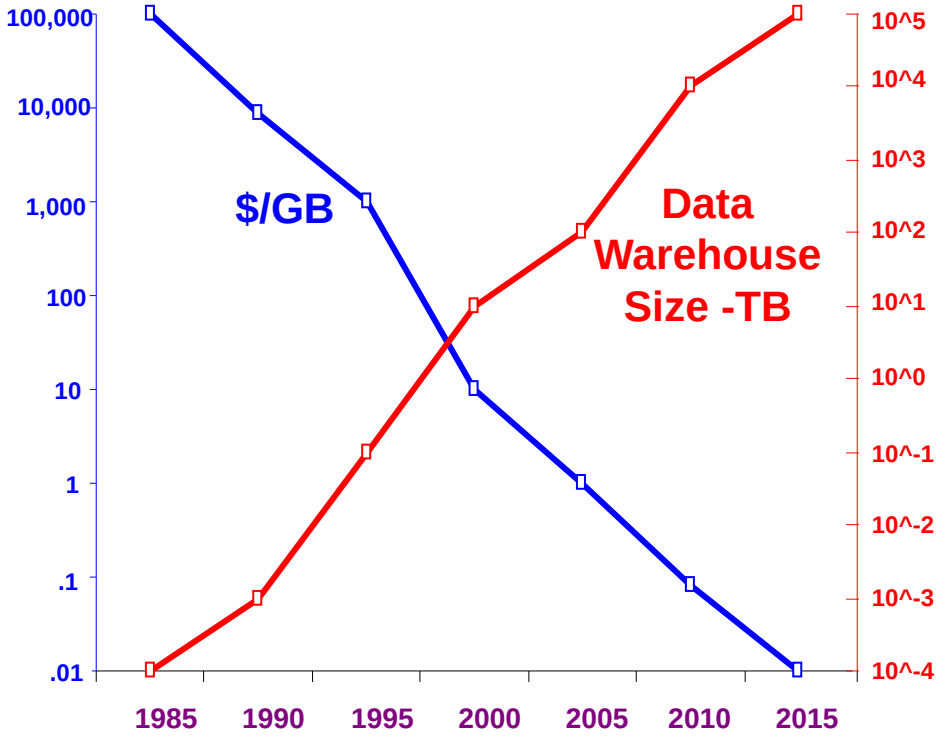

In-Memory Computing

IMEX RESEARCH.COM

Advances in Hardware

Multi-Core Architecture
(8 x 10core CPU per blade)

Parallel scaling across blades

One blade ~$50.000 = 1
Enterprise Class Server

64 bit address space – 2TB in current server boards

25GB/s data throughput

Cost-performance ratio rapidly declining

Advances in Software

Row and Column Store

Compression

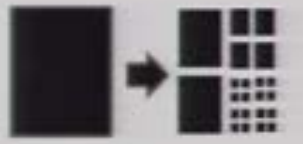

Partitioning

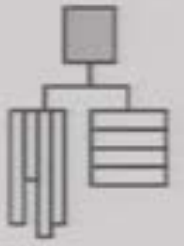

No Aggregate Tables

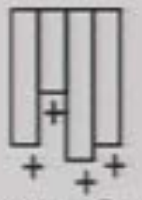

Insert Only

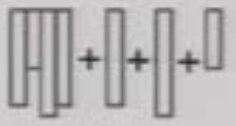

On-the-fly extensibility

Technology: Disk vs InMem DB Architecture

IMEX RESEARCH.COM

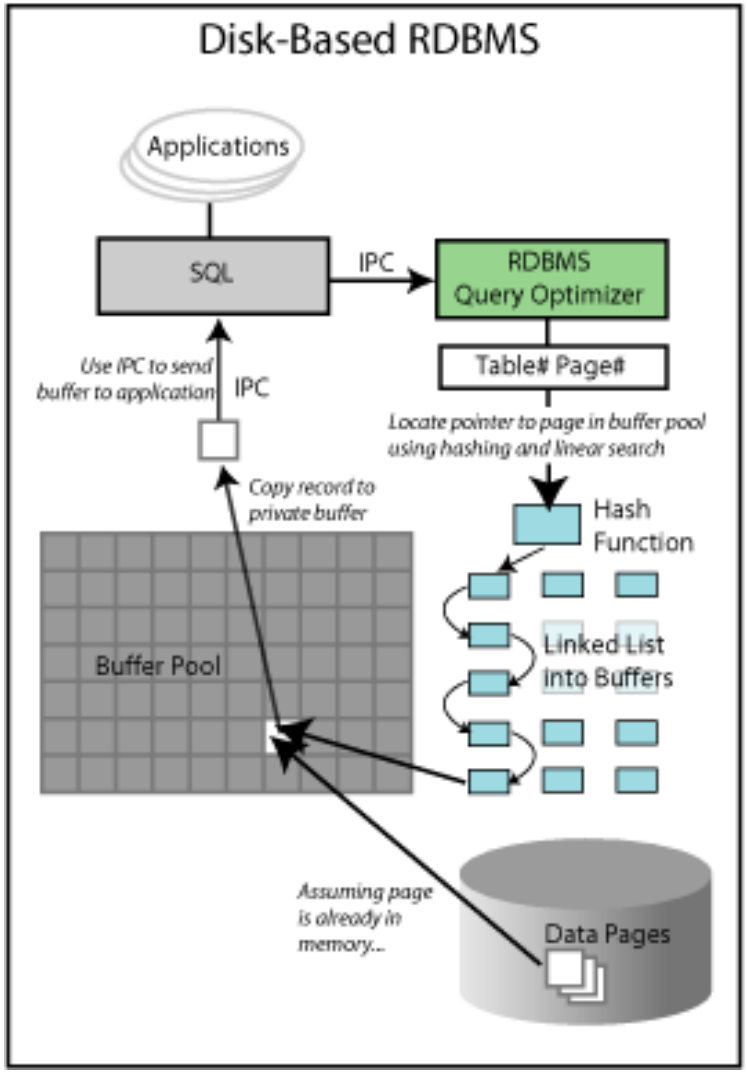

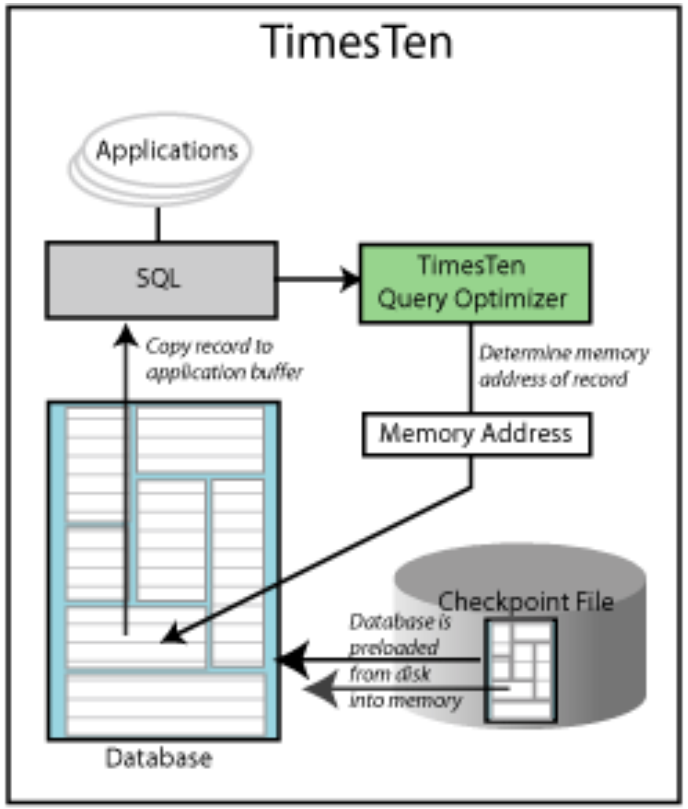

Technology: RDBMS vs. In-Memory DBMS

IMEX RESEARCH.COM

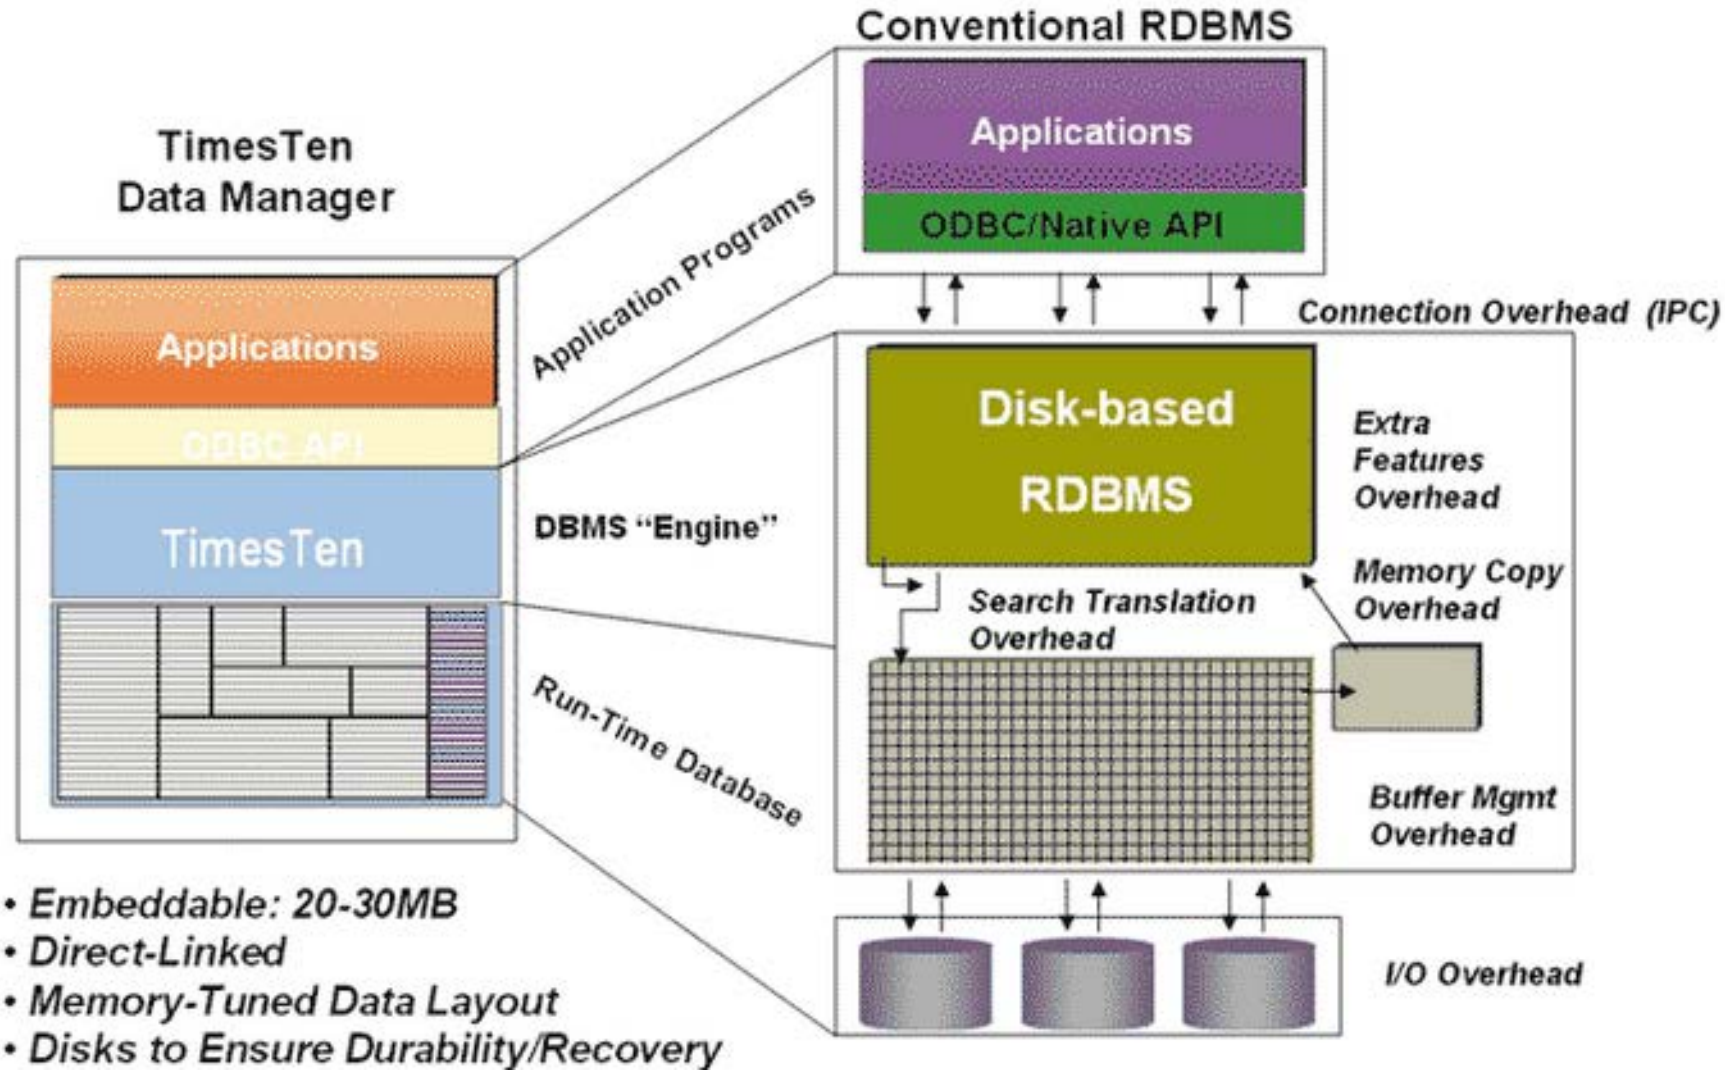

- *Embeddable: 20-30MB*
- *Direct-Linked*
- *Memory-Tuned Data Layout*
- *Disks to Ensure Durability/Recovery*

Technology:Legacy vs In-Memory Computing

IMEX RESEARCH.COM

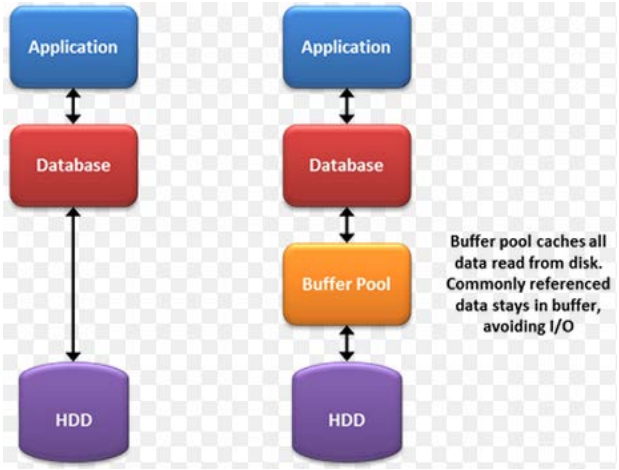

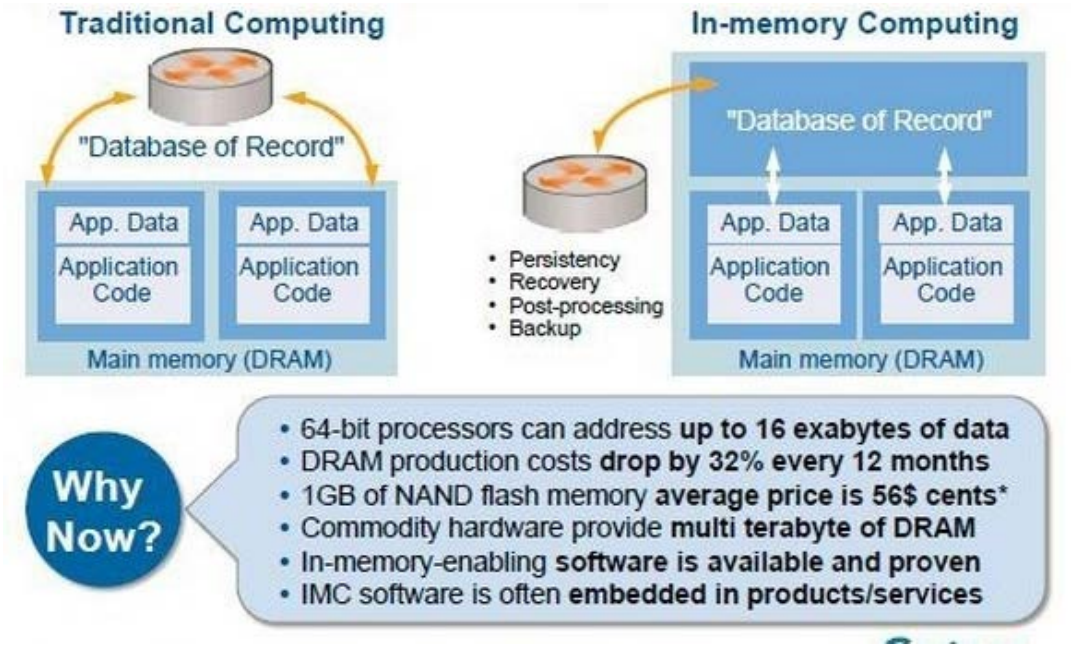

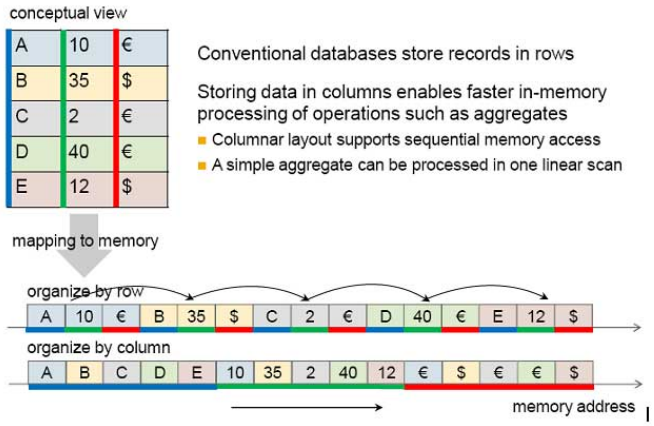

Oracle vs SAP vs IBM DB

IMEX RESEARCH.COM

Exadata Evolution

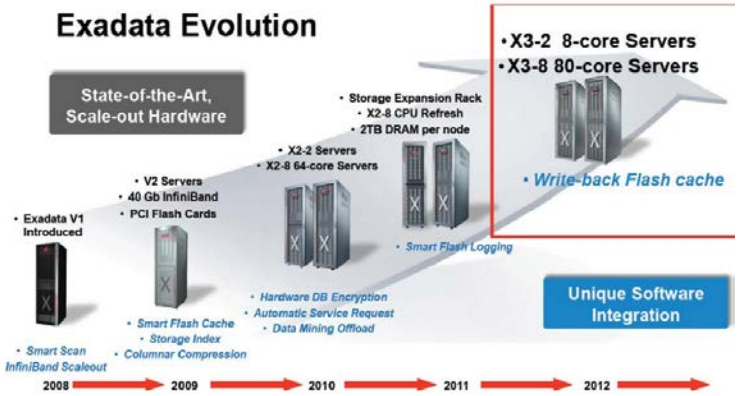

Exadata Architecture

Complete Database platform using standard servers for Compute and Storage

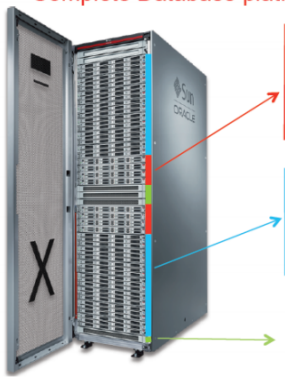

Scale-Out Database Servers
- 2-socket or 8-socket Xeon database servers
- Oracle Database, ASM, RAC; Linux or Solaris
- Standard Ethernet to data center

Scale-Out Intelligent Storage Servers
- 2-socket storage servers, Exadata Storage Software
- Up to 500 terabytes disk per rack
- 56 PCI Flash memory cards per rack

InfiniBand Network
- Unified internal connectivity (40 Gb/sec)

Exadata Hybrid Columnar Compression
Highest Capacity, Lowest Cost

Hardware Generation Advances

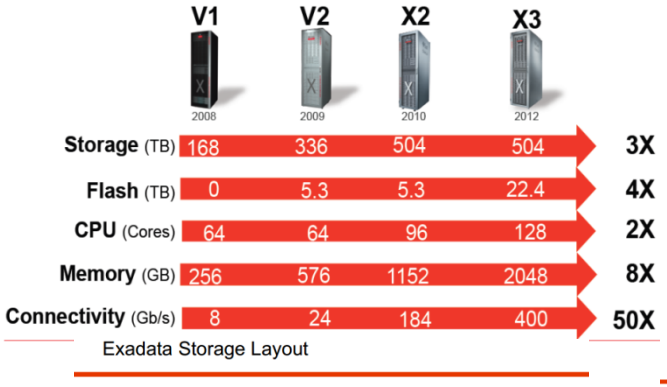

Exadata Storage Layout

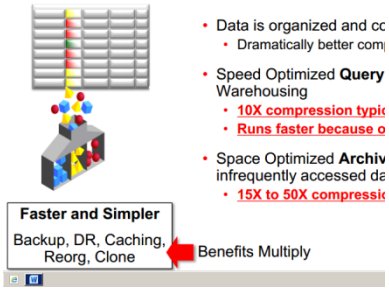

Exadata Smart Flash Cache
Extreme Performance OLTP

Exadata Summary

- Best for OLTP
  - Smart Flash Cache
  - 1 Million I/Os per Second
- Best for Warehousing
  - Intelligent Scale-Out storage
    - 10x faster queries
  - 10x Data Compression
- Best for Consolidation
  - Terabytes of Memory
  - Mix OLTP, DW, batch, reporting in single machine

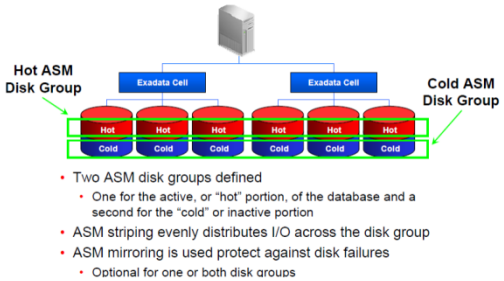

- Two ASM disk groups defined
  - One for the active, or "hot" portion, of the database and a second for the "cold" or inactive portion
- ASM striping evenly distributes I/O across the disk group
- ASM mirroring is used protect against disk failures
  - Optional for one or both disk groups

- Exadata has 5 TB / 22.3 TB of flash
  - 56 Flash PCI cards avoid disk controller bottlenecks
- Intelligently manages flash
  - Smart Flash Cache holds hot data
  - Gives speed of flash, cost of disk
- Exadata flash cache achieves:
  - Over 1 million IO/sec from SQL (8K)
  - Sub-millisecond response times
  - 50 GB/sec query throughput

Competition: Oracle DB Architecture

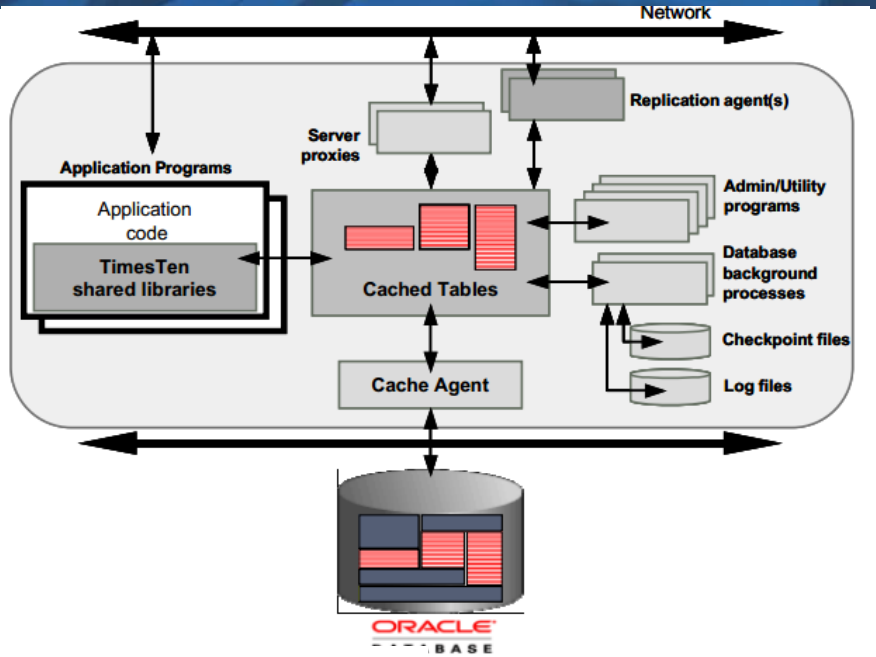

Exadata Smart Flash Cache – Write Operation

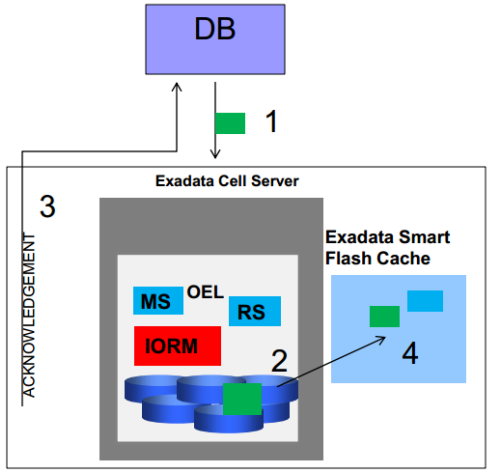

Exadata Smart Flash Cache – Read

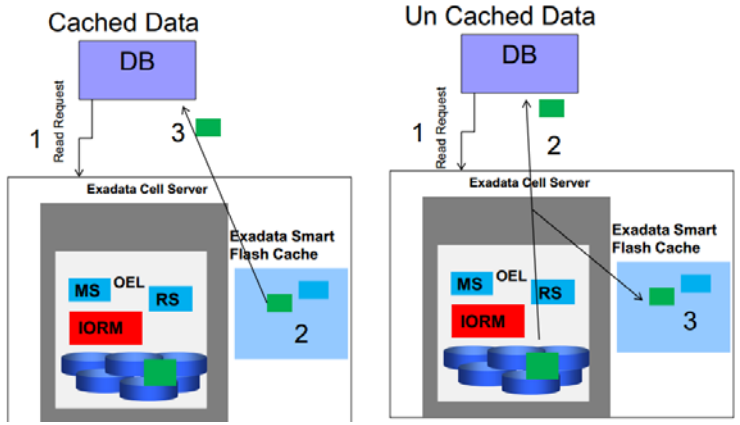

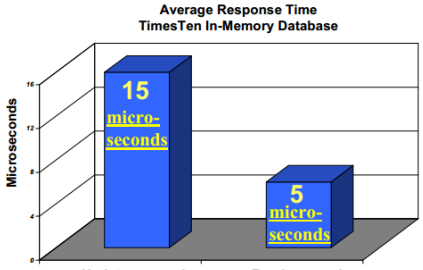

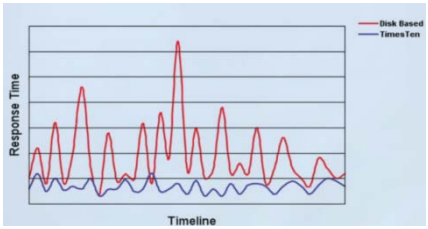

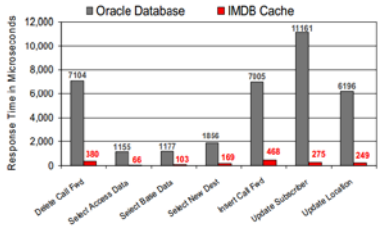

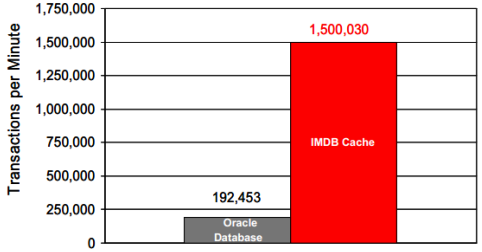

Competition: SAP/HANA (Multi-Applications)

A Converged DB System

- **In-memory database combining transactional data processing, analytical data processing, and application logic processing functionality in memory**.
- **A full DBMS with a standard SQL interface, high availability, transactional isolation and recovery (ACID properties)**
- **both row-based and column-based stores within the same engine** (row-based storage is good for transactional applications, while column-based storage is better for reports and analytics. Column-based storage compresses the better too.)
- **massively parallel execution using multicore processors**, SAP HANA optimizes the SQL which scales well with the number of cores. Aggregation operations by spawning a number of threads that act in parallel, each of which has equal access to the data resident on the memory on that node
- Additional functions - **freestyle search** (as SQL extensions). BI applications using MDX for Microsoft Excel & Consumer Services plus internal I/F for BusinessObjects
- **prepackaged algorithms in the predictive analysis library** of SAP HANA to perform advanced statistical calculations
- **built-in text support**, from its predecessor BI Accelerator that was based on the TREX search engine and Inxight functionality integrated into HANA text functions.

Competition: SAP/HANA (Multi-Applications)

IMEX RESEARCH.COM

- supports distribution across hosts, where large tables may be partitioned to be processed in parallel. DB “engine” of the SAP HANA Analytics appliance as well
- HANA’s combination of a row and column store is fundamentally different from any other database engine on the market today, which allows it to perform OLTP and analytics processing in memory, at the same time.
- Avoids CPU waiting info from Memory through its unique CPU-cache-aware algorithms and data structures that there is as much useful data in the CPU caches as possible,.
- it uses late materialization to decompress columnar structures as late as possible, or to run operations directly on the compressed data
- also sold as an appliance on Intel Xeon CPUs leveraging insights into Intel’s HyperThreading, Turbo Boost and Threading Building Blocks
- High Performance Analytic Appliance can perform large-scale data analyses on 500 billion records in less than a minute, taking analytics to an entirely new dimension
- represents a complete data warehouse in RAM, and as a result, much accelerated real-time analytics.
- .

Technology: In-Memory Computing

IMEX RESEARCH.COM

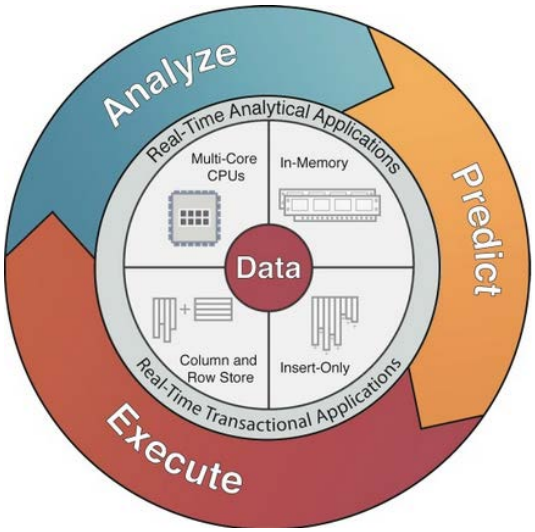

conceptual view

A	10	€
B	35	$
C	2	€
D	40	€
E	12	$

Conventional databases store records in rows

Storing data in columns enables faster in-memory processing of operations such as aggregates

- Columnar layout supports sequential memory access
- A simple aggregate can be processed in one linear scan

mapping to memory

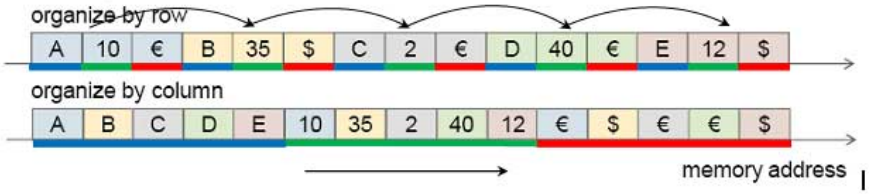

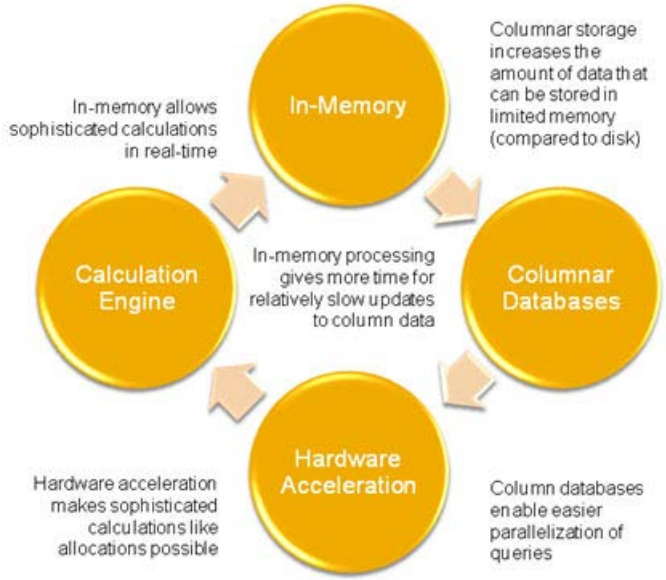

single HANA platform for all business processing

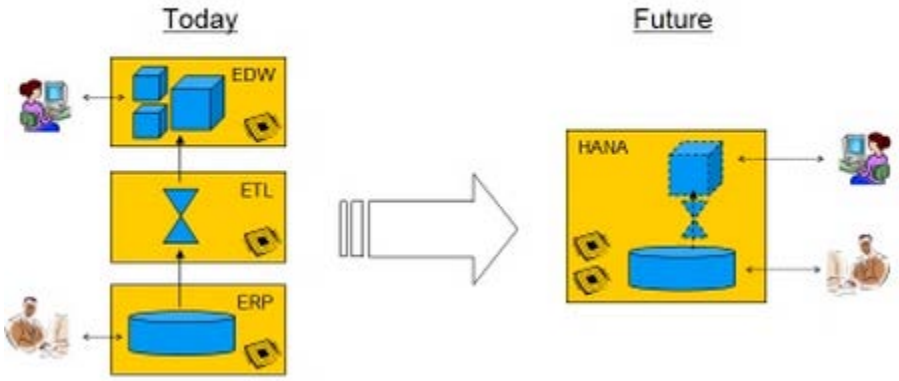

Trends: In-Memory Computing Adoption

IMEX RESEARCH.COM

SUSE HANA Certified Hardware

Pre-load SUSE Linux Enterprise Server for SAP Applications

IBM, HP, 1&1, amazon web services, FUJITSU, FUJITSU, DELL, IBM, Windows Azure, CISCO, HITACHI, HUAWEI, NEC

OS used on Top 500 Supercomputers

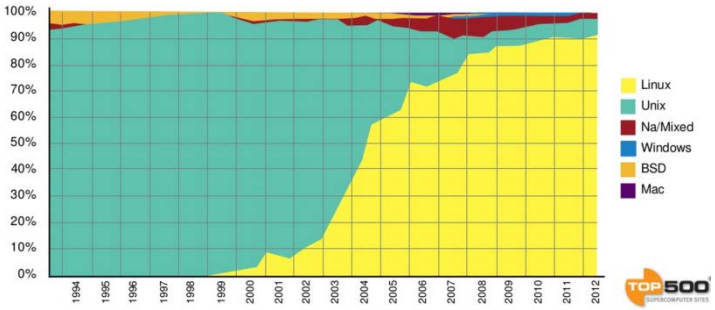

Unix Vs. Linux Trend Lines

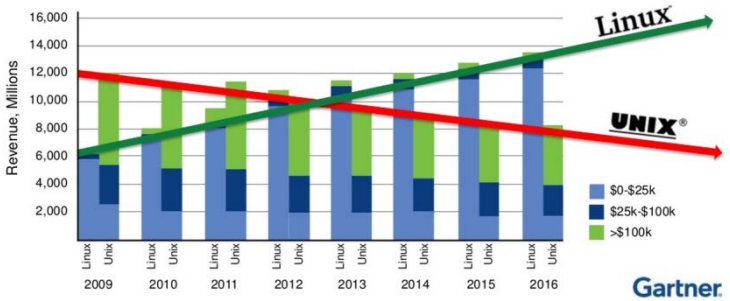

Big Data Technologies Planned

Organizations plan to use a blend of approaches

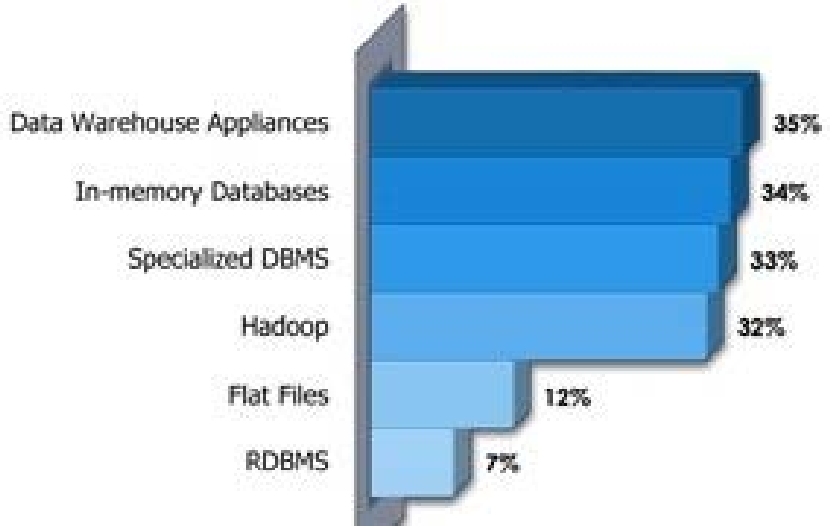

Trends: In-Memory DB Computing

IMEX RESEARCH.COM

Primary DB for Each Application

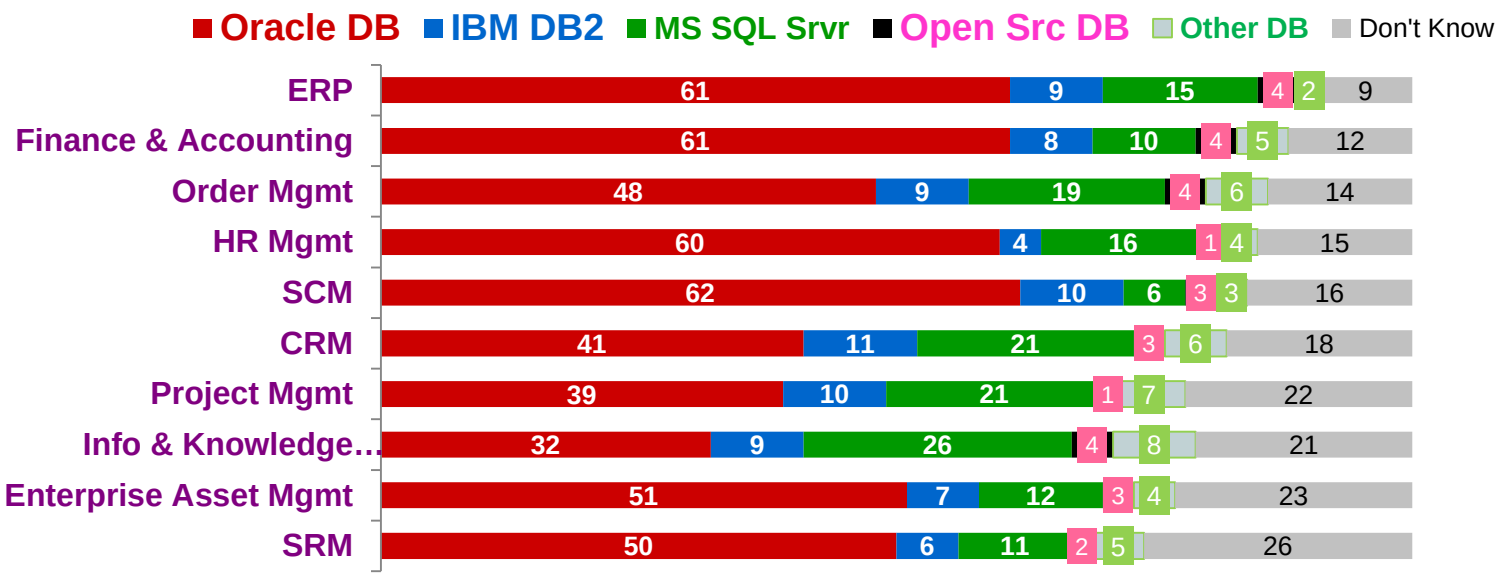

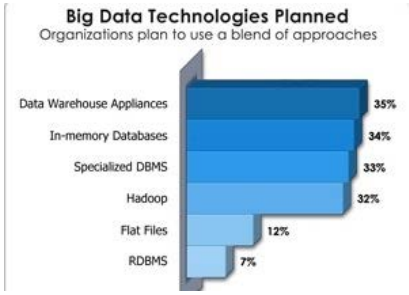

System Architecture
for
In-Memory Database

Anil Vasudeva
President & Chief Analyst
imex@imexresearch.com
408-268-0800

IMEX RESEARCH.COM