

Recent Innovations in Data Storage Technologies

Dr Roger MacNicol
Software Architect



Safe Harbor Statement

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

Recent Database Innovation at Oracle

In-Memory Technologies

Exadata

Seamless integration of Big Data

Fast Ingest for IoT

Autonomous Database

Hardware & Software in Silicon

Row Format Databases vs. Column Format Databases

Row



- **Transactions** run faster on row format
 - Example: Insert or query a sales order
 - Fast processing few rows, many columns


Column



- **Analytics** run faster on column format
 - Example : Report on sales totals by region
 - Fast accessing few columns, many rows

Until Now Must Choose One Format and Suffer Tradeoffs

In-Memory: Dual Format Database



- **BOTH** row and column formats for same table
- Simultaneously active and transactionally consistent
- OLTP uses proven row format
- Analytics & reporting use new in-memory Column format
- On demand loading
- No change to application!


In-Memory: Improves all Aspects of Analytical Queries

Data Scans



- Scan and Filter only the needed Columns
- Storage Indexes
- Vector Instructions
- In-Memory Expressions
- OSON: Compiled JSON

Joins



- Convert Star Joins into 10X Faster Column Scans
- Search large table for values that match small table
- Join Groups

In-Memory Aggregation



- Create In-Memory Report Outline that is Populated during Fast Scan
- Runs Reports Instantly

Exadata Vision

Dramatically Better Platform for All Database Workloads



- **Ideal Database Hardware** - Scale-out, database optimized compute, networking, and storage for fastest performance and lowest costs
- **Smart System Software** – specialized algorithms vastly improve all aspects of database processing: **OLTP, Analytics, Consolidation**
- **Full-Stack Integration** – Database-to-disk optimization, automation, testing, patching, and support to reduce operational costs

Identical On-Premises and Oracle Public Cloud

**Exadata Cloud
Service**

Exadata Smart System Software Highlights

Smart Analytics

- NVMe Flash bandwidth exceeds network
- => Move **queries to storage**, not storage to queries
- **100X** faster analytics



Smart OLTP

- **Special InfiniBand protocols and fastest PCIe flash** enable highest speed, lowest latency OLTP
- OLTP Storage Cache



Smart Storage

- Database-aware **Flash Caching** gives speed of PCI flash with capacity of disk
- **Hybrid Columnar Compression** reduces space usage by **10X**

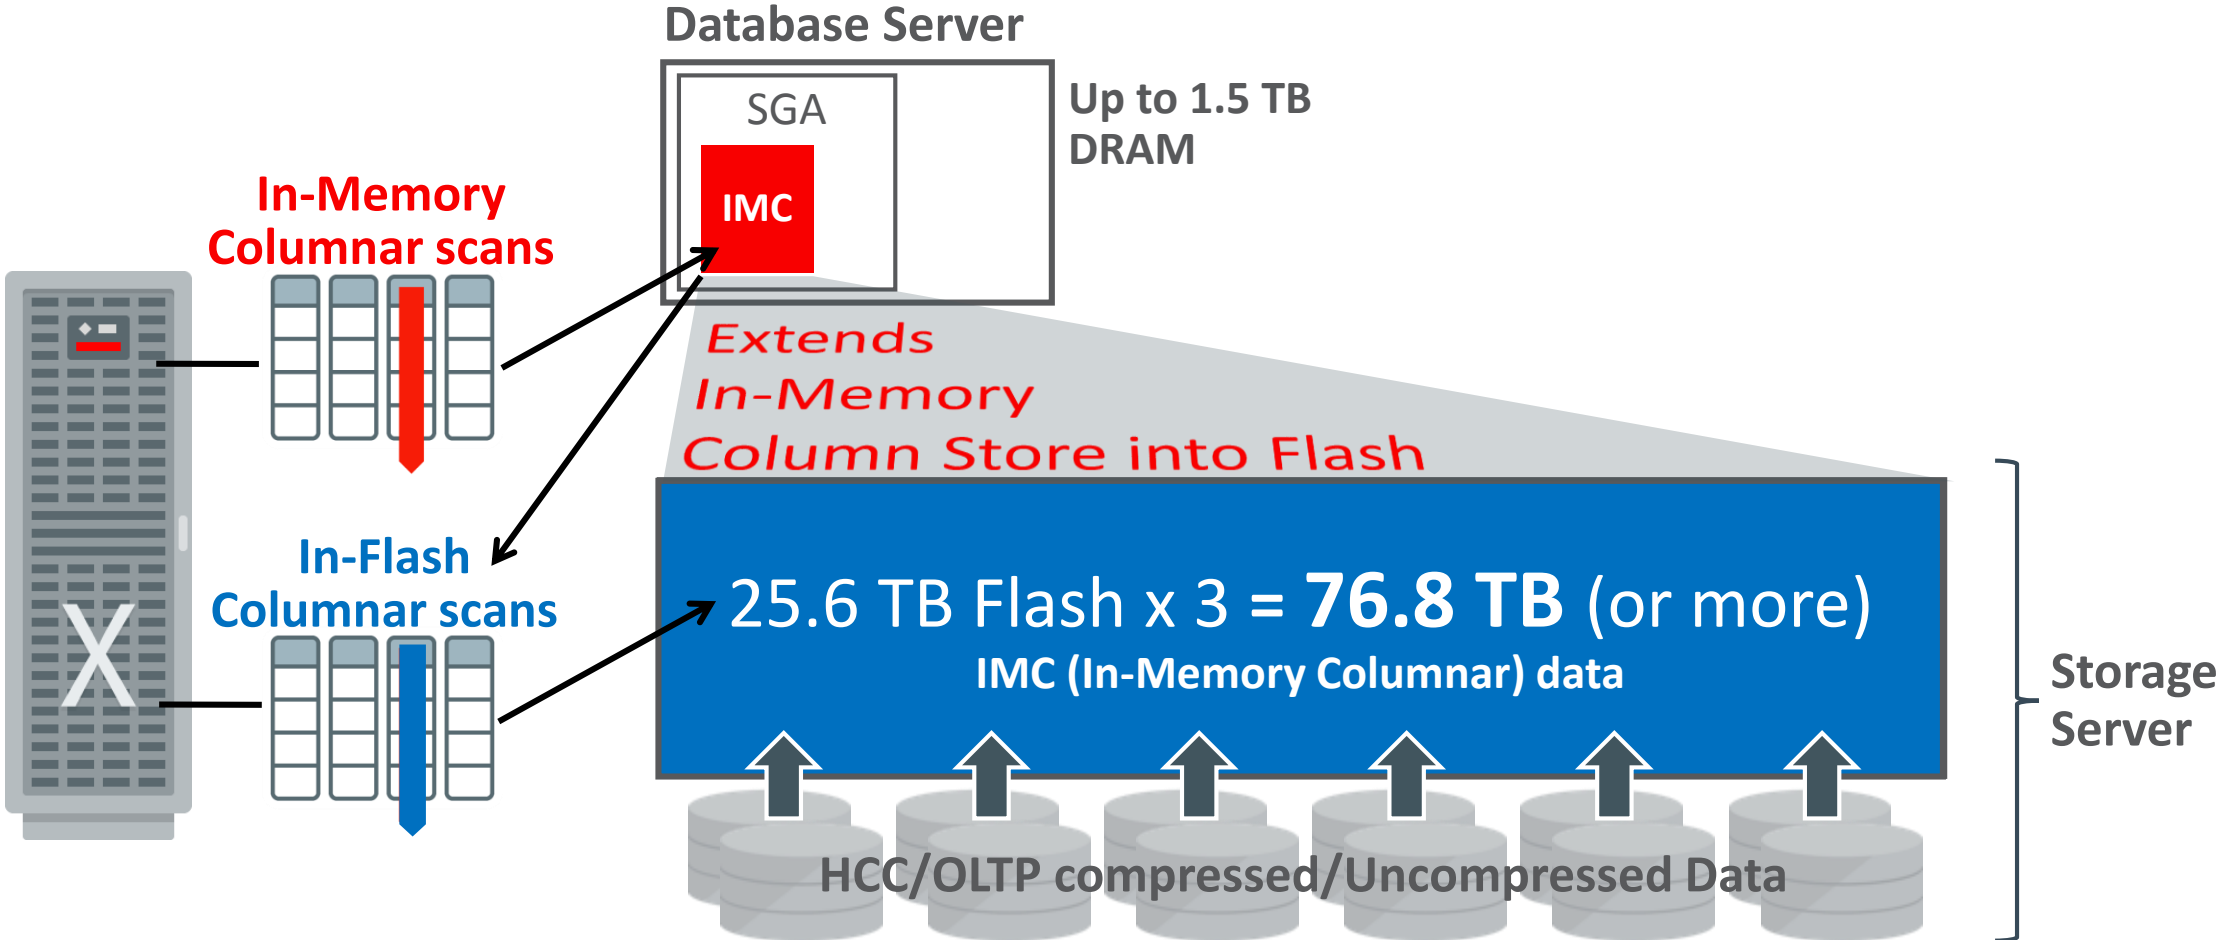
Smart Consolidation

- **Workload prioritization** of CPU, Net, and I/O enables **4X** more DB consolidation. Pluggable Databases.



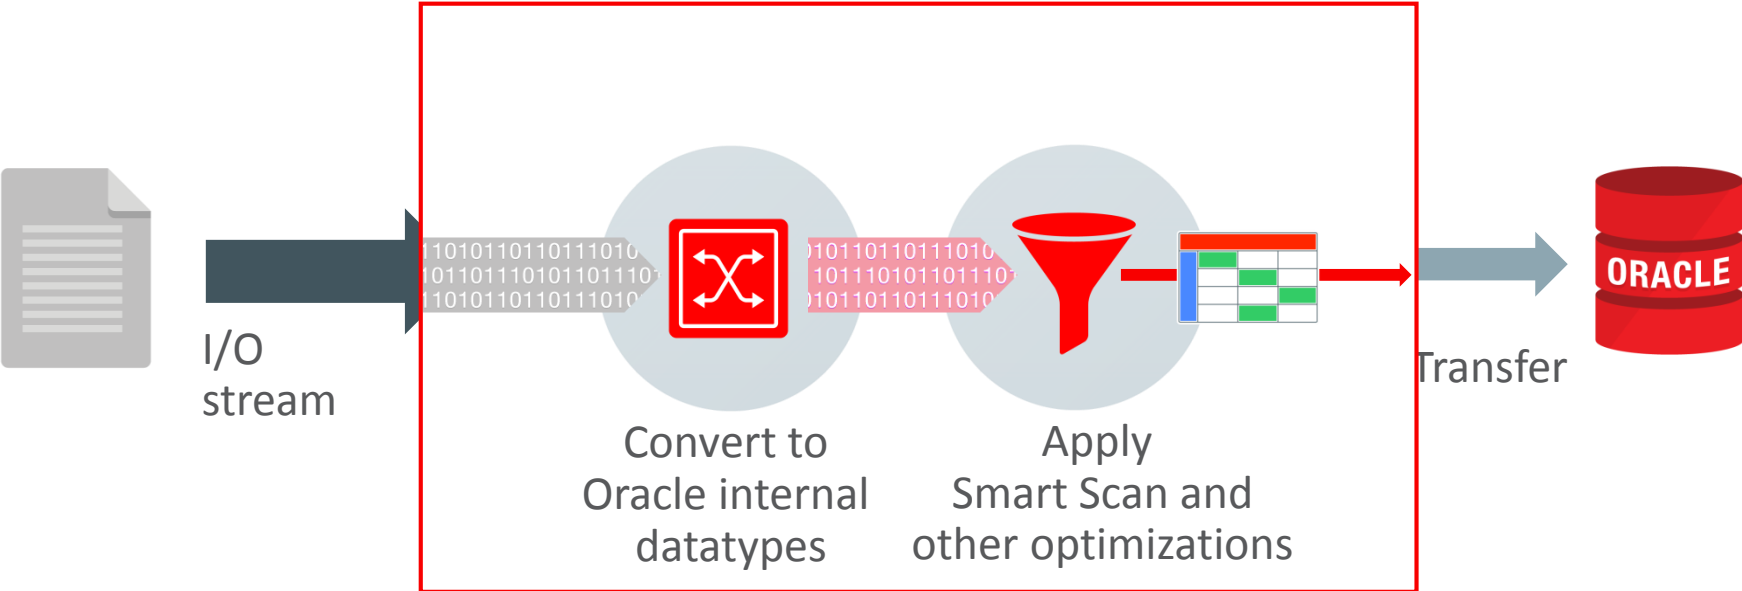
Exadata: In-Memory Performance on Larger Datasets

Columnar Flash Cache



Exadata Paradigm for Big Data Storage

Big Data SQL



















Autonomous Database
















Free the DBA to do more interesting work!

- Automatic In-Memory Management
 - Track heat and page tables or partitions into or out of In-Memory Area
- Automatically tier cooling data to more highly compressed formats
- Automatically tier cold data to cheaper storage
- Automatically manage NVMe for most efficient elimination of bottlenecks
- Extend Database OLTP Cache to tier to storage layer memory

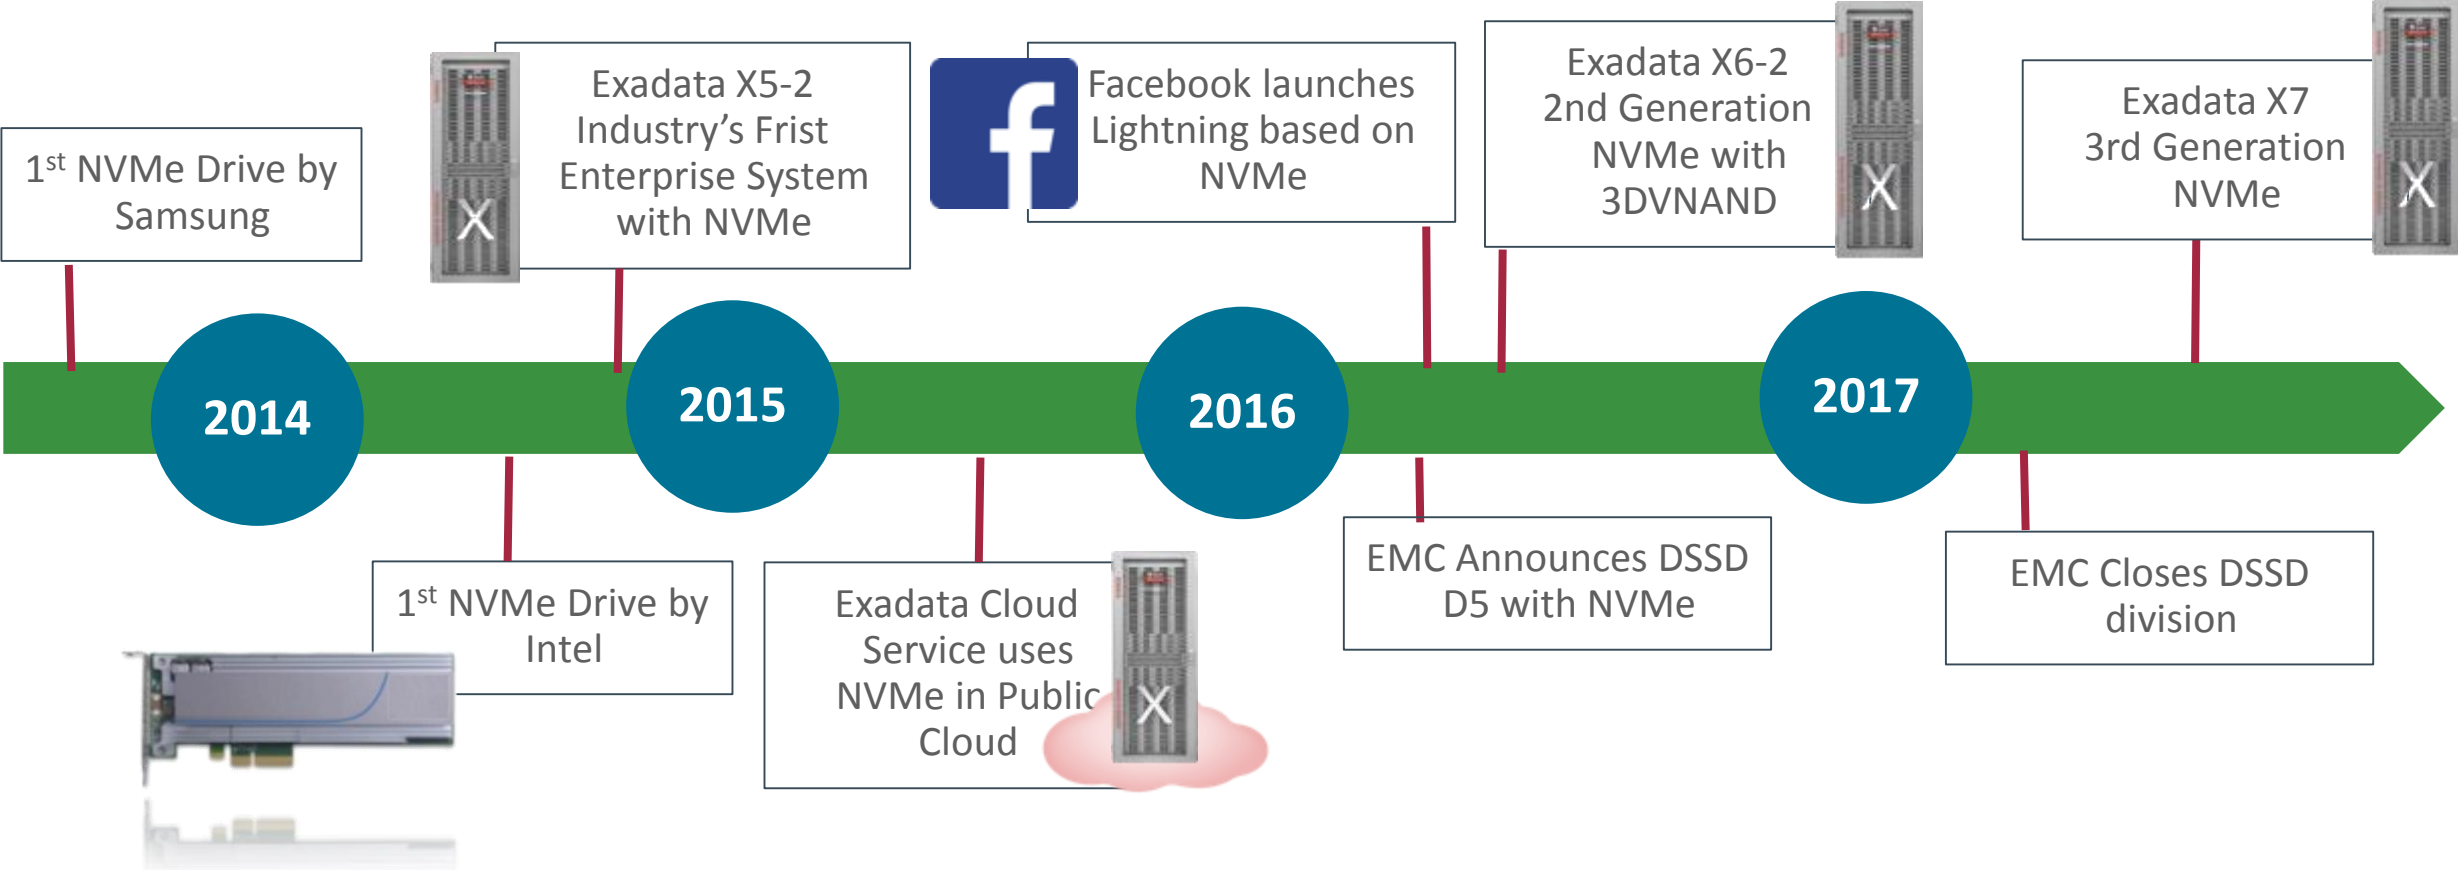
Pursuit of Performance: CPU

<u>2008</u>	<u>2009</u>	<u>2010</u>	<u>2012</u>	<u>2013</u>	<u>2014</u>	<u>2016</u>	<u>2017</u>
V1	V2	X2	X3	X4	X5	X6	X7
							
							
Harpertown 8 cores per server	Nahelam 8 cores per server	Westmere 12 cores per server	Sandy Bridge 16 cores per server	Ivy Bridge 24 cores per server	Haswell 36 cores per server	Broadwell 44 cores per server	Skylake 48 cores per server
Full Rack (8 servers)				192 cores	288 cores	352 cores	384 cores
Qtr Rack (2 servers)				48 cores	72 cores	88 cores	96 cores
Eighth Rack (2 servers x .5)				24 cores	36 cores	44 cores	48 cores

Pursuit of Performance: Flash


	<u>2008</u>	<u>2009</u>	<u>2010</u>	<u>2012</u>	<u>2013</u>	<u>2014</u>	<u>2016</u>	<u>2017</u>
	V1	V2	X2	X3	X4	X5	X6	X7
								
								
					1.8 GB/sec	2.5 GB/sec	5.4 GB/sec	5.5 GB/sec
Full Rack Scan Throughput					100 GB/sec	180 GB/sec	301 GB/sec	350 GB/sec
Qtr Rack Scan Throughput					22 TB/sec	30 GB/sec	64 GB/sec	75 GB/sec
Full Rack OLTP Read IOPS					2.8 M	4.1 M	5.6 M	5.97 M
Qtr Rack OLTP Read IOPS					570 K	1 M	1.1 M	1.2 M

Pursuit of Performance: Leading NVMe Adoption



New Hardware Paradigms: **SQL in Silicon**

SPARC M7



32 Database Accelerators (DAX)

- SIMD instructions were designed for HPC and Graphics, not SQL
- SPARC M7 Database Accelerators (DAX): **32 special cores for SQL**
- Oracle Zip (OZIP) bit decompress in DAX:
 - OZIP at runs at **> 120GB/sec** in DAX
 - **2x** capacity with speed of dictionary
- Columnar processing in hardware:
 - E.g. Find all values like “Russia”
 - Up to **170 Billion rows** per second!
 - Works on semantically compressed data

New Hardware Paradigms: **NVRAM**

- Launch Partner with Intel on 3D Xpoint
 - Worked closely with them on chipset, eADR, and CPU instructions for databases
- Current product uses NVRAM for “big memory”
- New NVRAM low-level libraries
 - Encapsulate correctness issues to avoid potential out-of-order flush
- Active research on innovative database storage techniques on NVRAM

Integrated Cloud

Applications & Platform Services

ORACLE®