Memory Performance With Shared Flash

Gurmeet Goindi
Exadata Product Management

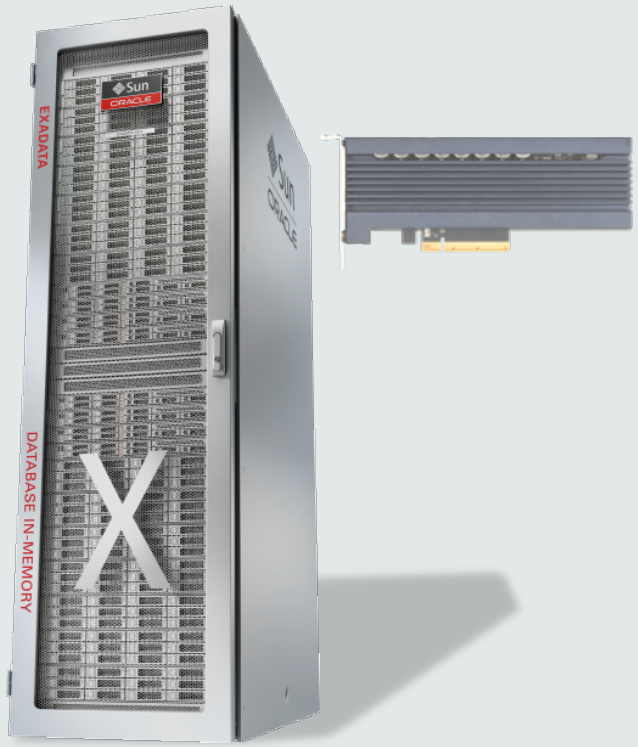

ORACLE

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.

Storage Arrays Fundamentally Bottleneck Performance

Sharing the Performance of Flash Across Servers is a Huge Challenge

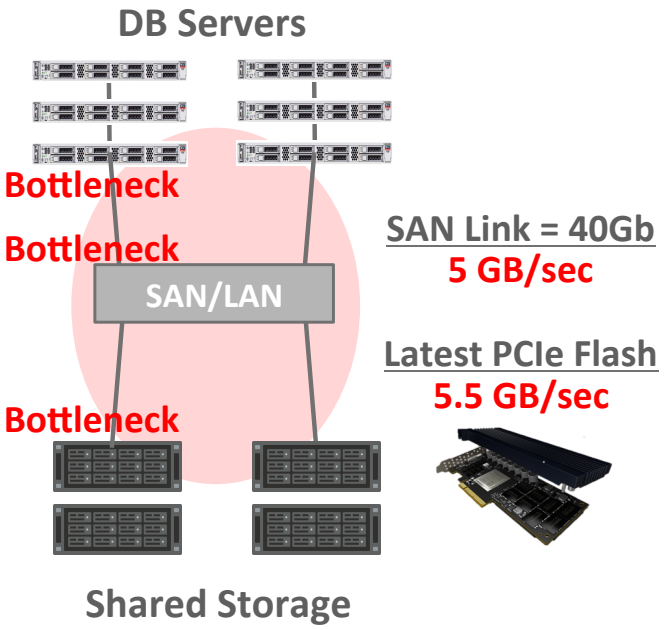

SAN Link = 40Gb
5 GB/sec

Latest PCIe Flash
5.5 GB/sec

- **Shared storage has many compelling advantages**
  - Much better space utilization, security, management, reliability
  - Enables DB consolidation, DB high availability, RAC scale-out
- **Sharing capacity is easy, sharing performance is hard**
  - Deliver performance of all shared flash drives to any server(s)
- **Flash performance has improved dramatically causing 100X bottlenecks across shared storage stack**
  - Speed of one flash card is now similar to fastest SAN or LAN link
  - A few flash cards deliver more throughput than
    - A storage array can output, a SAN/LAN can transfer, a server can input
  - Scale-out storage helps but does not solve the problem

ORACLE

Flash Performance is Wasted by Shared Storage Arrays

SAN Link = 40Gb
5 GB/sec

Latest PCIe Flash
5.4 GB/sec

- Recent improvements flash performance are causing 100X bottlenecks across shared storage stack
  - Speed of one flash card is now similar to fastest SAN or LAN link
  - Throughput of a few flash cards is too fast to transfer to servers

All-Flash Storage Array IO Path: *many steps, each adds latency and creates bottlenecks*

Exadata Achieves Memory Performance with Shared Flash

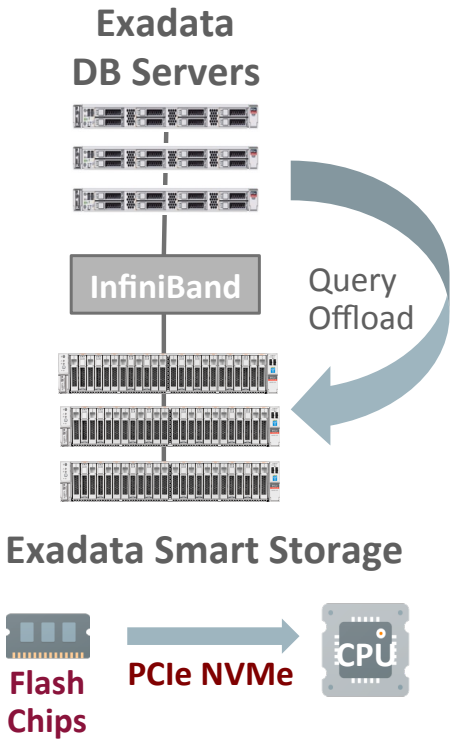

- Exadata X6 delivers 300GB/sec flash bandwidth to any server
  - Approaches 800GB/sec aggregate DRAM bandwidth of DB servers
- Must move compute to data to achieve full flash potential
  - Requires owning full stack, can't be solved in storage alone
- Fundamentally, Storage Arrays can share flash capacity but not flash performance
  - Even with next gen scale-out, PCIe networks, or NVMe over fabric
- Shared storage with memory level bandwidth is a paradigm change in the industry
  - Get near DRAM throughput, with the capacity of shared flash

ORACLE

Exadata Smart Scan
Move Queries to Data, Not Data to Queries

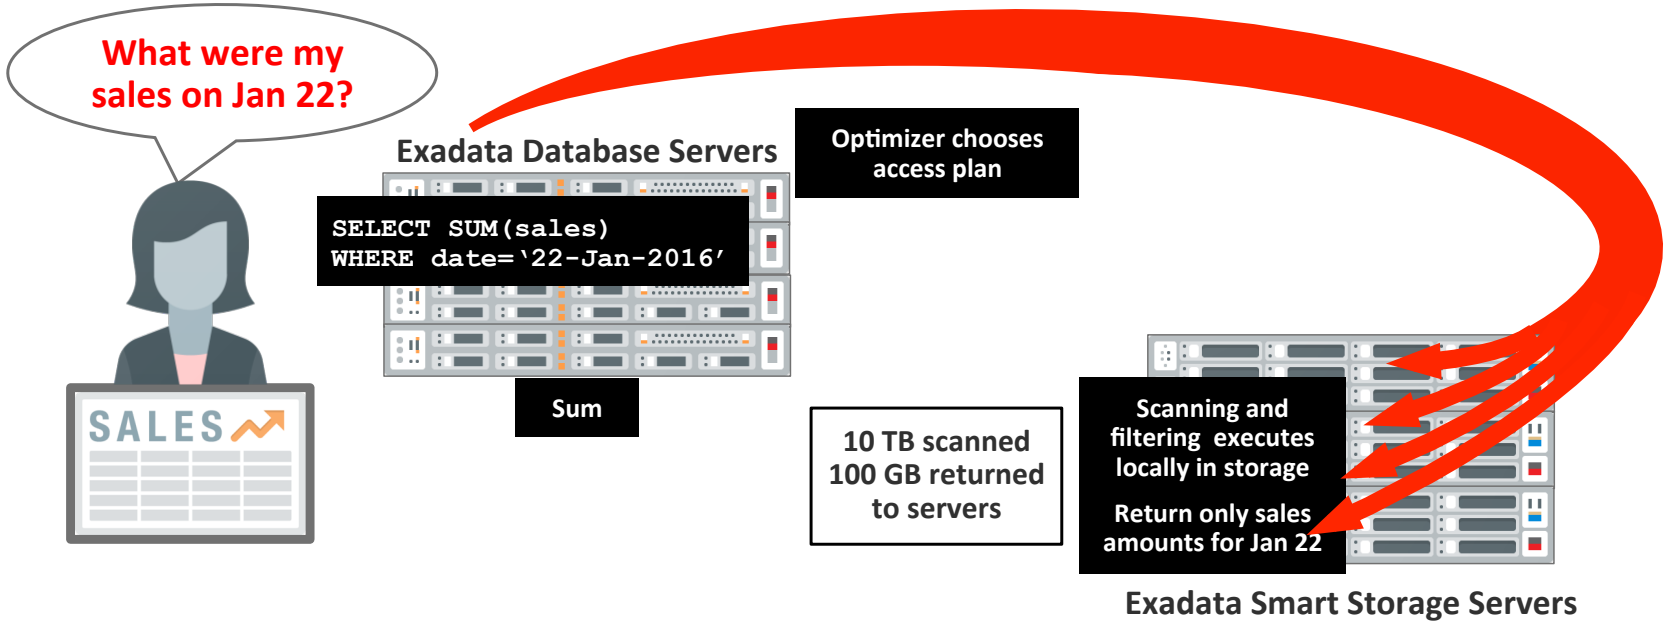

Only Exadata Achieves Full Performance of Shared Flash

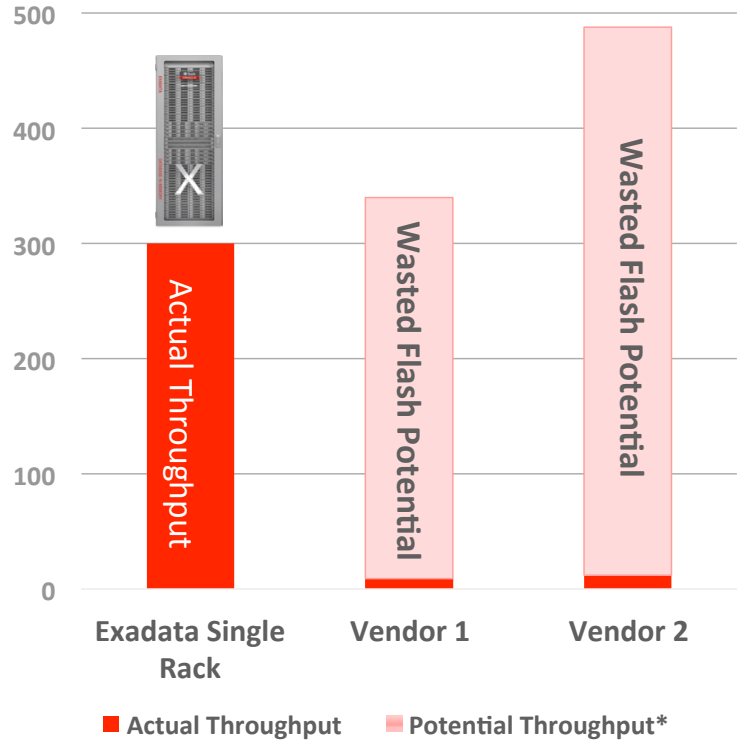

- **Leading All-Flash Storage Arrays achieve under 3% of potential flash throughput**
  - Vendor 1 – 132 MB/sec per flash drive
  - Vendor 2 – 120 MB/sec per flash drive
  - Spinning disk level throughput!
  - AND can't scale-out for higher performance
  - AND can't share even this slow performance due to bottleneck at server inputs
- **Exadata X6 achieves full flash throughput**
  - 5400 MB/sec per drive
- **Exadata also achieves much faster OLTP IOs**
  - 5.6 Million IOPs, 250us latency even at 2.4M IOs

*Potential Throughput based on number of flash devices

Exadata X6 Delivers Breakthrough DB IO Performance

301 GB/sec Analytic Throughput

5.6 Million 8K OLTP Read IOPS

5.2 Million 8K OLTP Write IOPS

250 us IO latency at 2.4 Million IOPS

Scales higher as racks are added

Performance of 1 Exadata Rack with 10 DB servers and 12 Extreme Flash storage servers

ORACLE

Preview: Redesigning Scan Offload for Memory Throughput

- With Exadata Flash throughput approaching memory throughput, SQL bottleneck moves from IO to CPU
- Exadata will automatically transform table data into In-memory DB columnar formats in Exadata flash cache
  - Dual format architecture extended from DRAM to flash
- Enables fast vector processing for storage server queries
  - Smart Scan results sent to DB using In-Memory Columnar format to reduce DB CPU usage
- Uniquely optimizes next generation flash as memory

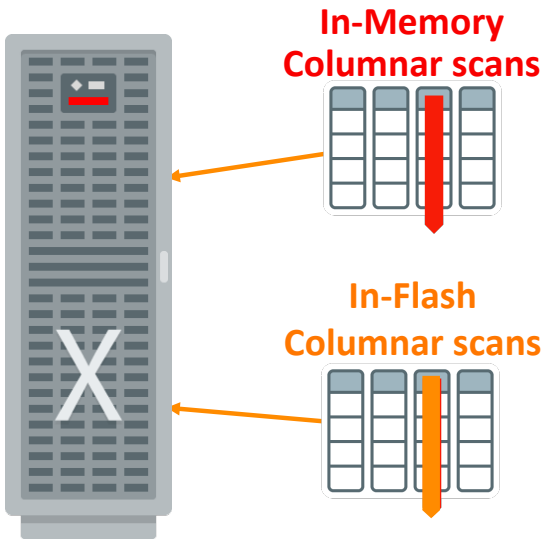

Production in 2016

Exadata Customer Case Studies

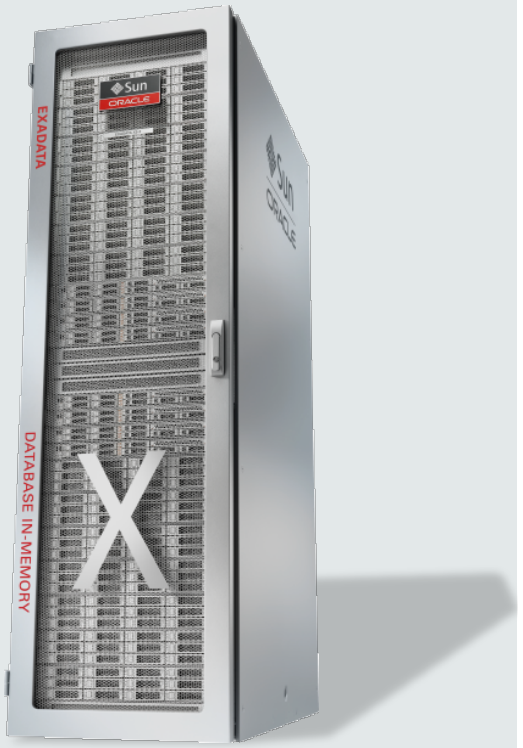

ORACLE

Proven at Thousands of Ultra-Critical Deployments since 2008

Half OLTP, Half Analytics, Many Mixed

- **Petabyte Warehouses**
- **Online Financial Trading**
- **Business Applications**
  - SAP, Oracle, Siebel, PSFT, …
- **Massive DB Consolidation**
- **Public SaaS Clouds**
  - Oracle Fusion Apps, Salesforce, SAS, …

4 OF THE TOP 5

BANKS, TELECOMS, RETAILERS RUN EXADATA

ORACLE

NTT docomo : MoBills（Mobile Billing System）

Benefits	Faster Billing Processing	Maximum Availability	Reduced Operational Cost	Reduced Introduction Cost	Data Center Cost Savings
"MoBills is a very important position as a mission-critical system to promote efforts toward the realization of "+d". Oracle Exadata is running very stable as a expected performance. We will continue to use the "Oracle Exadata" and we would like to establish a further advantage for our business." - Shimamura, Manager, Information System Department, NTT docomo	10X speedup 3 million SQL /sec	Local & Remote Standby	50%	25%	90% Space Reduction

Business Objectives

- Real-Time Billing Platform for 66 million customer
- Dramatically improve performance and availability
- Reduce cost and complexity

Solution

- Oracle Exadata X4-2 : 30 racks
- Oracle MAA (RAC / Active Data Guard - Local & Remote Standby database)

Pre-Exadata

Upgrade & Migration

- Real-Time Billing Processing
- High-end SMP Server
- + High-end Storage : 350 racks
- Storage Mirror Backup
- Storage Mirror Replication
- Oracle 9i Database Release 2

Exadata MAA

30 racks / Local & Remote Standby / RMAN backup

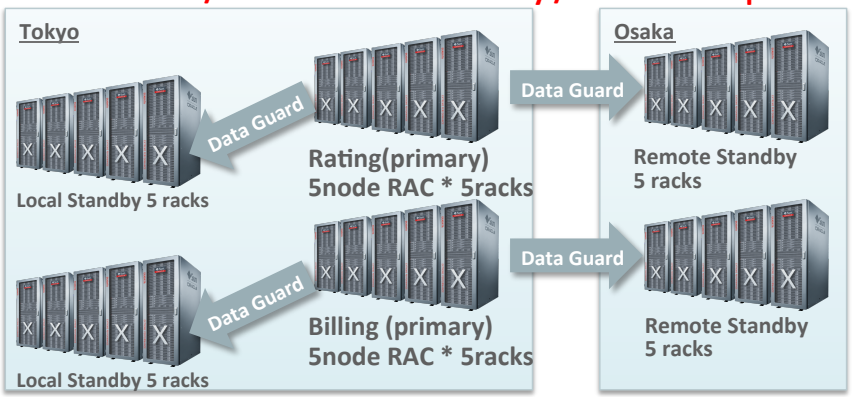

VocaLink: High-Volume OLTP

VOCALINK
safer payments, smarter partner

Benefits	Fast and Scalable Performance	Zero Unplanned Downtime	More Capacity	Zero Data Loss DR Protection
"Our critical electronic payments service has been live on Exadata since early 2011 with 100% uptime. The service reliably processes the transfer of billions of Euros per week and achieves subsecond response times for online enquiries." - Martin McGeough, Database Technical Architect			20M Txns/Day 	

Business Objectives

- Enable real-time enquiries against transaction data
- Process 20 million financial transactions/day
- Sub-second response time
- Zero data loss disaster recovery

Solution

- Three Qtr Rack Exadata X4-2 systems + Oracle MAA

Pre-Exadata
4 Prod Servers, 2 SANs

- Electronic payment system
- Mixed workload, high volume OLTP and batch processing

Exadata X4-2
Quarter Rack

- Performance boost with no application change using Exadata Smart Flash Cache and Smart Scan

Active Data Guard
Zero Data Loss
100 km (60 miles)

Exadata X4-2
Quarter Rack

- Zero RPO, 15 minute RTO
- Auto block repair
- Standby-first maintenance

ORACLE

Alpha Natural Resources: OLTP

Alpha Natural Resources

Benefits	Faster Applications	Reduced TCO	Maximum Availability	Data Center Cost Savings
"Oracle Exadata has enabled us to deliver exceptional service to our business users while reducing the cost involved . Exadata has been a big win for our business and for IT." - Saul Hernandez, CIO	5X+ Speedup	50%	99.95%	

Business Objectives

- Reduce cost and complexity
- Deploy Oracle E-Business Suite quickly
- Dramatically improve performance and availability

Solution

- X2-2 for production e-Business Suite
- X2-2 for DR/test/QA
- Database Appliance for dev

Pre-Exadata
(Hosted)

- IBM x86 servers
- Three O/S versions

Exadata X2-2
DB Consolidation

Data Guard

Data Guard

- Production
- E-Business Suite

Oracle Database Appliance
Dev/Test

- Disaster protection
- Test and QA

ORACLE


ORACLE®