

ORACLE®

Oracle Rdb customer experience in Japan

Satoshi Asano
Oracle Rdb Engineering
October, 2019

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.

Agenda

1 Overview customers in Japan

2 Customer references.

- Customer references – 1
- Customer references – 2
- Customer references – 3
- Customer references – 4

3 SSD & Oracle Rdb

A man with glasses and a denim shirt is gesturing while talking to a woman in a yellow top. They are sitting at a table with several papers featuring charts and graphs. A coffee cup and a smartphone are also on the table. The background is a blurred office environment.

Overview customers in Japan

Which area is Rdb used.

- Liquid Crystal Display(LCD) manufacturing
- Mobile phone system
- Trading commodities
- Restaurant chain
- Airlines
- Railway company
- Beer Brewer
- Warehouse Management System
- Semiconductor manufacturing
- More...

A man with glasses and a denim shirt is gesturing while talking to a woman in a yellow top. They are sitting at a wooden table with several papers featuring charts and graphs. A coffee cup and a smartphone are also on the table. The background is a blurred office environment.

Customer references - 1

Long time customer


Customer Reference -1

System – Overview

- Long time with us, more than 30 years.
- MIA(Multivendor Integration Architecture)
- Consists of three systems
 - EIS: Main system.
 - ACS: delivering system.
 - A-IMG: Store designs and image data(the smallest system)
- Software
 - OpenVMS
 - Oracle Rdb
 - ACMS xp -> ACMS
 - Others.
- Hardware
 - VAX -> Alpha -> Integrity

Customer Reference -1

MIA(Multivendor Integration Architecture)



Customer Reference -1

Database



- Provide services based on DATABASES.
- Corporate data.
- Approximately eight million corporate data entries gathered from all over Japan.
- Accumulate changes on day-to-day basis into local database, and main database has been modified once a month.
- 20% of records in the database has been changed in a year.
- Talk about the database more later.

Apr. 1986: This project started.

1987: First system started.




1989: VAX8810 added. **Active/Active cluster system**

<<<added new machines & storages>>>

1995: A&S started

1995: MLS and A-IMG in production, Both projects are MIA project.



1999: ACS started first Alpha system



: VAX

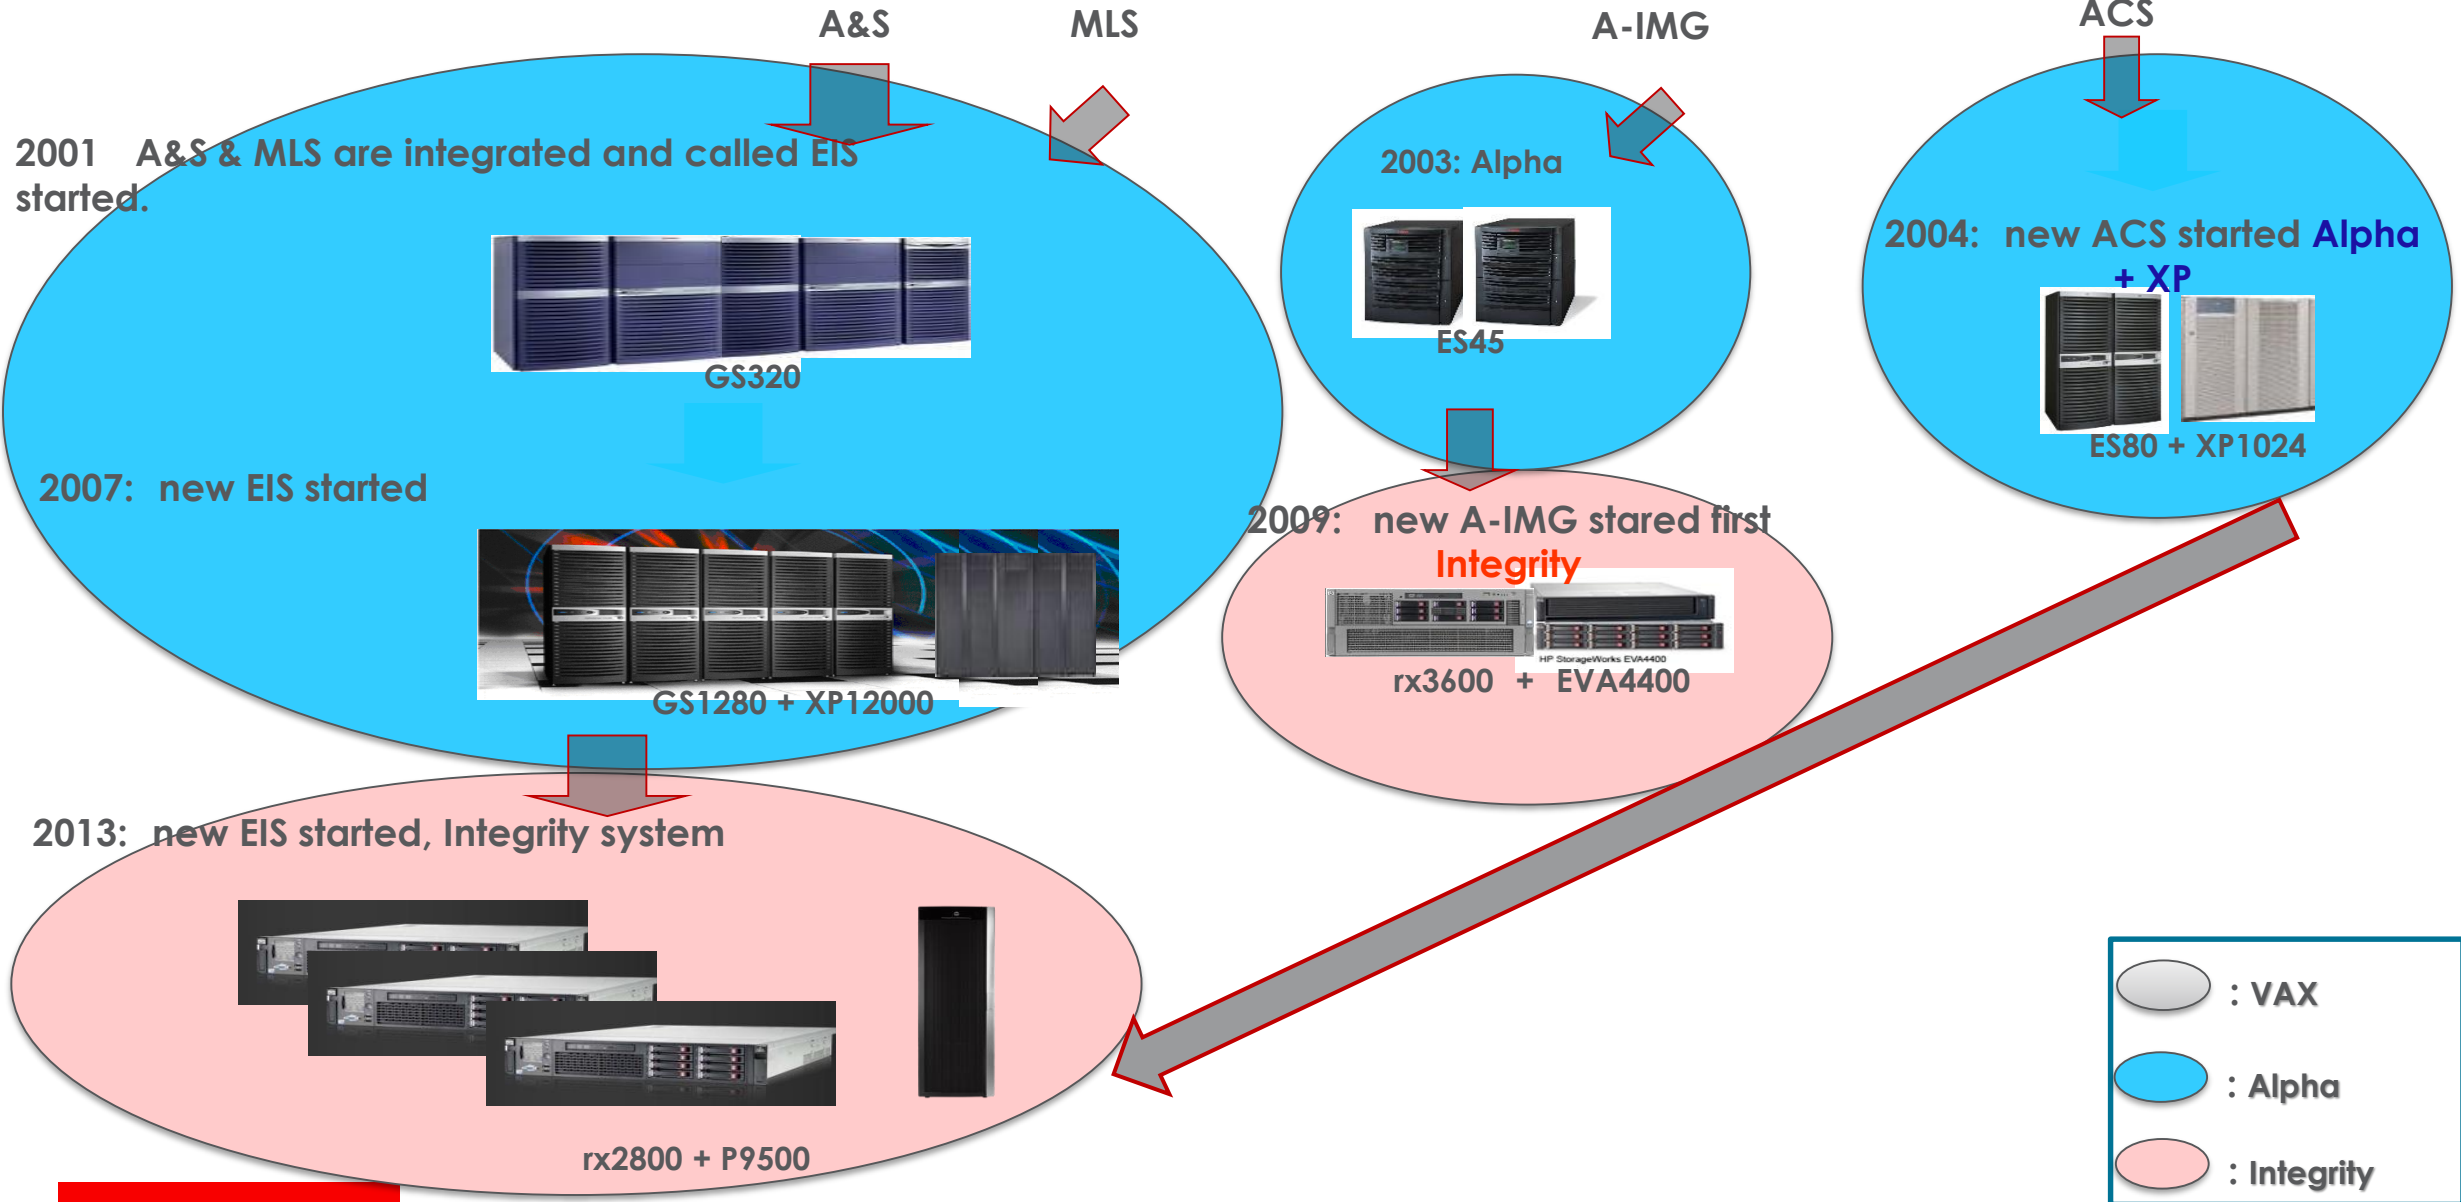
: Alpha

: Integrity

EIS

A-IMG

ACS



Customer Reference -1

Summary - First released MIA system on VAX

- MLS runs on VAX10700*4 cluster, A-IMG run on VAX4000 on 10 each area.
- OpenVMS, ACMSxp, CAP(MIA compliance, kind of middleware) and Rdb.
- VMS5.5-2M -> VMS6.1, RdbV4.2 -> Rdb6.1 and ACMSxp.
- Both are MIA compliance project.

Customer Reference -1

Summary - Alpha system

- In 2001, A&S and MLS system are integrated and called EIS system and run on GS320 Alpha machines. (VMS7.2-1, Rdb7.0-*, ACMSxp)
- MLS system divided Japan into 10 areas and each area has databases. When MLS integrated with A&S, 10 database merged with one big database.
- In 2004, ACS was ported to ES80*2 cluster + XP1024. (VMS7.3.2, Rdb7.1.* ACMSxp)
- Remove CAP portion, because it can be replaced with ACMS XP.
- In 2007, new EIS runs on the last Alpha EV7 GS1280*3 cluster + XP12000. With GS1280s, it was very good performance and robust!
(VMS7.3.2, Rdb7.1.* ACMSxp)

Customer Reference -1

Summary - Itanium system

- In 2013, EIS and ACS system were integrated and ported to Integrity rx2800*3 cluster + P9500 called EIS system(named from EIS), OpenVMS/Itanium does not support ACMSxp, so all source codes related to ACMSxp were ported to ACMS.
- VMS8.4, Rdb7.2* and ACMS


Customer Reference -1

Database

- 6 transactional(accessed by online) databases & 6 other databases in EIS.
- 10 transactional databases & 2 other databases in ACS.
- Total 1.7Tb
- Number of storage area: 34,000(16 transactional DBs).
- Number of 200 tables defined in EIS 6 databases.
- The biggest db in the system is 400G.
- Number of storage area: 5,000(The biggest DB).
- 120 databases are opened in one node.

Problems

Number of Bugs reported : 21



A man with glasses and a denim shirt is gesturing while talking to a woman in a yellow top. They are sitting at a wooden table with several papers, a coffee cup, and a smartphone. The papers contain various charts and graphs. The background is a blurred office environment.

Customer references - 2

Ongoing project

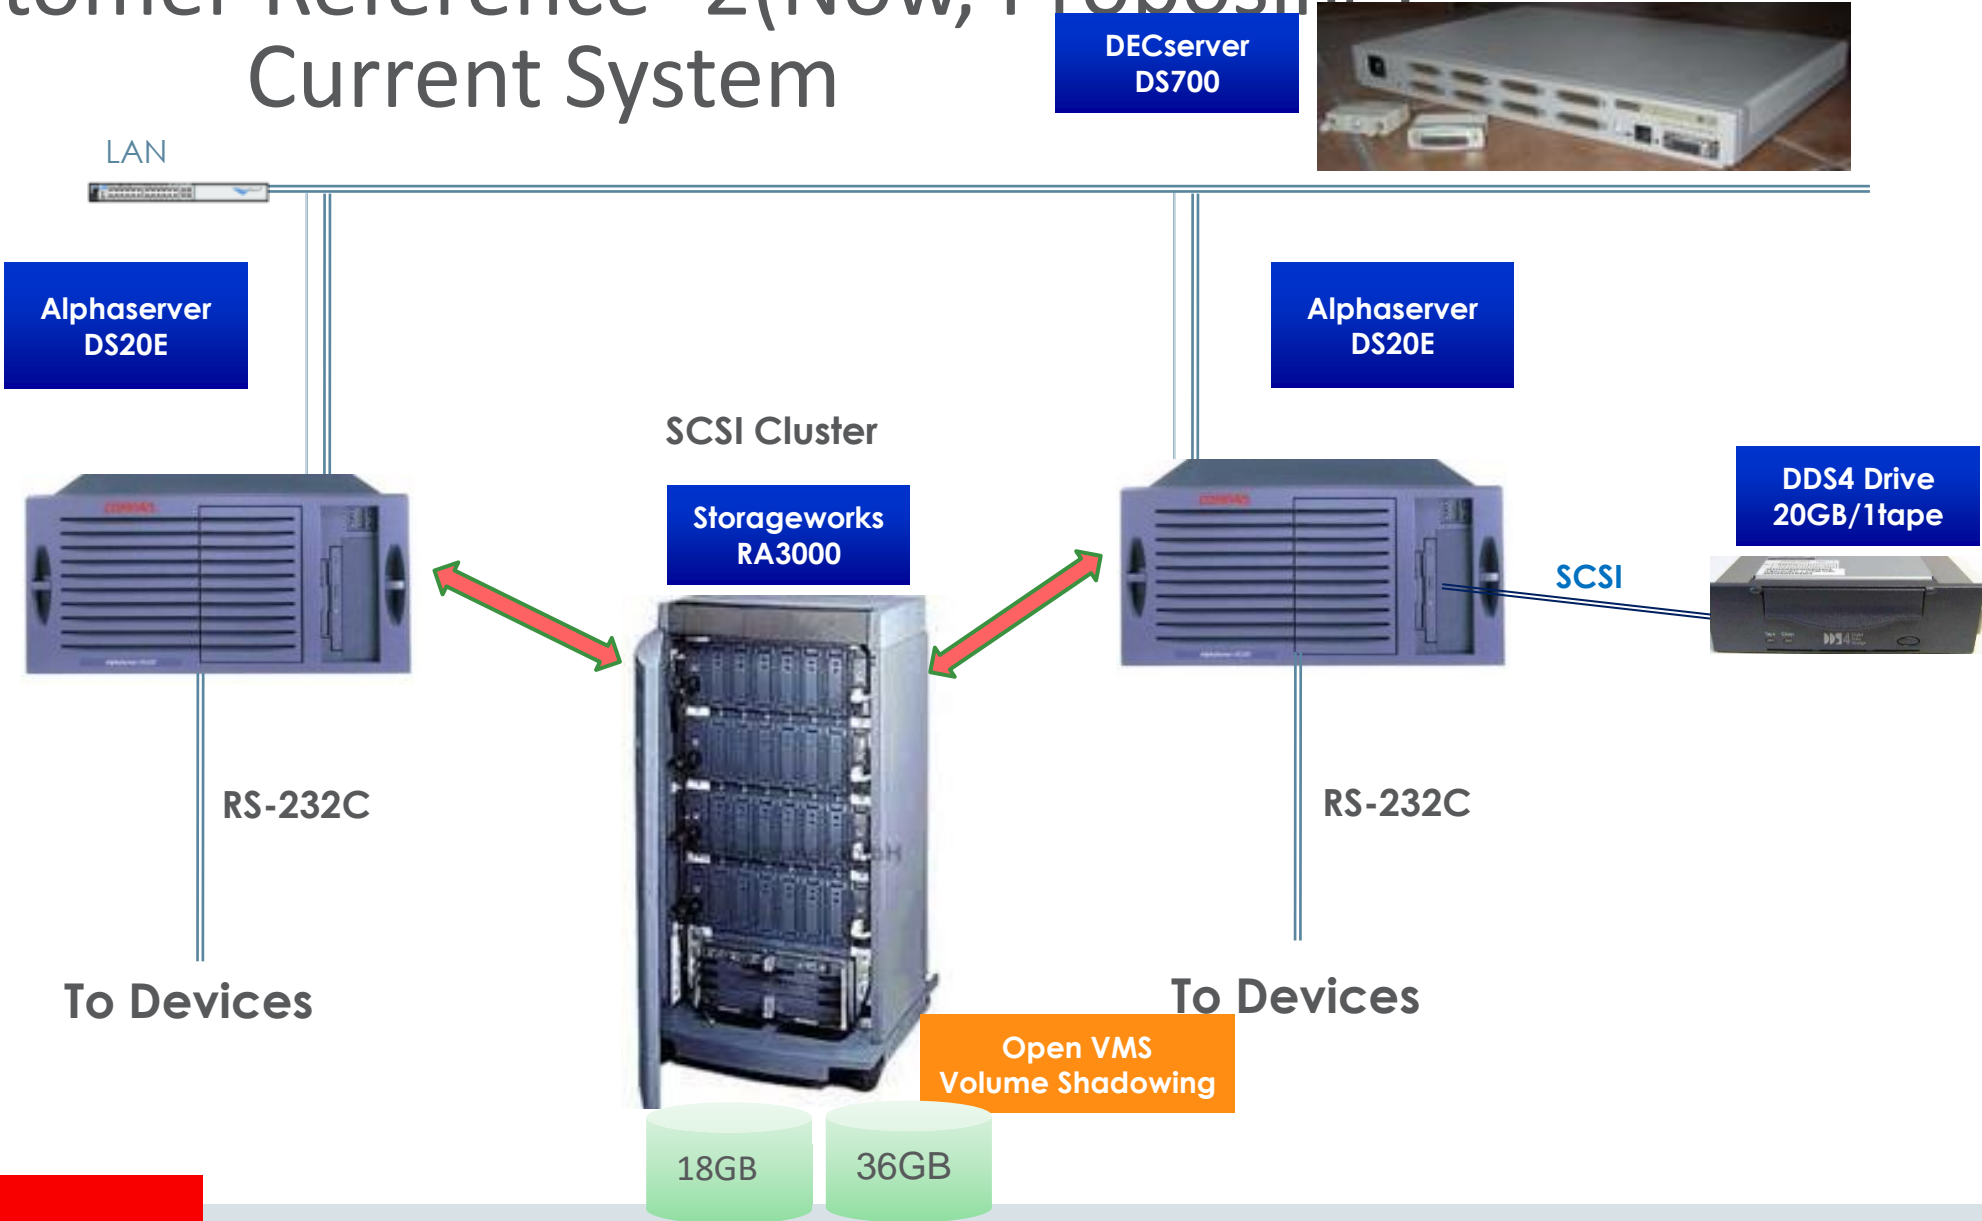
Customer Reference -2(Now, Proposing) System Overview

- Process management system of LCD manufacturing plant
 - Manage the state of each process from start to completion
 - Next, it decides which device to manufacture and automatically carries it to that device.
 - Online communication is between the conveying device and the manufacturing equipment.

Customer Reference -2(Now, Proposing) History

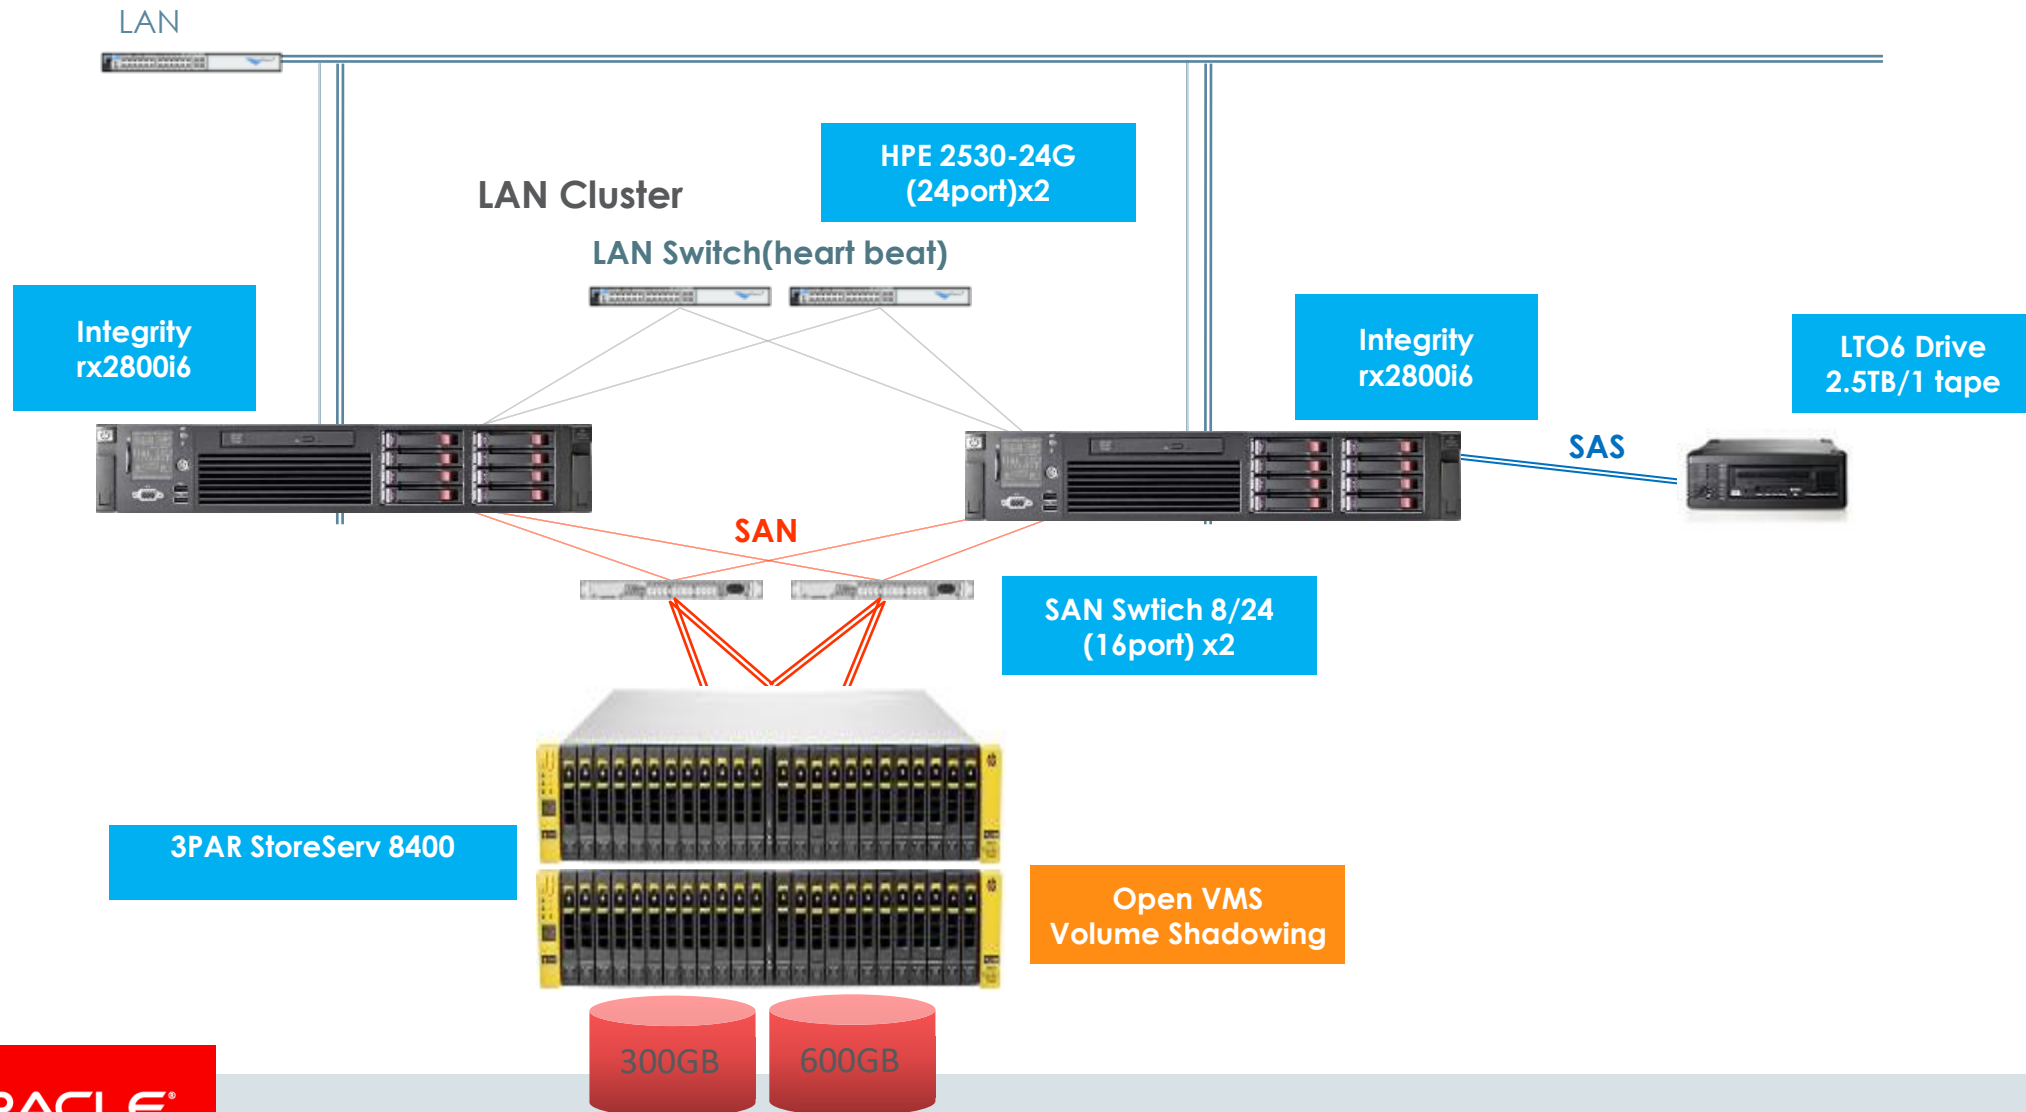
- LCD manufacturing process management system
- In 1992, in production with VAX/VMS(DEC Japan).
- In 1999, ported to Alpha/VMS(Compaq Japan).
- Until 2007 supported by HP Japan.

Customer Reference -2(Now, Proposing) Current System



Customer Reference -2(Now, Proposing)

Example of proposed System



Customer Reference -2(Now, Proposing)

List of Software

	Name of Product	Version(current)	Version(proposed)	
1	OpenVMS	7.2-1	8.4-2L1	Japanese OpenVMS
2	DECnet IV/MOP			Included in OpenVMS
3	LAT			Included in OpenVMS
4	TCPIP	J5.0	5.7 ECO5F	
5	Compaq C	6.4	7.4	
6	DECforms	3.2	4.2	
7	Japanese DCPS	1.8	2.8	
8	Oracle Rdb	7.0-5	7.3.2.4	
9	Oracle SQL/Services	7.0-5	7.3.3	
10	Oracle CDD/Repository	7.0	7.2.0.6.0	
11	Oracle MQ	4.0	5.0	
12	UDMS	?.?	—	Not available

Customer Reference -2(Now, Proposing) Issues

- UDMS – text base user interface which accessing Oracle Rdb and other files. There is no Itanium version.
- Article we found on internet, computer world published 4/15/1991.
- Alternative products
 - DECForms + C + SQL preprocessor(or moudle)
 - Java + JDBC
 - PHP
 - others

SYSTEMS & SOFTWARE

NEW PRODUCTS — SOFTWARE

Database management systems

Interactive Software Systems, Inc. has integrated its User Data Management System (UDMS), a user-oriented data management system designed for Digital Equipment Corp.'s VAX/VMS machines, with Ross Systems, Inc.'s business software. Integration is accomplished via a predefined data dictionary. The amalgamated product features enhanced report writing capabilities across all Ross Systems business applications. It also provides links to other applications with diverse databases, such as Oracle Corp.'s Oracle and DEC's RDB. Pricing for UDMS ranges from \$4,500 to \$54,000, depending on VAX CPU size. The Ross UDMS data dictionary is included with the product free of charge until June 1, 1991. Interactive Software Systems 7175 W. Jefferson Ave. Denver, Colo. 80235 (303) 987-1001

Applications packages

Software Engineering of Ameri-

a copy of Panedit, an interactive software utility package that facilitates interface building, and package libraries for writing Ada applications behind user interfaces. Pricing ranges from \$1,095 for an Intel Corp. 80386-based personal computer to \$4,300 for a large Sun Microsystems, Inc.

server. The product began shipping last month. Objective Interface Systems 1875 Campus Commons Drive Reston, Va. 22091 (703) 264-1900

ca, Inc. has announced an interactive job scheduling system that includes a TSO/ISPF interface with split-screen capability. Computer Scheduling and Reporting (CSAR)/MVS Release 3.1 features job control language preprocessing that supports Restructured Extended Executor-style functions. Cross-memory services, flexible frequency standards and tape pull list reporting

are a C of M ESA base P to \$2 Soft Am 200 Lak (516

Presenting All T Software

It's going to be quite a decade. Every employee, every department, every division will be expected to do more—with less. Including the data

across multiple operating systems and hardware platforms. The result is consistently higher levels of service. Improved response

balancin lation ar and autc CA-UN

Customer Reference -2(Now, Proposing) Issues

- No alternative factory line for test.
 - Cannot run old system and new system in parallel.
 - Cannot compare results between old and new
 - ➡ Emulator/Simulator
- Only 8 hours window for transition.
 - Alpha to Integrity.
 - Databases
 - Undo – just in case.
 - Etc.

A man with glasses and a denim shirt is gesturing while talking to a woman in a yellow top. They are sitting at a wooden table with several papers featuring charts and graphs. A coffee cup and a smartphone are also on the table. The background is a blurred office environment.

Customer references - 3

During a migration

Customer Reference -3 (outside Japan)

When migrating Alpha to Itanium

- Mobile phone carrier
- Migrate Alpha to Itanium a few years before.
- I attended the migration during weekend and following week.
- Before migration, they are testing on Itanium and compare the result with Alpha.
- The customer found the performance problem on a query which is a simple 2 table join.
- Took a few minutes on Alpha, more than three hours on Itanium.

Customer Reference -3(outside Japan)

When migrating Alpha to Itanium

- The customer is worry that they haven't found all of the sql statement that this problem might affect yet.
- It became show stopper.
- This customer said both Alpha and Itanium have the same configuration and both of them use read write transaction.
- Set transaction read only was a workaround, and it worked.
- The mystery still remained why the same query with same read write transaction on Alpha is much faster than on Itanium.
- Why...

Customer Reference -3(outside Japan)

When migrating Alpha to Itanium

- DECLARE TRANSACTION READ ONLY is defined in sqlini on Alpha, but not defined on Itanium machine.

A man with glasses and a denim shirt is gesturing while talking to a woman in a yellow top. They are sitting at a wooden table with several papers featuring charts and graphs. A coffee cup and a smartphone are also on the table. The background is a blurred office environment.

Customer references - 4

Something new?

Customer Reference -4

Something new or different

- Restaurant chain.
- About 400 restaurants in Japan.
- Running for factories, warehouse, POS system, salary and so on.
- Something new
 - Such as Cloud, Agile like DevOps and Out Source them.
- Can do it with OpenVMS and Oracle Rdb
 - X86 in near future.
 - Support many network protocols.
 - Computer languages.


SSD and Rdb

A man with glasses and a denim shirt is gesturing while talking to a woman in a yellow top. They are sitting at a wooden table with several papers containing charts and graphs. A coffee cup and a smartphone are also on the table. The background is a blurred office environment.

Characteristic of SSD



HDD




SSD

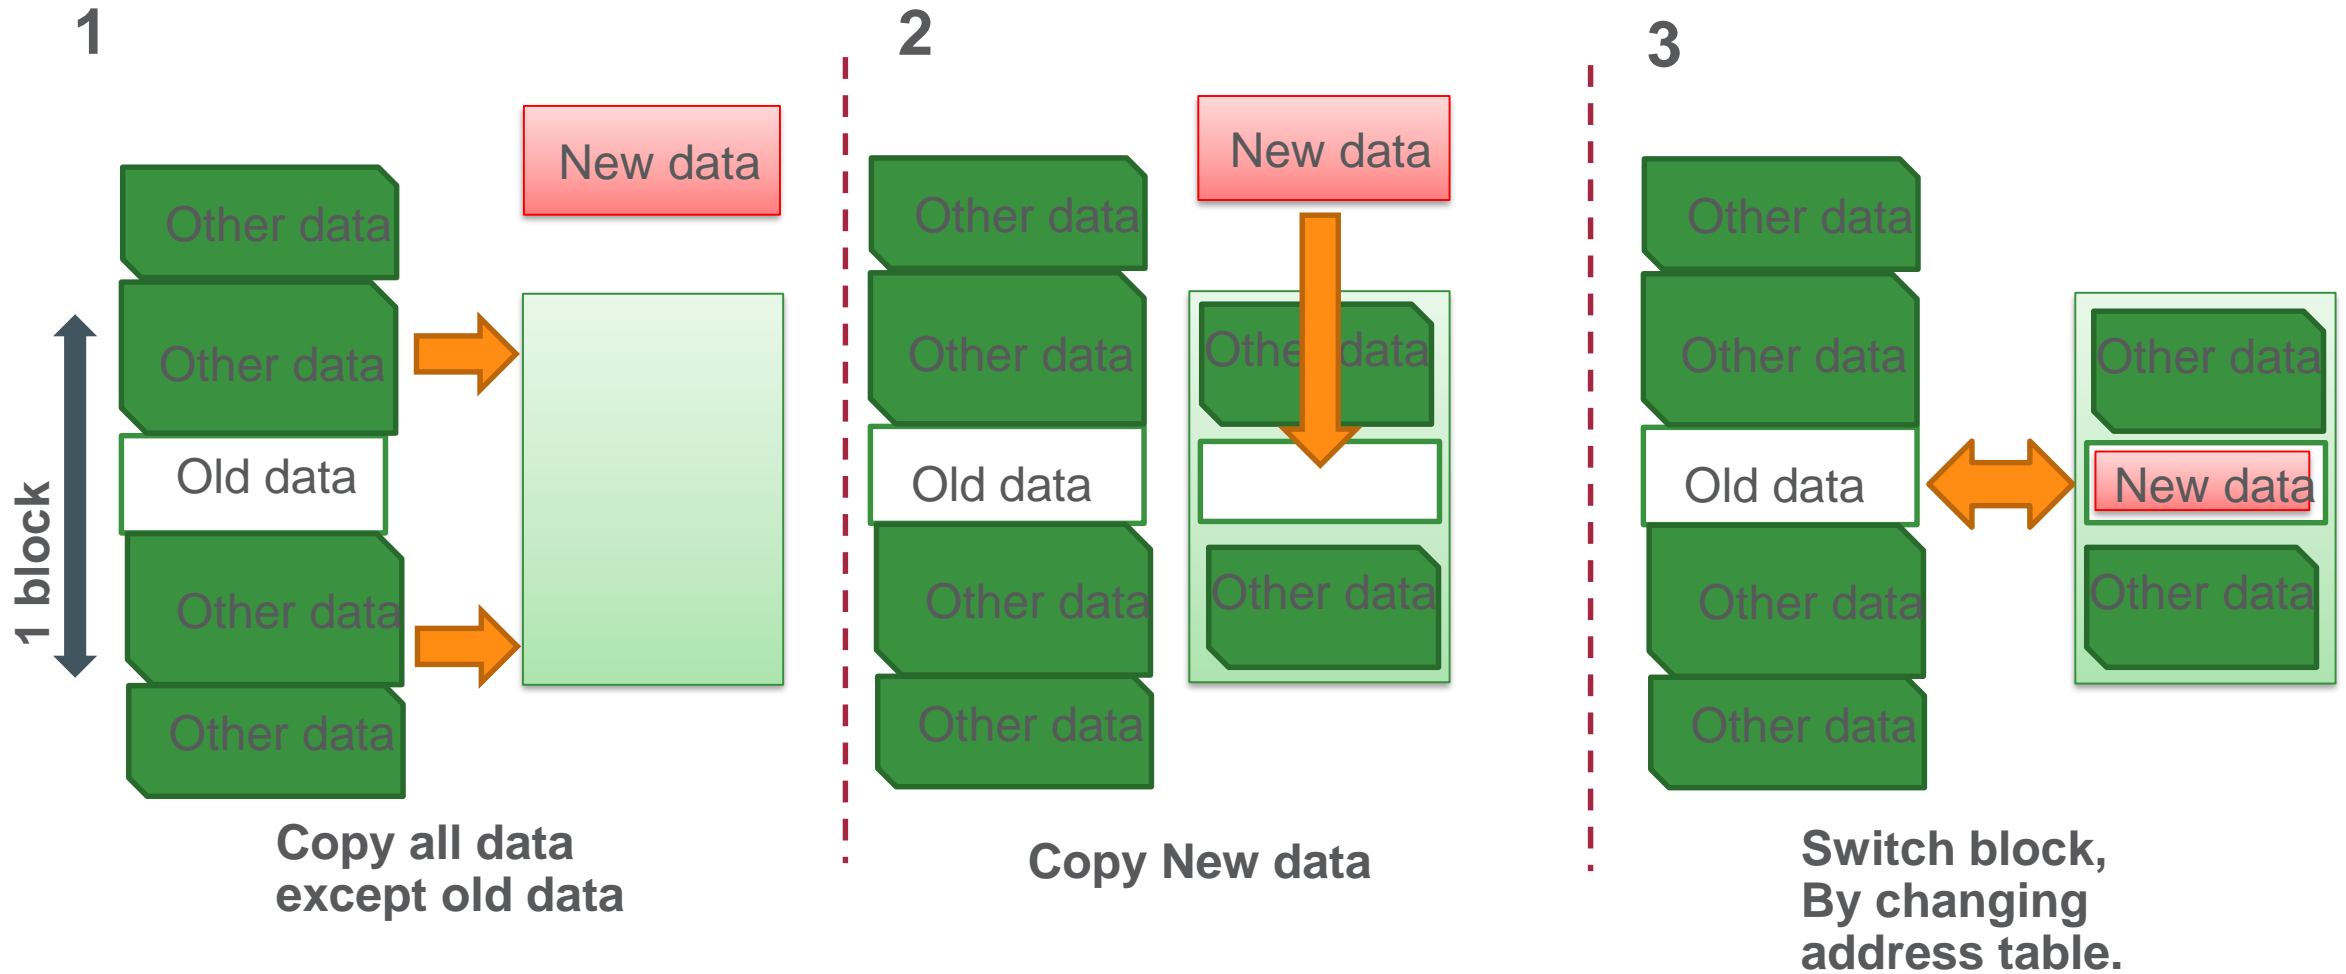
Characteristic of SSD

- No seek time
- Unlike HDD, read is much faster than write.
- Unlike HDD, random access is fast.
- Less power draw than HDD
- Short life span(it's getting longer...)
- Update data is slow
 - Data can not overwritten.
 - need to copy whole block(silicon substrate) to other place such as SSD or buffer.
 - then, replace old data with new data,
 - and switch block by changing address table.
 - delete old block, sooner or later.

Data change(update) : HDD




Data Change(update): SSD



Data Change: SSD

4



- Delete the old block to make a empty space.
- Or leave the old block, and later it will be a empty space by garbage collection.

Test machine Configuration

- **Comparison with three machines**

- [[[SSD-equipped machine]]]**

- Server: HP Integrity rx2800 i4 (2.53GHz/32.0MB) 16CPU, 128GB memory

- Storage: 3PAR(All Flash) 20840(Cache 816GB, 112 SSDs), FC 8Gbps

- Internal Disk: 4spindles(1.2TB) under Smart Array P410i, Disk 15,000rpm

- [[[3PAR, HDD-equipped machine]]]**

- Server: HP Integrity rx2800 i2 (1.60GHz/5.0MB) 8CPU, 32GB memory

- Storage: 3PAR 7400(Cache 8GB, Disk 15,000rpm, 22spindles), FC 8Gbps


- [[[EVA5000-equipped machine]]]**

- Server: HP Integrity rx2800 i4 (2.53GHz/32.0MB) 8CPU, 64GB memory


- Storage: EVA5000(Disk 15,000rpm, 24 spindles), FC 2Gbps

* : 3PAR All Flash, following slides except in the graph is referred as 3PAR (SSD).

3PAR StoreSrev SSD VS. 3PAR StoreSrev HDD



3PAR StoreSrev SSD VS. EVA5000 HDD



TPC-B

- One of TPC benchmark test.
- Generate a large amount of I / O to the database.
- It's already obsolete, but still useful.
- Implemented in C and SQL Pre-Compiler
- There are 4 tables in the database.


テーブル名	ACCOUNTS	BRANCHES	TELLERS	HISTORY
Operation	UPDATE/SELECT	UPDATE	UPDATE	INSERT
Num of rec. *1	6,400,000	64	640	0 *2

*1: These records were loaded before the test.

*2: Number of records were inserted was the same as number of transactions.

- Process multiplicity from 1 to 64, all tables except HISTORY were partitioned by KEY, and each storage area partitioned by key was corresponding with each process, so basically no conflict between processes.
- Run 4 minutes, then calculate TPS(transaction/sec).

Comparison of Hash Index



Summery

- Overall, 3PAR(SSD) is the faster than others.
- Especially , Hashed Index on SSD is effective very much.
- Other tests such as no AIJ, AIJ, Loading data and Recovery, it looks SSD is faster than others, but the test done under the different condition. So should do enough test before it get in production.
- Should consider to change of DB structure in order to take advantage of SSD's characteristics
 - Hashed from Sorted index
 - Most of activities to the database is read, and rare to write activity, should consider whole database or part of it locates on SSD.
- Actual use of SSD in production, should determine the location based on the results of adequate test.

Q&A



Safe Harbor Statement

The preceding 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.

Integrated Cloud

Applications & Platform Services

ORACLE®