Modernizing Oracle RAC Database on IBM Power

Practical guidance on getting current!

Table of contents

Overview

Challenge

Customers need to upgrade Oracle DB and IBM Power servers to assure continued support and manage costs.

Solution

This paper provides an overview of how to efficiently upgrade from Oracle 12c Release 1 Real Application Cluster (RAC) to 19c RAC, and from IBM Power8 to IBM Power10.

Executive overview

The objective of this white paper is to provide guidance for joint customers of IBM® and Oracle on how to upgrade from IBM Power8® to IBM Power10, and from Oracle RAC 12cR1 to Oracle RAC 19c. For reasons of long-term support, cost of maintenance, security, and overall efficiency, it is important to assist customers in planning their move to the most current long-term support version of the IBM server and Oracle Database.

To prepare this paper, the test team selected two workloads that are representative of customer OLTP and Analytics scenarios. Each workload was run on Oracle 12cR1 RAC using two Power8 processor-based servers. Both the Oracle database and the Power8 processor-based server were configured as per best practices, including the application of current service to IBM AIX® and Oracle Database.

With the workload running on the Power8's, the team then utilized the Oracle RAC capability to non-disruptively add the two Power10 nodes to the cluster, and then remove the two Power8 Systems allowing the workload to continue execution during the move to the new systems.

The test team recommends that all the plans to upgrade Oracle Database, and the systems it is running on, start with the IBM and Oracle software being upgraded to current levels.

Disclaimer

The results shown in this paper are for education purposes only. The results **do not** represent the full potential capability of IBM Power10 processor-based systems, Oracle Database 12c Release 1, Oracle Database 19c, and IBM FlashSystem® 5200. The results were derived from configurations that used default values, and generally accepted best practices, without any intense tuning on AIX, Oracle Database 12c Release 1, Oracle Database 19c, or the storage area network (SAN) server. The results would vary on different Power8 and Power10 processor-based systems and for different types of applications with differing workload characteristics.

Software

- Oracle RAC 12c R1(12.1.0.2) and 19c
- IBM AIX 7.2
- IBM AIX 7.3

Hardware

- IBM Power S824
- IBM Power S1024
- IBM FlashSystem 5200

Oracle RAC Database and IBM Power configurations

Using two node Oracle Real Application Cluster (RAC) database 12c Release 1 (12.1.0.2.0), each workload was driven to high system utilization on a Power8 logical partitions (LPARs) sized to meet the required throughput objectives. The team then moved the workload to Power10 LPARs that were sized to have half the number of Power8 cores.

The Grid Infrastructure and database were migrated from Power8 processor-based systems to Power10 processor-based systems by adding Power10 nodes into the existing Real Application Cluster. Before adding the Power10 nodes to the cluster, they were prepared with all pre-requisite tasks to make them eligible to be members of the Oracle Real Application Cluster.

The Power10 nodes were added to the existing Real Application Cluster by using "*addnode.sh*" script. After the Power10 nodes were added, the Power8 nodes were removed from the cluster, leaving only the remaining Power10 nodes in the cluster. This method of migration demonstrated the online migration of Oracle RAC from Power8 to Power10 nodes without interrupting the availability of the Oracle Database for the user workloads. Below are the high-level steps for adding and deleting the RAC nodes.

Oracle RAC "a*ddnode.sh*":

1. Execute the prerequisite steps on new node and run the `cluvfy` validation check.

   ```
   $ cluvfy comp peer -n <new_node> -refnode <existing_node>
   $ cluvfy stage -pre nodeadd -n <new_node> -fixup -verbose
   ```

   Note: Use latest cvu, download it from below link
   https://www.oracle.com/database/technologies/cvu-downloads.html

   Note: Using old cluvfy may report - Reference data is not available for release "12.1" on the Operation System Distribution "AIX7.2"

2. Run "addnode.sh" script from directory $GRID_HOME/addnode of any existing nodes.

   ```
   $GRID_HOME/addnode/addnode.sh -silent "CLUSTER_NEW_NODES={<new_node1>,<new_node2>}"
   "CLUSTER_NEW_VIRTUAL_HOSTNAMES={<new_node1_vip>,<new_node2_vip>}"
   ```

   Note: If prechecks failures need to be ignored then use options "`-ignoreSysPrereqs -ignorePrereq`" for addnode.sh script.

   Run root scripts when prompted.

3. For Oracle DB Home addition, use same "addnode.sh" script from directory $ORACLE_HOME/addnode of any existing nodes.

   ```
   $ORACLE_HOME/addnode/addnode.sh -silent "CLUSTER_NEW_NODES={<new_node>}"
   "CLUSTER_NEW_VIRTUAL_HOSTNAMES={<new_node_vip>}"
   ```

4. Add the new instances to the cluster database using *dbca* by selecting "Instance Management".

Deleting node in Oracle RAC:

1. Remove the DB instances that are running on nodes that are going to be deleted using *dbca* tool.

2. On deleting node update the DB Home inventory.
```
$ORACLE_HOME/oui/bin/runInstaller -updateNodeList ORACLE_HOME=<DB
HOME> "CLUSTER_NODES={<deleting_node>}" CRS=TRUE -local
```

3. Deinstall oracle DB home using -local option run from deleting node.
```
$ORACLE_HOME/deinstall/deinstall -local
```

4. Update inventory in remaining nodes of cluster.
```
$ORACLE_HOME/oui/bin/runInstaller -updateNodeList ORACLE_HOME=<DB
HOME> "CLUSTER_NODES={<node1>,<node2>}"
```

5. From deleting node deconfigure the clusterware.
Run as root user <GRID_HOME>/crs/install/rootcrs.sh -deconfig -force -verbose

6. From running nodes remove the clusterware configuration of deleting node.
Run as root user "crsctl delete node -n <deleting_node>

7. On deleting node update the Grid Home inventory.
```
$ORACLE_HOME/oui/bin/runInstaller -updateNodeList ORACLE_HOME=<GRID
HOME> "CLUSTER_NODES={<deleting_node>}" CRS=TRUE -local
```

8. Deinstall oracle Grid home using -local option run from deleting node.
```
$GRID_HOME/deinstall/deinstall -local
```

9. Update inventory in remaining nodes of cluster.
```
$ORACLE_HOME/oui/bin/runInstaller -updateNodeList ORACLE_HOME=<GRID
HOME> "CLUSTER_NODES={<node1>,<node2>}"
```

Refer to below link for more details regarding Addition and Deletion of nodes from Oracle RAC cluster
https://docs.oracle.com/database/121/CWADD/GUID-929C0CD9-9B67-45D6-B864-5ED3B47FE458.htm#CWADD1167

The test team noticed below known issue while adding a node in 12cR1 cluster and followed the workaround mentioned in the My Oracle Support (MOS) note,

Doc ID 2718587.1: root.sh failing with CLSRSC-293: Error: validation of OCR location 'NO_VAL' failed

The following Figure 1 shows the flow of migrating two node Oracle Real Application Cluster database from Power8 LPARs to Power10 LPARs, upgrading RAC into 19c, and upgrading AIX72 to AIX 7.3.

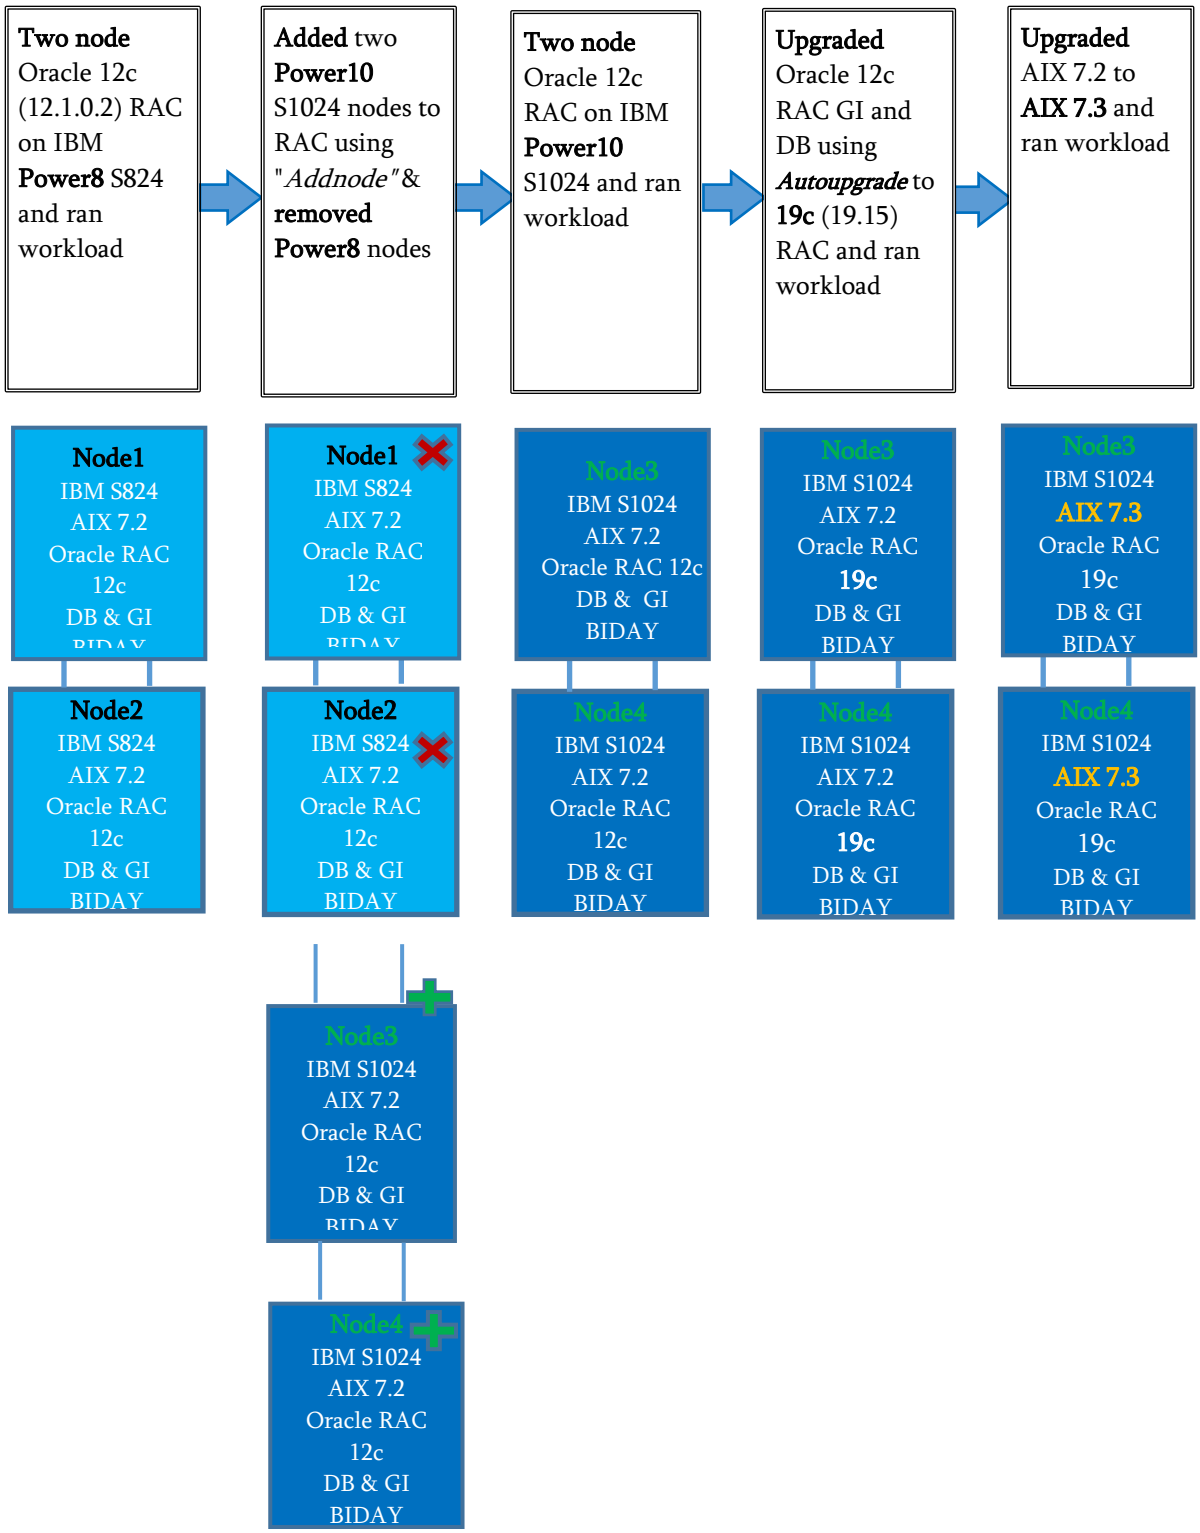

Figure 1. Oracle RAC Grid Infrastructure and Database Migration Flow diagram

After migrating the Oracle 12c R1 RAC into Power10 nodes, the baseline execution of validation tests was repeated, an upgrade to 19c RAC Gird Infrastructure (version 19.15) was completed, an upgrade to 19c RAC database using the Oracle Database *AutoUpgrade* utility was completed, and workload execution was resumed. When workload was again ramped up, the same high utilization was achieved, and the throughput was recorded.

The test team used latest version of autoupgrade.jar file and for more information regarding AutoUpgrade tool refer below link:
https://docs.oracle.com/en/database/oracle/oracle-database/19/upgrd/about-oracle-database-autoupgrade.html

Brokerage – OLTP workload

This OLTP workload simulated the transactions and database of a stock brokerage firm. The test team created a database size of approximately 1 TB using 50,000 customers. The team modified the transaction mix with a read/write ratio of approximately 90/10 so that more CPU operations can be performed.

For the Oracle 12cR1 RAC baseline, the 2-node RAC database was installed on two IBM Power S824 servers (which is based on IBM Power8 architecture) each with a logical partition (LPAR) using 24 cores. The LPAR placement of the cores and memory in Power8 was done to provide better alignment. When the team performed the runs, it was observed that the CPU utilization reached approximately 100%.

The following table provides the hardware and software configuration (returned by using the `prtconf` command). This configuration was used for the sample Brokerage Online Transaction Processing (OLTP) workload.

Configuration	2 x IBM Power8 (S824)	2 x Power10 (S1024)
System model	IBM,8286-42A	IBM,9105-42A
Processor type	PowerPC_POWER8	PowerPC_POWER10
Number of cores	24	12
Clock speed	3.5 GHz	3.6 to 3.99 GHz
Memory	500G	500G
Firmware version	FW860.B0 (SV860_240)	FW1020.00 (NL1020_067)
Sockets	2	1
Chips per socket	2	2
SMT (Default values used)	4	8
OS level	AIX 7.2 TL05 SP04	AIX 7.2 TL05 SP04 / AIX7.3 TL0 SP02
OS Patch	IJ41092	IJ41092 (IJ41094 AIX 7.3)
Oracle Grid Home Version	12.1.0.2.0 + PSU 33829718	12.1.0.2.0 + PSU 33829718, 19.15 RU
Oracle DB home version	12.1.0.2.0 + PSU 33829718	12.1.0.2.0 + PSU 33829718, 19.15 RU
Oracle one-off patch	NA	NA
Workload users	216	216

Table 1. Hardware and software configuration

Partition placement in Power servers for OLTP workload

The *lssrad* tool was used to get the partition placement details. On this Power8 processor-based system, the processors and memory spread across two sockets on each S824 as shown in Figure 3.

```
# lssrad -va
REF1   SRAD        MEM       CPU
0
          0   121740.44       0-23
          1   125704.44       24-47
1
          2   124749.00       48-71
          3   124998.00       72-95
#
```

```
# lssrad -va
REF1   SRAD        MEM       CPU
0
          0   121740.44       0-23
          1   125688.44       24-47
1
          2   124749.00       48-71
          3   124998.00       72-95
#
```

Figure 2. Output of lssrad command on the LPAR1 and LPAR2 in Power8

On the Power10, the required capacity was estimated as 12 cores per node, and this resized LPAR was able to fit on a single socket on each S1024 as shown in Figure 4.

```
# lssrad -va
REF1   SRAD        MEM       CPU
0
          0   248595.38       0-47
          1   248479.00       48-95
1
          2       0.00
#
```

```
# lssrad -va
REF1   SRAD        MEM       CPU
0
          0   248595.38       0-47
          1   248479.00       48-95
1
          2       0.00
#
```

Figure 3. Output of lssrad command on the LPAR1 and LPAR2 in Power10

Capacity comparison of Power10 with Power8 running OLTP workload

Initially the OLTP workload was executed on two Power8 S824 systems, where each LPAR having 24cores with SMT4. Later as mentioned in migration flow diagram Figure 1, using *addnode.sh* script added two Power10 LPARs each having 12cores with SMT8 and then deleted the running Power8 LPARs from RAC cluster. The OLTP workload was executed against 12cR1 RAC database which was running on Power10 nodes each having 12cores. It was observed that the 12 cores per node Power10 processor-based system with the same number of workload users as 24 cores per node Power8 processor-based system, reached CPU utilization of approximately 100%. Next, test team upgraded both Grid Infrastructure and database home stack to 19c from 12cR1 and executed workload on 19c two node RAC database running on Power10 LPARs. Below Figure 4 illustrates the system diagram containing both Power8 and Power10 servers.

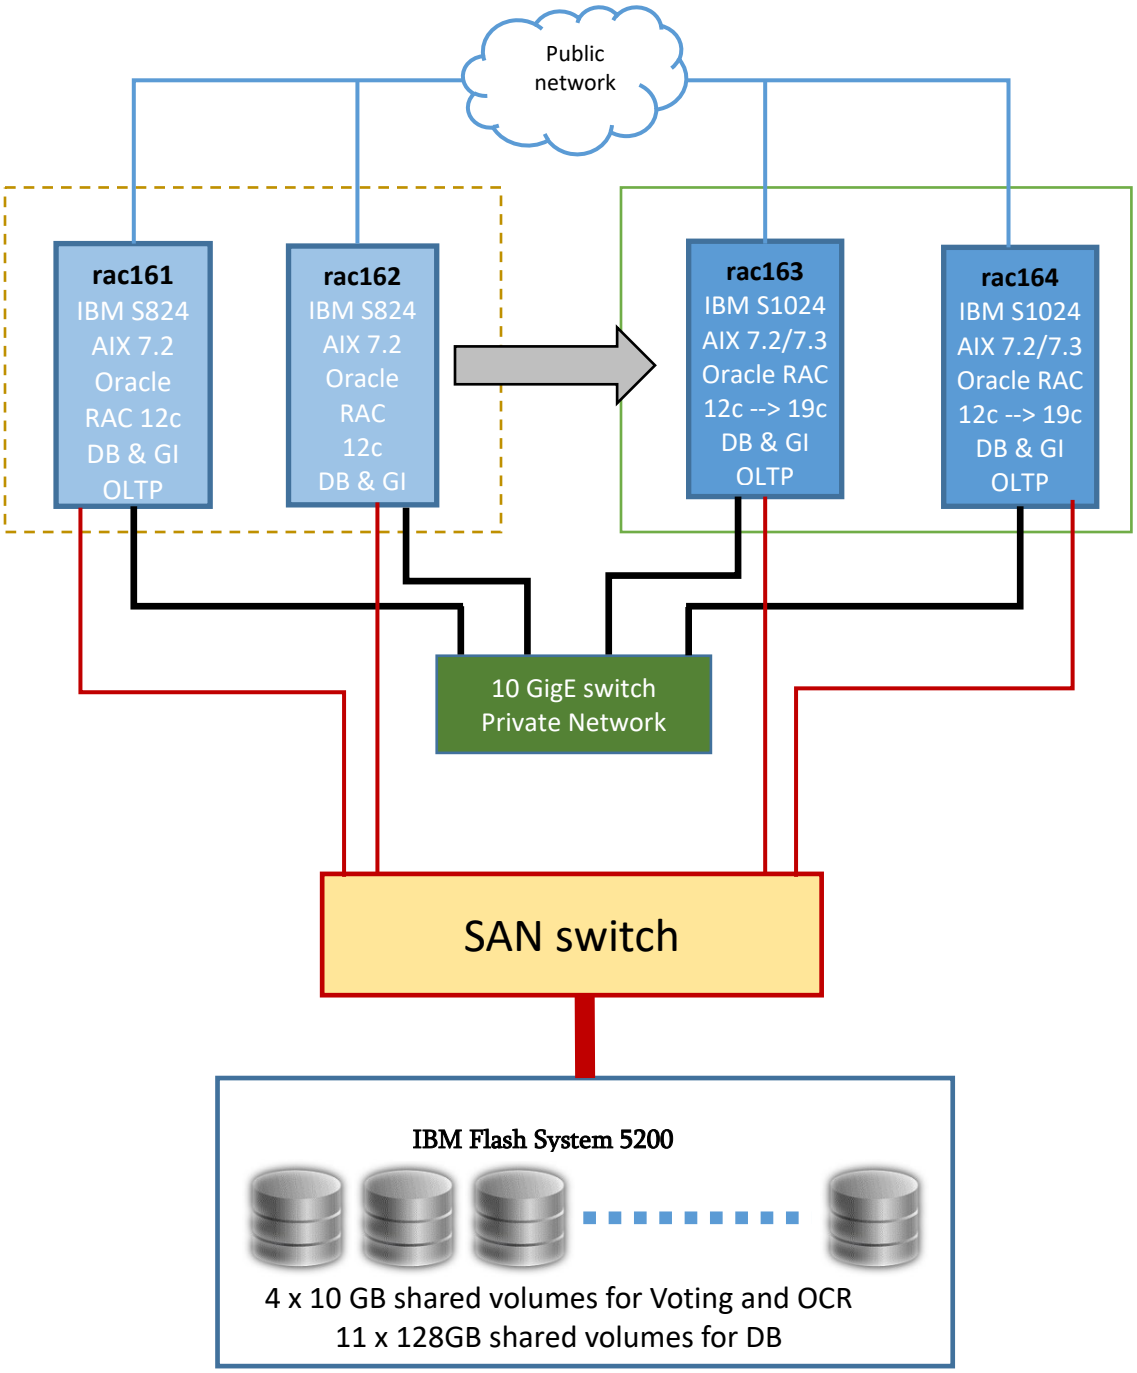

Figure 4. OLTP Workload System Setup

The Figure 5 shows the relative throughput (TPS) per core for the runs performed on Power8 and on Power10. The following observations can be made from Figure 5:

- The execution plans(plan_hash_value) for workload SQLs did not change when the test team moved the 12cR1 RAC Database storage volumes from Power8 to Power10. The improvement in throughput were observed when the test team restarted execution of 12cR1 on the Power10 processor-based systems.

- On Power10 partitions, some of the workload SQLs changed their execution plans when the team upgraded the database from 12cR1 to 19c. But this did not substantially change the throughput.
- After the upgrade, when the test team changed the database `compatible` parameter to `19.0.0` from `12.1.0.2.0` and made no other configuration changes, the execution plans of the workload did not change.

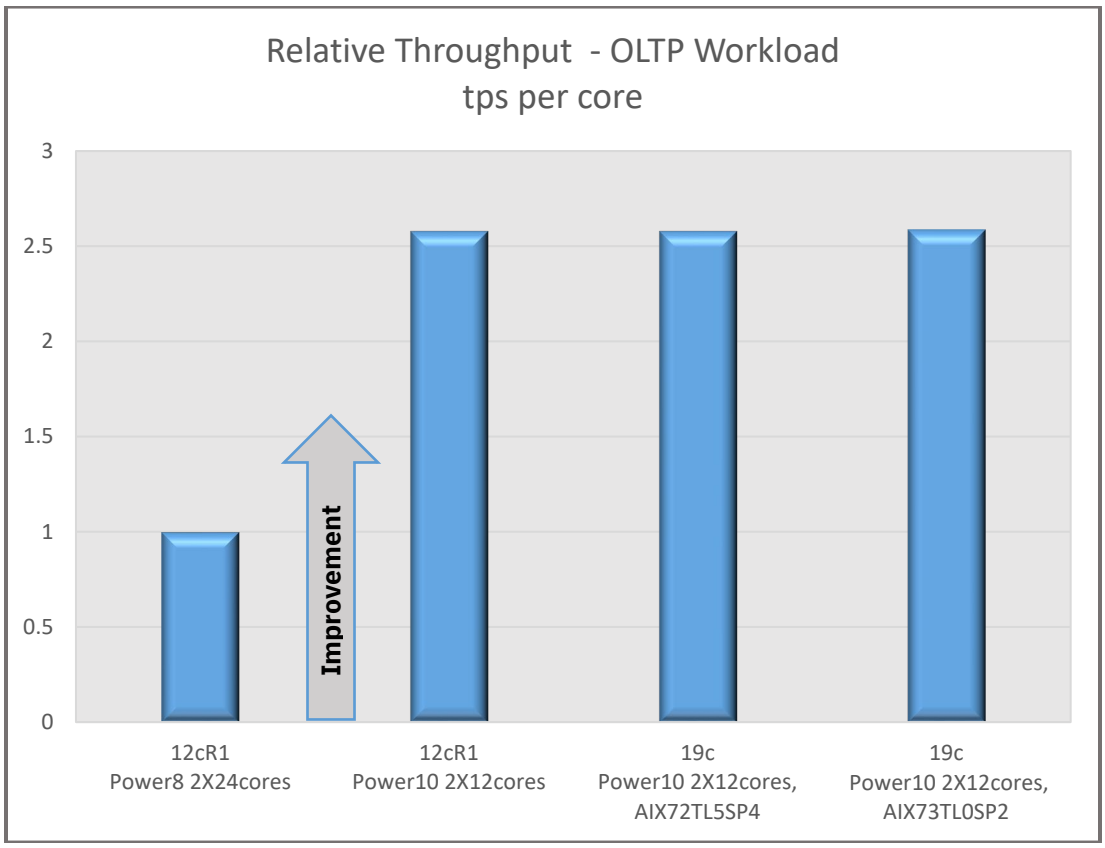

Figure 5. Throughput per core

It has been observed that the movement from Power8 to Power10 provided a benefit of ~2.6 times more throughput per core (delivered by Power10) and improvements delivered in both database versions 12cR1 and 19c. Generational improvements in the Power10 architecture increased the number of cores on a chip and improved per core throughput, allowing the partition to be placed on a single socket. This improved throughput and reduced the space and power requirements for the workload.

Benefits of AIX 7.3

Testing included upgrading to AIX 7.3, which 19c is certified to run on. We found that this upgrade did not significantly influence the capacity available, however it allows customers to make use of the additional features provided that can improve scalability, usability and tools. More details on these benefits are in the AIX 7.3 section.

BIDAY analytical workload

BIDAY is an IBM developed analytical workload that models different types of analysis that businesses can use in their business intelligence (BI) workload. The workload has one terabyte (1 TB) of raw data representing retail sales tracking loaded into Oracle Real Application Cluster Database.

The BIDAY workload consist of a set of 26 queries with complexity ranging from simple to very complex queries. The workload can be scaled by executing 1, 2, 4, 8, and more concurrent users. The BIDAY schema has one of the tables, named **Sales_Fact**, as a fact table which is single Range-Hash partitioned and filled with nine billion rows.

During the test, the set of all 26 queries were run sequentially in the same strict order and the team presented each concurrent execution of this set as an additional user. For example, for two users two copies of each query set are run, and for eight users eight copies of each query set are run concurrently. The execution of the queries is serialized, and therefore, all the users are executing the same query set. In the two node Oracle RAC database, the driver for running the queries was started from one of the RAC nodes.

The configuration (returned by using the `prtconf` command) shown in Table 2 was used with an IBM AIX LPARs for evaluating the capacity of the Oracle RAC Database on IBM Power servers.

Configuration	**2 x IBM Power8 (S824)**	**2 x IBM Power10 (S1024)**
System model	IBM, 8286-42A	IBM,9105-42A
Processor type	PowerPC_POWER8	PowerPC_POWER10
Number of cores	24 (dedicated) per system	12 (dedicated) per system
Clock speed or WOF* range	3.5 GHz	3.6 to 3.99 GHz
Memory	512 GB	512 GB
Firmware version	FW860.B0 (SV860_240)	FW1020.00 (NL1020_067)
Sockets	2	1
Chips per socket	2	2
SMT (default values used)	4	8
OS level	AIX 7.2 TL05 SP04	AIX 7.2 TL05 SP04 / AIX7.3 TL0 SP02
OS Patch	IJ41092	IJ41092 (IJ41094 AIX 7.3)
Oracle Grid Infrastructure	12.1.0.2.0 + PSU 33829718	12.1.0.2.0 + PSU 33829718, 19.15 RU
Oracle Database	12.1.0.2.0 + PSU 33829718	12.1.0.2.0 + PSU 33829718, 19.15 RU
Oracle one-off patch	NA	NA

Table 2. Configuration used with an IBM AIX LPAR for 2 node Oracle RAC database

The BIDAY workload queries with high-level complexity are CPU-intensive and can saturate 48 dedicated cores in Power8 and 24 dedicated cores in Power10 easily with four or more users concurrently running the queries.

12 dedicated cores with an LPAR per system in Power10 with a total of 24 cores is half the number of cores used in Power8 LPARs.

* WOF refers to Workload Optimized Frequency. IBM Power9 and Power10 processor-based scale-out and scale-up servers implement Workload Optimized Frequency as a new feature of the energy management (EnergyScale) technology.

Partition placement in Power servers for BIDAY queries

The CPU cores and memory resources assigned to the LPARs were aligned to be placed in a single socket in the Power10 processor-based server and in two sockets in a Power8 processor-based server. The assigned memory was almost equally shared to the cores. The output of the cores and memory assignment can be viewed using the following AIX command.

```
$ lssrad -av
REF1   SRAD        MEM       CPU
0
          0   127401.88       0-23
          1   127239.00       24-47
1
          2   127239.00       48-71
          3   127216.00       72-95
```

```
$ lssrad -av
REF1   SRAD        MEM       CPU
0
          0   127401.88       0-23
          1   127239.00       24-47
1
          2   127239.00       48-71
          3   127216.00       72-95
```

Figure 6. LPAR cores and memory placement in Power8 LPAR1 and LPAR2

```
# lssrad -av
REF1   SRAD        MEM       CPU
0
          0   254603.38       0-47
          1   254455.00       48-95
1
          2       0.00
```

```
# lssrad -av
REF1   SRAD        MEM       CPU
0
          0   254603.38       0-47
          1   254455.00       48-95
1
          2       0.00
```

Figure 7. LPAR cores and memory placement in Power10 LPAR1 and LPAR 2

The BIDAY workload queries ran on Oracle RAC Database version 12.1.0.2.0 with 1, 2, 4, and 8 users with the `parallel_degree_limit` DB parameter set to a number which matched the number of logical processors available in the LPAR. The `parallel_degree_policy` parameter was set to AUTO. For example, the LPAR in the IBM Power8 processor-based system was set up with 24 dedicated cores with SMT4 enabled. In total, 96 logical processors are available to the Oracle Database for executing the queries with up to 96 parallel processes per RAC node.

Next, the test team added two Power10 LPARs with the RAC using "*addnode*" command as mentioned in the earlier section of this document. After successfully adding two new nodes running on Power10 LPARs to the cluster, the nodes running on Power8 LPARs were removed from the cluster. Adding new Power10 nodes to the cluster and removing Power8 nodes from the cluster migrated the cluster completely to Power10 LPARs without down time.

Next, the BIDAY queries were run with 1, 2, 4, and 8 users and the team collected the elapsed time of each of the users. Because Power10 executes in the SMT8 mode by default, with the 12 dedicated Power10 cores the `parallel_degree_limit` DB parameter remained the same number 96 to which matched the logical processors available in the LPAR in the Power10 processor-based system. The `parallel_degree_policy` parameter remained to be set to AUTO as well.

The team installed Oracle Grid Infrastructure and Database 19c version 19.15.0.0.0 on another set of filesystems for an out-of-place upgrade from Oracle Database 12c (12.1.0.2.0) to 19.15.0.0.0. They used the Oracle Database AutoUpgrade tool with the latest *autoupgrade.jar* file. After successfully completing the Grid Infrastructure and database upgrade steps, the BIDAY workload queries were executed on the

Oracle Database version 19.15.0.0.0 with the same values assigned to the database parameters, `parallel_degree_limit=96` and `parallel_degree_policy=AUTO` and the total elapsed time from staring the set of queries to completion of them was collected.

The Oracle RAC setup for 2 nodes at a high level shown in the following diagram, Figure 8.

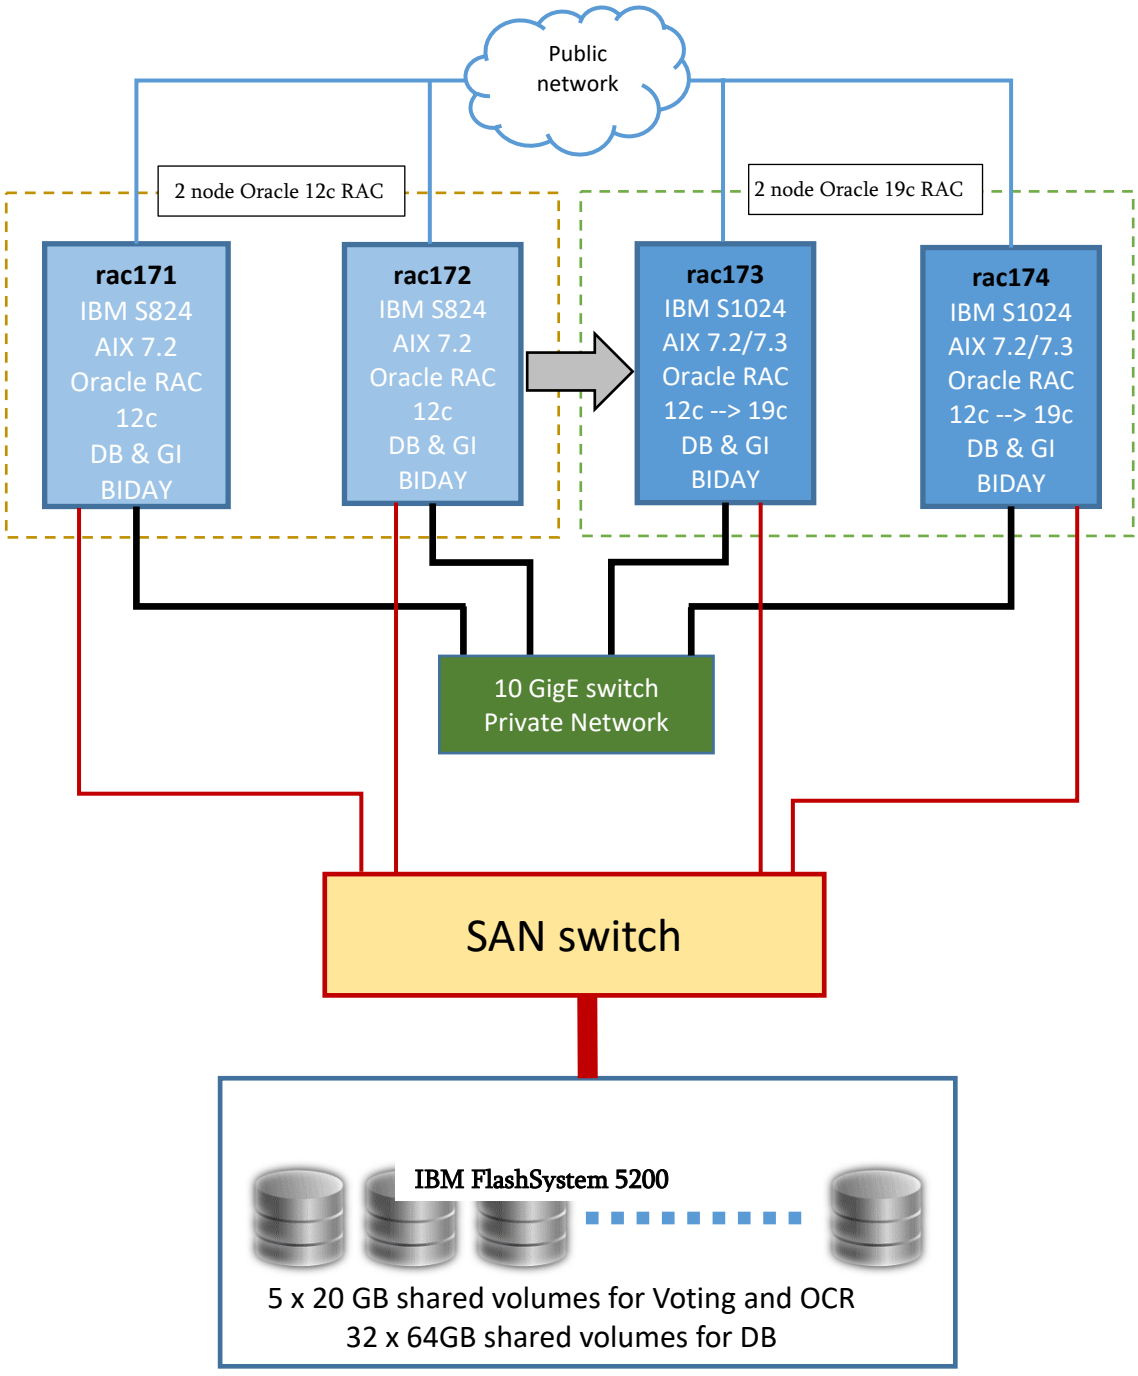

Figure 8. Oracle RAC setup diagram for BIDAY

The LPARs were interconnected with 2 x 10GigE network interfaces with a private 10 GigE network switch. The SAN connectivity from each of the LPARs to the IBM FlashSystem 5200 can provide a theoretical bandwidth up to 6.4 GB/sec. Each Power8 node used four 16 Gigabit Fibre Channel connections, and each Power10 node used two 32 Gigabit Fibre Channel connections.

Capacity comparison of Power10 with Power8 running BIDAY queries

The IBM Power10 processor-based system with 24 cores provided more capacity than the Power8 processor-based system with 48 cores for running business analytics type of queries with the following results.

Capacity metrics	Power8 48 cores (24c + 24c) Oracle DB 12c R1 RAC (12.1.0.2.0) AIX7.2	Power10 24 cores (12c +12c) Oracle DB 12c R1 RAC (12.1.0.2.0) AIX7.2	Power10 24 cores (12c +12c) Oracle DB 19c (19.15) RAC AIX7.2	Power10 24 cores (12c +12c) Oracle DB 19c (19.15) RAC AIX7.3	Power10 24 cores (12c +12c) Oracle DB 19c (19.15) RAC In-Memory AIX7.2	Power10 24 cores (12c +12c) Oracle DB 19c (19.15) RAC In-Memory AIX7.3
Relative per core improvement factor	1.0x	2.6x	3.6x	3.6x	132.4x	140.9x

Table 3. Capacity comparison of Power10 with Power8

CPU utilization and I/O throughput

The result shown in Table 3 was obtained for the BIDAY queries ran on row format data by eight users, and an average value of 96% and a maximum value of 100% CPU utilization was observed. I/O bottlenecks were not observed in the SAN volumes and adapters. The queries ran on In-Memory format data by 8 users consumed an average value of 67% and a maximum value of 100% CPU cycles.

The SAN connectivity between the AIX LPARs and the IBM FlashSystem 5200 had enough bandwidth to read the data for processing. The following I/O throughput (as shown in Table 4) was observed while executing the queries on row format data. The higher data rates (storage I/O) of Oracle Database 19.15 are reflected in the improved query time compared to Oracle Database 12c R1. The optimizer in the Oracle Database 19.15 was using different execution plan for some of the complex queries compared to Oracle Database 12c R1.

System	Total number of cores in the LPAR	Database version	Relative average Storage I/O Bandwidth
Power8 processor-based	48	12c R1 (12.1.0.2.0)	1.0x
Power10 processor-based	24	12c R1 (12.1.0.2.0)	1.3x
Power10 processor-based	24	19c (19.15.0.0.0)	8.0x

Table 4. I/O throughput

Oracle Database In-Memory feature for BIDAY queries

The BIDAY workload can leverage Oracle Database 19c In-Memory feature which has been enhanced since 12cR1. With this option and 19c, it is also possible to run with native data types such as BINARY_FLOAT,

which was used. When executed with the In-Memory feature, the schema loads the nine billion row table, named **Sales_Fact**, into the In-Memory area with *priority critical DISTRIBUTE BY ROWID RANGE* and the default compression level. While the fact table was loaded into the In-Memory area, data was distributed among two nodes almost with the rate of 50%. The entire Sales_Fact table with an original size of 726 GB was compressed to less than 200 GB. This means the real memory for the table, and some additional metadata was contained in the resized *SGA_TARGET* of 400 GB.

While executing the queries with the In-Memory feature enabled, no significant SAN I/O throughput was observed.

While comparing the effective capacity of the scenarios, a metric which takes into account the improvement in core capacity as well as improvements in query elapsed time is required. For this purpose, the test team used a metric computed as *(cores x queries elapsed time in seconds) / users*.

Each Power10 LPARs was used with 12 dedicated cores and the each Power8 LPARs was used with 24 dedicated cores. The Power10 processor-based system was observed to improve the BIDAY queries up to 2.6 times for row type queries with the Oracle 12cR1 RAC database. The capacity increase from Power8 to Power10 when moving to the 19c database was 3.6 times for row type queries and 141 times for InMemory type queries when normalized to a per core basis.

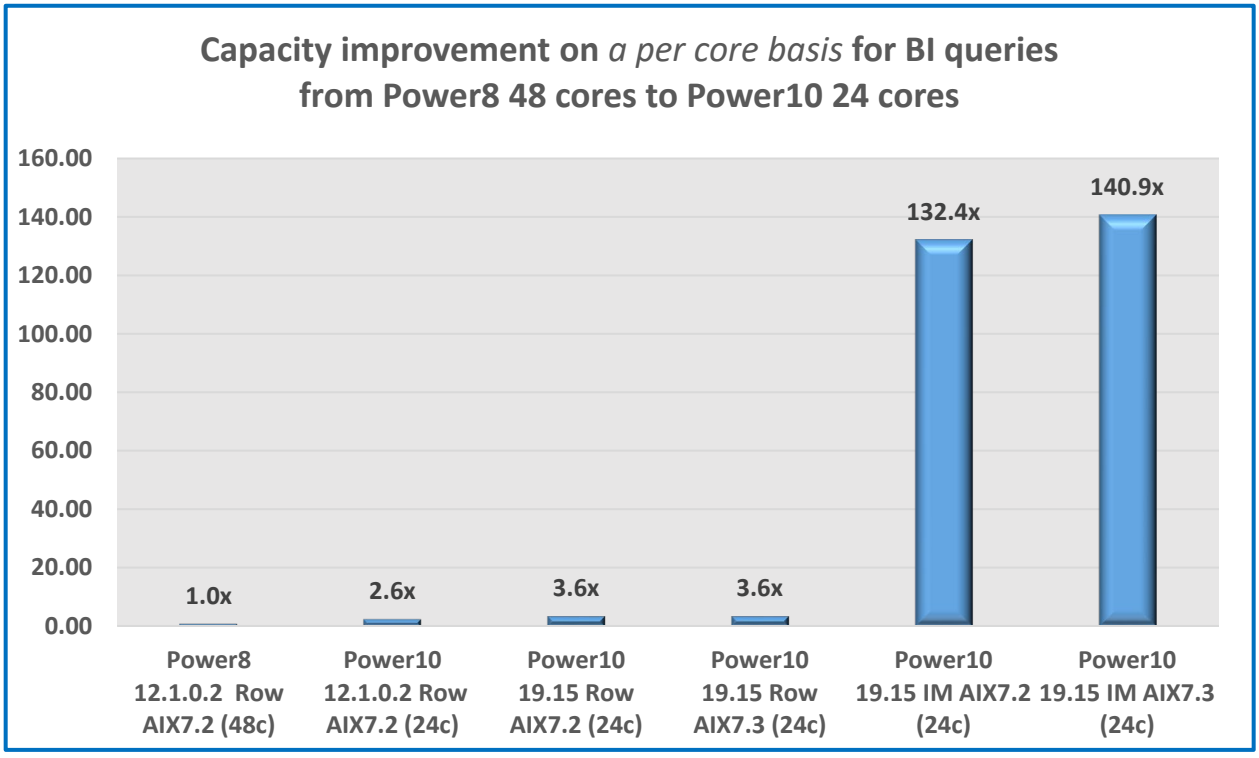

Figure 10. Capacity changes on a per core basis from Oracle DB 12cR1 to 19.15

Oracle Database and AIX tuning parameters for BIDAY

The Oracle Database server instance was not tuned for the best possible capacity. Most of the parameters hold the database default values except the following parameters:

```
PARALLEL_DEGREE_LIMIT = <number of logical threads available in the
LPAR>
PARALLEL_DEGREE_POLICY = AUTO
SGA_TARGET=260G
```

In-Memory only:

```
SGA_TAREGT=400G
INMEMORY_SIZE=150G
INMEMORY_OPTIMIZED_ARITHMETIC=ENABLE
```

For more detailed information about DB parameters and their values, refer to the “Appendix” section.

Benefits of AIX 7.3

Testing included upgrading to AIX 7.3, which 19c is certified to run on. When making this upgrade we also changed the from 12cR1 to 19c compatibility. This change allowed the optimizer to select better methods for some of the queries resulting in additional improvements to capacity. In addition, this allows customers to make use of the additional features provided that can improve scalability, usability and tools. More details on these benefits are in the AIX 7.3 section.

BIDAY workload summary

The BI workload (BIDAY) queries running with two node Oracle RAC Database version 19.15 on the newly introduced IBM Power10 processor-based systems with a total of 24 cores shows 3.6 times more capacity on a per core basis for row type data, and 141 times improvement with In-Memory data compared with the result of row type queries on IBM Power8 processor-based systems for a total of 48 cores.

Summary

This paper explained how proper planning can help users of earlier Oracle database versions to update to the current version on the new IBM Power10 processor-based server with minimal disruption. The steps to accomplish includes:

- Updating the Oracle database release to the current patch set level.
- Migrating the Oracle database Real Application Cluster from Power8 to Power10 LPARs
- Using Oracle tools that provide compatibility, including *AutoUpgrade*.
- Updating and upgrading the AIX operating system and firmware to the current levels.
- Adding Oracle 12cR1 and 19c features including Database In-Memory incrementally.

Benefits of AIX 7.3

- Up to 240 cores and 1920 HW threads per LPAR on Power10
- Up to 128 TB JFS2 filesystem and file size
- Up to 2X increased asyncIO IOPs scaling
- Improved fork()/exec() scaling
- Ready for the new IBM OpenXL compilers
- Enabled for SW exploitation of Power10 MMA
- Power10 optimized memcpy()
- TCP CUBIC support for improved performance with high latency networks
- ”Out of the box” ready for HW GZIP with all components installed by default
- Ready for Ansible with Python 3.9
- bash 5.1 as an available shell
- pigz compression utility with Power HW gzip compression
- Stronger default security policy, including 255-character passwords
- Reduced boot time for multi-terrabyte lpars
- Enhanced Dynamic LPAR performance for adding memory and compute to live running LPARs
- AIX Active Memory Expansion now defaults to 64KB pages on Power10 systems

Additional information can be found at: https://www.ibm.com/docs/en/aix/7.3?topic=whats-new

Appendix

Oracle Database parameters used for BIDAY workload

The parameter values for Oracle Database 12.1.0.2.0 and 19.15 are same for the row type except `<instancename>.thread` and `<instancename>.undo_tablespace` parameters.

```
biday1.__oracle_base='/u01/base'#ORACLE_BASE set from environment
biday2.__oracle_base='/u01/base'#ORACLE_BASE set from environment
*.audit_file_dest='/u01/base/admin/biday/adump'
*.audit_trail='db'
*.cluster_database=true
*.compatible='12.1.0.2.0'
*.control_files='+BIDAYDATA/BIDAY/CONTROLFILE/current.263.1110039509'
*.db_block_size=8192
*.db_create_file_dest='+BIDAYDATA'
*.db_name='biday'
*.diagnostic_dest='/u01/base'
*.dispatchers='(PROTOCOL=TCP) (SERVICE=bidayXDB)'
biday2.instance_number=2
biday1.instance_number=1
*.open_cursors=400
*.optimizer_adaptive_reporting_only=TRUE
*.optimizer_dynamic_sampling=0
*.parallel_degree_limit='96'
*.parallel_degree_policy='AUTO'
*.parallel_force_local=FALSE
*.parallel_max_servers=3000
*.parallel_min_servers=512
*.parallel_threads_per_cpu=1
*.pga_aggregate_target=60g
*.processes=8960
*.recyclebin='OFF'
*.remote_login_passwordfile='exclusive'
*.resource_manager_plan=''
*.sga_target=260g
*.star_transformation_enabled='TRUE'
biday2.thread=2
biday1.thread=1
biday1.undo_tablespace='UNDOTBS3'
biday2.undo_tablespace='UNDOTBS4'
```

The Database In-Memory parameters for the Oracle Database version 19.15.0.0.0,
For the Oracle database 19c with AIX7.2, the compatible parameter was unchanged, and it was set to 12.1.0.2.0. After the AIX7.2 upgraded to AIX7.3, the compatible parameter was set to 19.0.0.

```
biday3.__oracle_base='/u01/base'#ORACLE_BASE set from environment
biday4.__oracle_base='/u01/base'#ORACLE_BASE set from environment
*.audit_file_dest='/u01/base/admin/biday/adump'
*.audit_trail='db'
*.cluster_database=true
*.compatible='12.1.0.2.0'
*.control_files='+BIDAYDATA/BIDAY/CONTROLFILE/current.263.1110039509'
*.db_block_size=8192
*.db_create_file_dest='+BIDAYDATA'
*.db_name='biday'
*.diagnostic_dest='/u01/base'
*.dispatchers='(PROTOCOL=TCP) (SERVICE=bidayXDB)'
*.inmemory_max_populate_servers=128
*.inmemory_size=161061273600
*.inmemory_optimized_arithmetic='ENABLE'
biday3.instance_number=3
biday4.instance_number=4
*.nls_language='AMERICAN'
*.nls_territory='AMERICA'
*.open_cursors=400
*.optimizer_adaptive_reporting_only=TRUE
*.optimizer_dynamic_sampling=0
*.parallel_degree_limit='96'
*.parallel_degree_policy='AUTO'
*.parallel_force_local=FALSE
*.parallel_max_servers=4000
*.parallel_min_servers=512
*.parallel_threads_per_cpu=1
*.pga_aggregate_target=60g
*.processes=8960
*.recyclebin='OFF'
*.remote_login_passwordfile='exclusive'
*.sga_target=400g
*.star_transformation_enabled='TRUE'
biday3.thread=3
biday4.thread=4
biday3.undo_tablespace='UNDOTBS5'
biday4.undo_tablespace='UNDOTBS6'
```

Oracle Database parameters used for Brokerage OTLP workload

The parameter values for Oracle Database 12cR1 and 19.15 are the same except that for 12cR1 the `compatible` parameter is set to `12.1.0.2.0`.

```
*.audit_file_dest='/u01/base/admin/tpce/adump'
*.audit_trail='db'
*.cluster_database=true
*.compatible='19.0.0'
*.control_files='+DATATPCE/TPCE/CONTROLFILE/current.257.1109924369'
*.db_block_size=8192
*.db_create_file_dest='+DATATPCE'
*.db_domain=''
*.db_name='tpce'
*.diagnostic_dest='/u01/base'
*.dispatchers='(PROTOCOL=TCP) (SERVICE=tpceXDB)'
tpce2.instance_number=2
tpce1.instance_number=1
*.open_cursors=300
*.parallel_max_servers=600
*.parallel_min_servers=8
*.pga_aggregate_target=32212254720
*.processes=5000
*.remote_login_passwordfile='exclusive'
*.resource_manager_plan=
*.sga_target=429496729600
tpce2.thread=2
```

Get more information

For more information on using Oracle Database on IBM Systems, you can contact ibmoracle@us.ibm.com

For more information on using Oracle AutoUpgrade, refer to the following websites:

- Learn how to use AutoUpgrade to simplify your upgrade tasks
  https://docs.oracle.com/en/database/oracle/oracle-database/19/upgrd/using-autoupgrade-oracle-database-upgrades.html

- Database upgrade guide
  https://docs.oracle.com/en/database/oracle/oracle-database/19/upgrd/index.html

Acknowledgments

IBM acknowledges Oracle engineering staff who provided guidance on the development and review of the white paper.

About the authors

Wayne Martin is the IBM Systems Technology Solutions Manager responsible for the technology relationship between IBM and the developers of Oracle Corporation's Database and Fusion Middleware for all IBM server brands. His responsibilities include driving the mutual understanding between IBM and Oracle on technology innovations that will generate benefits for mutual customers, managing the process of getting that technology implemented in products and ensuring that availability of the products to customers is timely. Wayne has held a variety of technical and management roles at IBM that have focused on driving enhancements of ISV software that uses IBM's mainframe, workstation, and scalable parallel products.

Bhargavaram Akula is a Technical Consultant with IBM India, Hyderabad. He collaborates with the specialists at the IBM Oracle International Competency Center based in Foster City and Redwood Shores (California, US) working on Oracle product certifications on IBM Power running AIX.
He has extensive experience with Oracle products.

Dennis Massanari is a Senior Software Engineer with the ISV Technical Enablement team in IBM Systems. He focuses on IBM Power running Oracle Database.

Bret Olszewski is an experienced Senior Technical Staff Member working on technical enablement programs for IBM Power. Bret has long experience on the Power architecture, with particular emphasis on performance of Enterprise software and overall systems design. Bret has a BSCS degree from the University of Minnesota.

Ravisankar Shanmugam is an Advisory Software Engineer with the ISV Technical Enablement team in the IBM Systems group based in Foster City, California. He writes technical collateral, leading proof of concept, and performance evaluation projects for IBM Power servers with Oracle products.

IBM Systems
3039 Cornwallis Road
RTP, NC 27709

Produced in the United States of America

Please recycle