



SAP on Exadata X8M-2



SAP ECC with Oracle Database on Exadata X8M-2 POC Results.

Carlos Eduardo Benetti – Principal Sales Consultant
André Mercanti – Master Principal Sales Consultant
April 15, 2020 | Version 1.1
Copyright © 2020, Oracle and/or its affiliates

PURPOSE STATEMENT

This document provides an overview of a POC with SAP ECC 6.0 EHP4 on a real production workload, the features and enhancements included in Exadata X8M-2 with PMEM for Oracle Database 19c. It is intended solely to help you assess the business benefits of it on a SAP landscape and to plan your I.T. projects.

DISCLAIMER

This document in any form, software or printed matter, contains proprietary information that is the exclusive property of Oracle. Your access to and use of this confidential material is subject to the terms and conditions of your Oracle software license and service agreement, which has been executed and with which you agree to comply. This document and information contained herein may not be disclosed, copied, reproduced or distributed to anyone outside Oracle without prior written consent of Oracle. This document is not part of your license agreement nor can it be incorporated into any contractual agreement with Oracle or its subsidiaries or affiliates.

This document is for informational purposes only and is intended solely to assist you in planning for the implementation and upgrade of the product features described. 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 in this document remains at the sole discretion of Oracle.

Due to the nature of the product architecture, it may not be possible to safely include all features described in this document without risking significant destabilization of the code.

TABLE OF CONTENTS

- Source System Information: 3
- Source Architecture: 4
- Target POC Information: 4
- POC Architecture: 5
- POC Target Needs: 6
- POC Target Achieved: 7
 - Step 1 - Preupgrade step 11
 - Step 2 – Run Pre Upgrade Fixups 11
 - Step 4 – Execute de upgrade command 12
 - Step 5 – Execute Post Upgrade Fixups. 12
 - Step 6 – After complete, check logs on directory to check the upgrade results. 12

INTRODUCTION

Exadata helps to eliminate common bottlenecks found in traditional architectures due to its integrated hardware and software development. In addition to this, we have in the Exadata X8M-2 two new features that help SAP workloads in Oracle Database to perform at the maximum in performance: a third level of tierization that allows faster access to data with performance compatible with memory and an interconnection of all components in a 100Gb/s network.

In a system like SAP ECC, where we can find types of mixed workloads, this architecture provided by Exadata is essential to keep both jobs executed in a performatic batch chain, as well as maintaining user transactions with a satisfactory response time to the various business needs.

In this sense, we work on a test involving a workload and mass of real data, associated with a test scenario close to the main needs listed by one of our customers.

GENERAL LANDSCAPE INFORMATION

The POC focus was the performance of the database in order to eliminate several performance problems that the client's current architecture had been facing.

But for that, we needed to simulate the same type of processing power in the application layer with the same capacity to receive requests in the same proportion as the production environment.

Below we have the traditional architecture adopted by the client today:

Source System Information:


- ✓ SAP ECC 6.0 EHP4 / Netweaver 7.01
- ✓ Oracle Database 11.2.0.4 with RAC and ASM
- ✓ Kernel 721 patch level 600
- ✓ 7 Physical Application Servers
- ✓ 2 Physical Database Instances
- ✓ Red Hat Enterprise Linux 6
- ✓ Database Size: 20TB

The diagram illustrates the source system architecture, divided into two main layers: the Application Layer and the Database Layer. The Application Layer consists of three server models: PowerEdge M630, PowerEdge M620, and ProLiant DL360 Gen9. The Database Layer consists of ProLiant DL580 Gen9 servers. Performance metrics for SAP are provided for both layers.

Server Model	Processor	Cores	vCPUs	RAM	Layer	SAP Performance
PowerEdge M630	Xeon E5-2660 v3	20	40	256GB	Application Layer 128 Cores – 256 vCPUs – 1.792TB RAM ~373.873 SAPS	
PowerEdge M620	Xeon E5-2670 v2	20	40	256GB		
ProLiant DL360 Gen9	Xeon E5-2637 v4	8	16	256GB		
ProLiant DL580 Gen9	Xeon E7-8867 v4	72	144	512GB	Database Layer 144 Cores – 288 vCPUs – 1.024TB RAM ~427.552 SAPS	

The resources are divided into 2 different sites to maintain the availability required by the business. The application and database servers are distributed among the 2, so that if one of the sites is unavailable, the other will assume the demand:

Source Architecture:



POC LANDSCAPE INFORMATION

To perform the POC, we used the Exadata X8M-2 Quarter Rack for the Database layer and the integrated server platform Private Cloud Appliance (PCA), both available at the Oracle Solution Center in São Paulo.

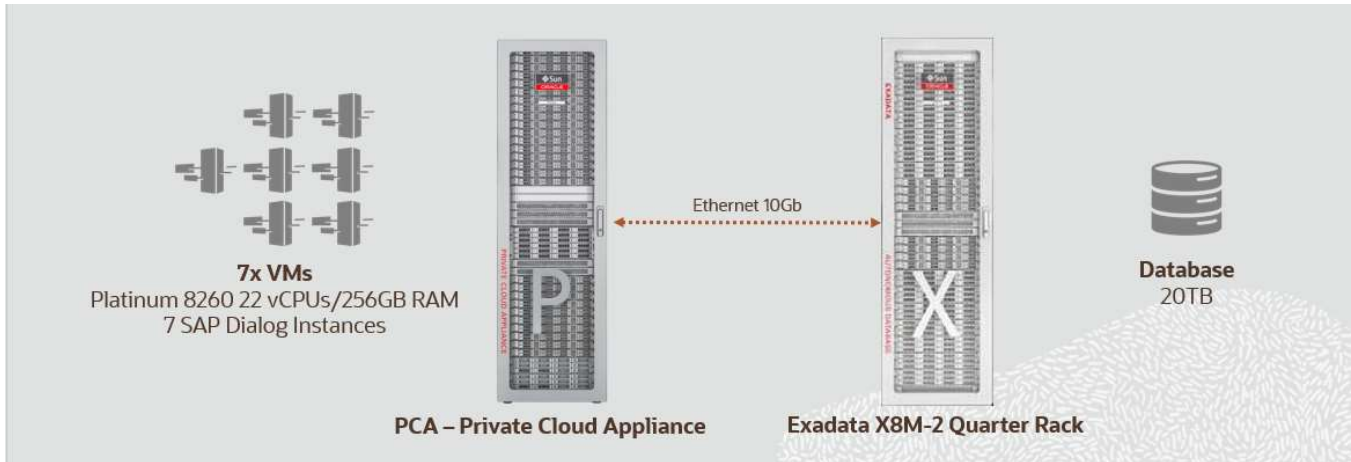
At PCA we created 7 machines with the same number of vCPUs and memory available in the original environment. To accommodate the client's 20TB of DB backup and the installation media needed to recreate the application servers, we use an external disk area provided by the ZS7-2 storage, all connected to a 10Gb/s Ethernet network.

Target POC Information:

- ✓ SAP ECC 6.0 EHP4 / Netweaver 7.01
- ✓ Oracle Database 11.2.0.4 and 19c with RAC and ASM
- ✓ Kernel 721_EXT patch 1300
- ✓ 7 VM Application Servers
- ✓ 2 Physical Database Instances
- ✓ Oracle Linux 7
- ✓ Database Size: 20TB



POC Architecture:



MAPPING AND EXECUTION PLANNING

The most important thing in the test is an adequate survey of what should be tested, how it should be tested, the steps to reproduce them and the expected achievement target.

Basically, three test pillars were listed based on critical business processes that need improvement and which today has an important impact on the other production chains that are part of a daily, weekly or monthly routine.

Night Batch Chain:

- ✓ Most important and complex job, dialog and RFC processes, macro divided as:
- ✓ Credit Analysis - 7:00 pm to 8:00 pm: A maximum limit of 90 dynamically scheduled jobs
- ✓ Pre-Allocation - 8:00 pm to 9:00 pm: Sequence of more than 100 Jobs in parallel, divided into 2 distribution centers, broken down by brand and retail
- ✓ Second Phase Allocation and WM - 22:30 to 03: 00hs: Programs that at times create parallelism in dialog
- ✓ Boxing Process - 19:00 to 03: 00hs

CO/PA Apportionment:

- ✓ Standard program scheduled as a sequential job with 2 different cycles

SD Apportionment:

- ✓ Customized program that checks the data to be processed and distributes parallel Jobs on Application Servers, making some dialog access

To set up the entire environment and reproduce the test scenarios, we adopted the following strategy:


Macro Execution Plain:

- ✓ Requisition of the necessary equipment in the Oracle Solution Center
- ✓ Provide all necessary documentation
- ✓ Copy the Database RMAN Backup on the stage area
- ✓ Copy the necessary SAP medias on the stage area
- ✓ Creating Database Instance on Exadata based on the RMAN backup
- ✓ Apply SAP Bunddle Patches and SAP Notes 1020260 / 1171650
- ✓ Collect statistics, including dictionary and system statistics
- ✓ Installing ASCS Instance and additional Dialog Instances on the App Servers
- ✓ Equalization of SAP Instance Profiles according original system ECP
- ✓ Execution of the first round tests on DB version 11.2.0.4
- ✓ Execution of the second round tests on DB version 11.2.0.4
- ✓ Upgrade Database to Oracle 19c
- ✓ Execution of the first round tests on DB version 19c
- ✓ Execution of the second round tests on DB version 19c
- ✓ Collect the necessary information and evidence
- ✓ Present Results

TARGET AND ACHIEVEMENTS

As a target for this POC, we had a goal of achieving as below.

POC Target Needs:



A 50% reduction in execution times is a difficult task to be performed since only the Database layer was in question.

POC Target Achieved:

Before talking about the results, it is important to highlight the amount of SAPS that we use in the POC versus that delivered by the original environment. Adding the amount of SAPS provisioned by the Application Servers in the PCA and Exadata's DB Nodes, we had, on average, 30% less computational capacity to carry out the necessary jobs and transactions.

SAPS Comparative – Original Environment x POC Environment



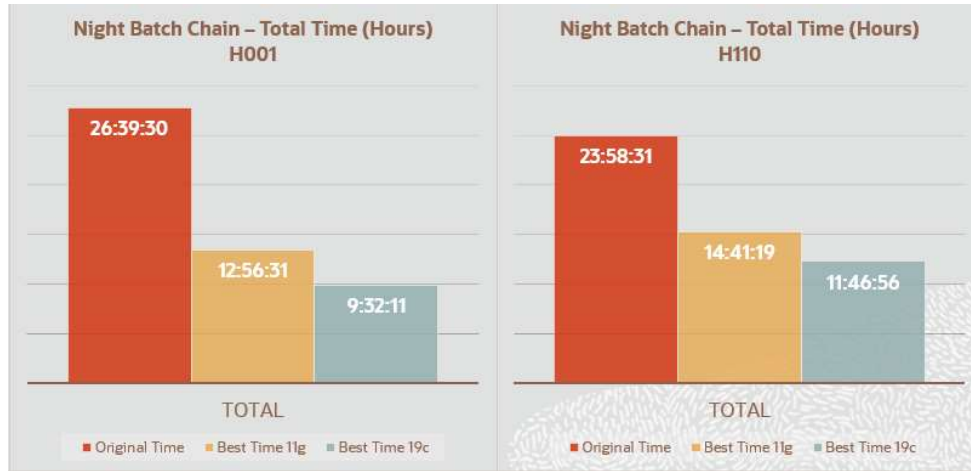
Even with less SAPS, we achieved a good result and reached the expected reduction goal in the execution times of the processes. In the 3 pillars of tests carried out, we achieved considerable results that in general can be seen in the graphs below.

Best Times Achieved – General



Below is the details of each test pillar.

Best Times Achieved – Night Batch Chain




Best Times Achieved – Night Batch Chain per Process



Best Times Achieved – Apportionments




Exadata Features on Workload – Oracle Database 19c



Some issues are important and should be addressed as why the test results on Oracle Database 11g are different from those obtained with Oracle Database 19c and because the Exadata X8M-2 behaved differently. Basically, with the use of Oracle Database 19c on Exadata X8M-2 we have the use of a third tier of information storage, in front of flash and behind the DB Nodes memory based on the new technology of Persistent Memory Data & Commit Accelerator that make use of Intel® Optane modules, associated with RDMA over Converged Ethernet (RoCE) 100GbE.

Exadata Storage Servers transparently add Persistent Memory Accelerator in front of flash memory and the Database uses RDMA instead of I/O to read remote PMEM, which is automatically tiered and shared across databases Using as a cache for hottest data.



This same technology also allows us to lowering log write latency, that is critical for OLTP performance. Fast log write gives fast transaction commit time and any log write slowdown causes commit backlog. Automatic Commit Accelerator allow databases a “one-way” RDMA writes to PMEM on multiple Storage Servers, bypasses network and I/O software, interrupts, context switches, etc.

Below we can see this information on AWR:

Cache Savings

- Disk write savings (overwrites) - writes absorbed by flash cache that would have otherwise gone to disk
- Database Cache Hit% - for the database, not restricted to an instance
- Cell Cache Hit% - for the cells, not restricted to this database or instance
- Cell Memory Cache Hit% - for the cells, not restricted to this database or instance

Database Flash Cache Hit %	1.2
Database Flash Cache Temp Read Hit %	95.84
Database Flash Cache Temp Write Hit %	100.00
Database Memory Cache Hit%	
Database PMEM Cache Hit%	99.27
Database PMEM Cache RDMA Hit%	91.19
Cell Flash Cache OLTP Hit %	96.45
Cell Flash Cache Scan Hit %	100.00
Cell Memory Cache Hit%	
Cell PMEM Cache Hit%	99.70
Disk Write savings/s	122.28
Large Writes/s	250.90

Smart IO

- These statistics are collected by the cells and are not restricted to this database or instance
- Storage Index - bytes saved by storage index and percentage of requested bytes saved by storage index
- Flash Cache - bytes read from flash cache and percentage of requested bytes read from flash cache
- Offload - bytes processed by the cells and not returned to the database
- Passthru - bytes returned as-is to the database (for reasons other than high cell cpu) and percentage of requested bytes returned as-is to the database
- Reverse Offload - bytes returned as-is to the database due to high cell cpu and percentage of requested bytes returned as-is to the database
- Ordered by Total MB Requested desc

Cell Name	MB Requested			Storage Index		Flash Cache		Offload		Passthru		Reverse Offload	
	% Total	Total	per Sec	MB	% Optimized	MB	% Optimized	MB	% Efficiency	MB	% Passthru	MB	% ReverseOffload
Total (4)		294,203.38	81.36	80,103.95	27.23	214,159.77	72.79	264,052.90	89.75				
ed01celadm01	25.51	75,065.18	20.76	20,470.29	27.27	54,417.62	72.49	67,312.74	89.67				
ed01celadm04	25.23	74,231.85	20.53	20,141.70	27.13	54,026.21	72.78	66,756.22	89.93				
ed01celadm02	24.68	72,599.50	20.08	19,769.19	27.23	53,219.18	73.31	65,117.05	89.69				
ed01celadm03	24.58	72,306.84	20.00	19,722.77	27.28	52,496.76	72.60	64,866.89	89.71				

PMEM Cache Space Usage

- These statistics are collected by the cells and are not restricted to this database or instance
- Space is at the time of the end snapshot
- Ordered by Space (GB) desc

Cell Name	Space Allocated	% OLTP	
		%Clean	%Dirty
Total (4)	3.73TB	99.99	
ed01celadm02	956.60GB	99.99	
ed01celadm01	956.43GB	99.99	
ed01celadm03	955.50GB	99.99	
ed01celadm04	954.71GB	99.99	

PMEM Cache Internal Writes

- These statistics are collected by the cells and are not restricted to this database or instance

Cell Name	Population Write Requests			Population Write MB		
	Total	per Sec		Total	per Sec	
Total (4)	233,768		64.65	3,891.55		1.08
ed01celadm03	66,029		18.26	1,060.96		0.29
ed01celadm01	58,139		16.08	1,007.21		0.28
ed01celadm02	55,111		15.24	903.10		0.25
ed01celadm04	54,489		15.07	920.27		0.25

PMEM Cache User Reads

- These statistics are collected by the cells and are not restricted to this database or instance
- ordered by hits desc

Cell Name	Read Requests						Read MB			
	Hits	Hits per Sec	%Hit	Misses	Misses per Sec	Hits MB	Hits MB/s	Misses MB	Misses MB/s	
Total (4)	3,262,657	902.28	99.70	9,695	2.68	312,741.50	86.49			
ed01celadm02	959,403	265.32	99.76	2,268	0.63	79,445.98	21.97			
ed01celadm03	798,391	220.79	99.72	2,231	0.62	77,458.57	21.42			
ed01celadm04	790,927	218.73	99.69	2,457	0.68	78,555.52	21.72			
ed01celadm01	713,936	197.44	99.62	2,739	0.76	77,281.43	21.37			

All this optimization can be seen in the latency highlighted below:

Top Databases By Requests - Details

- Request details for the top databases by IO requests

DB Name	DBID	IOs/s	Small Requests						Large Requests									
			Reqs/s			Latency		Queue Time		Reqs/s			Latency		Queue Time			
			Total	Flash	Disk	Flash	Disk	Flash	Disk	Total	Flash	Disk	Flash	Disk	Flash	Disk		
██████		1,895.08	236.76	182.54	54.22	50.45us	100.99us		2.23us			1,658.32	1,623.45	34.87	317.87us	934.16us	33.00ns	1.95us

Throughput of Flash and Disk:

Top Databases by IO Throughput

- The top 10 databases by IO Throughput are displayed
- (*) indicates current database. Current database is always displayed.
- %Captured - % of Captured DB IO throughput
- Total - total IO throughput or IO requests (Flash + Disk)
- Ordered by IO Throughput desc

DB Name	DBID	IO Throughput (MB)					IO Requests			
		%Captured	Total MB	per Sec	Flash	Disk	Total Requests	per Sec	Flash	Disk
██████	3287276164	95.76	487,132.11	134.72	374,654.67	112,477.44	6,852,602	1,895.08	6,530,474	322,128

Upgrade Oracle Database 19c – Details

Note: You can download latest preupgrade script from the document:

How to Download and Run Oracle's Database Pre-Upgrade Utility (Doc ID 884522.1)

Step 1 - Preupgrade step

```
/u01/app/oracle/product/19.0.0.0/dbhome_1/jdk/bin/java -jar
/u01/app/oracle/product/19.0.0.0/dbhome_1/rdbms/admin/pre_up/preupgrade.jar TERMINAL TEXT
```

The Pre-Upgrade Tool (preupgrade.jar) creates the following files:

PREUPGRADE SUMMARY

```
/oracle/DATABASE_NAME/11204/cfgtoollogs/ DATABASE_NAME /preupgrade/preupgrade.log
```

```
/oracle/ DATABASE_NAME /11204/cfgtoollogs/ DATABASE_NAME /preupgrade/preupgrade_fixups.sql
```

```
/oracle/ DATABASE_NAME /11204/cfgtoollogs/ DATABASE_NAME /preupgrade/postupgrade_fixups.sql
```

Step 2 – Run Pre Upgrade Fixups

Log into the database and execute the preupgrade fixups.

```
sqlplus "/ as sysdba"
```

```
@/oracle/DATABASE_NAME/11204/cfgtoollogs/DATABASE_NAME/preupgrade/preupgrade_fixups.sql
```

Step 3 – Shutdown Database and Start Database in Upgrade Mode in the new Oracle Home

```
sqlplus "/ as sysdba"
```

```
startup upgrade
```

Step 4 – Execute de upgrade command

```
$ORACLE_HOME/bin/dbupgrade
```

Step 5 – Execute Post Upgrade Fixups.

Log into the database and execute the postupgrade fixups.

```
sqlplus "/ as sysdba"
```

```
/oracle/DATABASE_NAME/11204/cfgtoollogs/ DATABASE_NAME /preupgrade/postupgrade_fixups.sql
```

Step 6 – After complete, check logs on directory to check the upgrade results.

```
$ORACLE_HOME/cfgtoollogs/DATABASE_NAME/upgrade20200115130408
```

upg_summary.log

```
Oracle Database Release 19 Post-Upgrade Status Tool      01-15-2020 13:21:4
Database Name: DATABASE_NAME

Component                Current          Full            Elapsed Time
Name                      Status          Version        HH:MM:SS

Oracle Server             UPGRADED       19.5.0.0.0     00:08:39
Oracle Real Application Clusters  VALID         19.5.0.0.0     00:00:00
Oracle XML Database      VALID         19.5.0.0.0     00:00:00
Datapatch                 00:01:03
Final Actions             00:04:16
Post Upgrade              00:00:26

Total Upgrade Time: 00:16:00

Database time zone version is 23. It is older than current release time
zone version 32. Time zone upgrade is needed using the DBMS_DST package.

Grand Total Upgrade Time:      [0d:0h:18m:7s]
```

catupgrd_datapatch_upgrade.log

```
SQL Patching tool version 19.5.0.0.0 Production on Wed Jan 15 13:15:46 2020
Copyright (c) 2012, 2019, Oracle. All rights reserved.

Log file for this invocation: /oracle/DATABASE_NAME/cfgtoollogs/sqlpatch/sqlpatch_60605_2020_01_15_13_15_46/sqlpatch_invocation.log

Connecting to database...OK
Gathering database info...done
Bootstrapping registry and package to current versions...done
Determining current state...done

Current state of interim SQL patches:
  No interim patches found

Current state of release update SQL patches:
  Binary registry:
    19.5.0.0.0 Release_Update 190909180549: Installed
  SQL registry:
    No release update patches installed

Adding patches to installation queue and performing prereq checks...done
Installation queue:
  No interim patches need to be rolled back
  Patch 30125133 (Database Release Update : 19.5.0.0.191015 (30125133)):
    Apply from 19.1.0.0.0 Feature Release to 19.5.0.0.0 Release_Update 190909180549
  No interim patches need to be applied

Installing patches...
Patch installation complete. Total patches installed: 1

Validating logfiles...done
Patch 30125133 apply: SUCCESS
  logfile: /oracle/DATABASE_NAME/cfgtoollogs/sqlpatch/30125133/23151502/30125133_apply_DATABASE_NAME_2020Jan15_13_16_04.log (no errors)
SQL Patching tool complete on Wed Jan 15 13:16:48 2020
```

dba_registry details

COMP_NAME	VERSION	STATUS
Oracle Database Catalog Views	19.0.0.0.0	VALID
Oracle Database Packages and Types	19.0.0.0.0	VALID
Oracle Real Application Clusters	19.0.0.0.0	VALID
JServer JAVA Virtual Machine	19.0.0.0.0	VALID
Oracle XDK	19.0.0.0.0	VALID
Oracle Database Java Packages	19.0.0.0.0	VALID
OLAP Analytic Workspace	19.0.0.0.0	VALID
Oracle XML Database	19.0.0.0.0	VALID
Oracle Workspace Manager	19.0.0.0.0	VALID
Oracle Text	19.0.0.0.0	VALID
Oracle Multimedia	19.0.0.0.0	VALID
Spatial	19.0.0.0.0	VALID
Oracle OLAP API	19.0.0.0.0	VALID
Oracle Label Security	19.0.0.0.0	VALID
Oracle Database Vault	19.0.0.0.0	VALID

The database upgrade version 11.2.0.4 to 19C is very fast (in this case 20 minutes) and simple as we can see here. It is very important before you upgrade your database version, to check your application compatibility and always validate the database performance using Oracle tools like RAT.

Note: All the steps described to perform the database upgrade were done by command lines and scripts provide from Oracle Support.

CONNECT WITH US

Call +1.800.ORACLE1 or visit oracle.com.
Outside North America, find your local office at oracle.com/contact.

 blogs.oracle.com

 facebook.com/oracle

 twitter.com/oracle

Copyright © 2020, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

This device has not been authorized as required by the rules of the Federal Communications Commission. This device is not, and may not be, offered for sale or lease, or sold or leased, until authorization is obtained.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. 0120

White Paper Title
May, 2020
Author: [OPTIONAL]
Contributing Authors: [OPTIONAL]

