

ExaCS Database Backup and Restore with Object Storage Performance Observations

Produced by:

Oracle Maximum Availability Architecture Systems Development - Server Technologies

January 2022

Agenda

- 1 Goals
- 2 Configuration Details
- 3 Backup and Restore Throughput Summary
- 4 Reading the Result Charts
- 5 Test Results and Analysis with Full Rack ExaCS – 368 OCPUs
- 6 Test Results and Analysis with Quarter Rack ExaCS – 92 OCPUs
- 7 Test Results and Analysis with Quarter Rack ExaCS - 48 OCPUs
- 8 Test Results and Analysis with Quarter Rack ExaCS - 24 OCPUs
- 9 Setup Recommendations: Appendix A (Backup) & B (Restore)

Goals

Provide backup and restore performance results based on default settings and optional higher RMAN parallelism

Results from MAA and Cloud development

Configuration Details

RMAN/Backup/Restore

- Backup location: Object Storage Service
 - Using SWIFT-BASED object storage configuration (libopc module version 12.2.0.2)
- Level 0 backups with section size 64GB, filesperset 1¹
- Compression: Tested with OFF vs RMAN LOW¹
- Backup channels use SCAN service to spread across RAC¹
 - For these observations, backups channel allocation was based on OCPU count
 - Refer to Appendix A for recommended BACKUP channel counts and modification details.
- Restore RMAN channels use one database node¹
 - For these observations, the restore was executed outside the DBaaS tooling: A manual restore run block was configured to use SCAN instances or TNS alias across the RAC with 2 * backup channels.
 - Refer to Appendix B for manual run block RESTORE examples.

¹Current defaults for DBaaS tooling

Configuration Details

Shape 1: Environment configuration – ExaCS Full Rack (X7-2)

- Database: 8 Node RAC CDB (RDBMS 19.7)
- With 40 PDBs: Used space of ~70+TB : TDE : No RDBMS compression
- OLTP workload: Total of ~98K+ TPS against the PDBs – running from two clients.
- 2 iterations: @ 368 OCPUs

Shape 2: Environment configuration – ExaCS Quarter Rack (X7-2)

- Database: 2 Node RAC CDB (RDBMS 19.7)
- With 2 PDBs: Used space of ~7 TB : TDE : No RDBMS compression
- OLTP workload: Total of ~5K+ TPS against the PDBs – running from one client.
- 3 iterations each: @ 92, 48 & 24 OCPUs

Backup and Restore Throughput Summary

Backup Throughput Summary

Configuration	Low Impact ¹ Approximate < 5% CPU overhead	Medium Impact ¹ Approximate < 10% CPU overhead	High Impact ¹ Approximate < 25% CPU overhead
ExaCS Full Rack 368 CPUs	14 TB/hour ²	19.6 TB/hour	33.0 TB/hour
ExaCS ¼ Rack 92 CPUs	5.4 TB/hour	9.3 TB/hour	13.7 TB/hour
ExaCS ¼ Rack 48 CPUs	2.8 TB/hour	5.3 TB/hour	9.2 TB/hour
ExaCS ¼ Rack 24 CPUs	1.3 TB/hour	2.6 TB/hour ³	4.8 TB/hour ³

Restore Throughput Summary

Configuration	Low Impact ¹ Approximate < 10% CPU overhead	Medium Impact ¹ Approximate < 25% CPU overhead	High Impact ¹ Approximate < 50% CPU overhead
ExaCS Full Rack 368 CPUs	28.0 TB/hour	31.0 TB/hour	31.0 TB/hour
ExaCS ¼ Rack 92 CPUs	12.8 TB/hour	17.8 TB/hour	24.1 TB/hour ²
ExaCS ¼ Rack 48 CPUs	8.3 TB/hour	14.1 TB/hour	18.3 TB/hour
ExaCS ¼ Rack 24 CPUs	4.4 TB/hour	8.8 TB/hour	14.0 TB/hour

¹ Low, Medium & High impacts are based on channels allocation in Appendix A & B

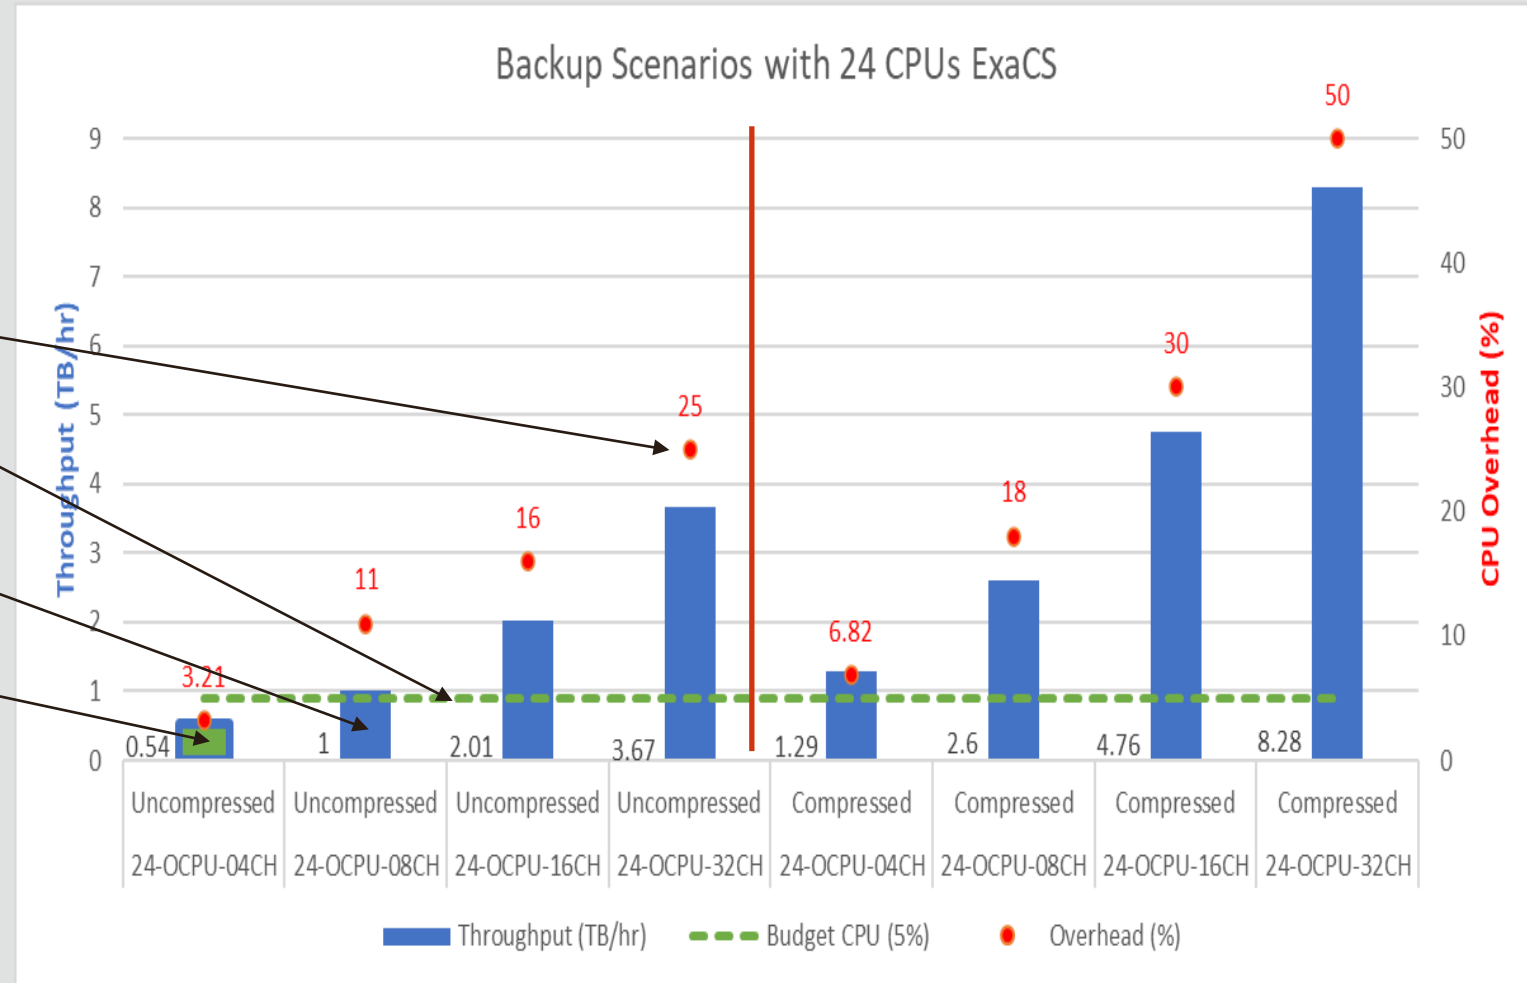
² Throughput numbers for these cells are not reflected in the result charts

³ Impact exceeds CPU Overhead, see result charts

Reading the Result Charts

Each test was performed three times to eliminate any possible outlier.

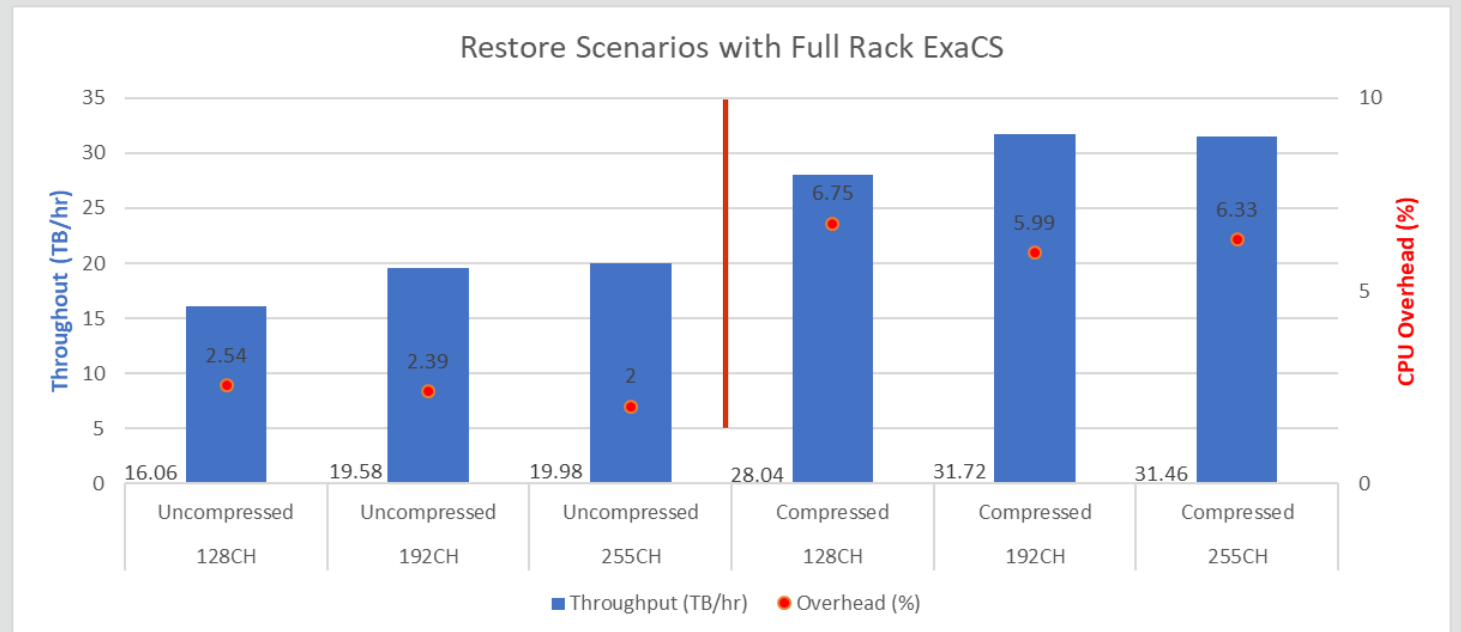
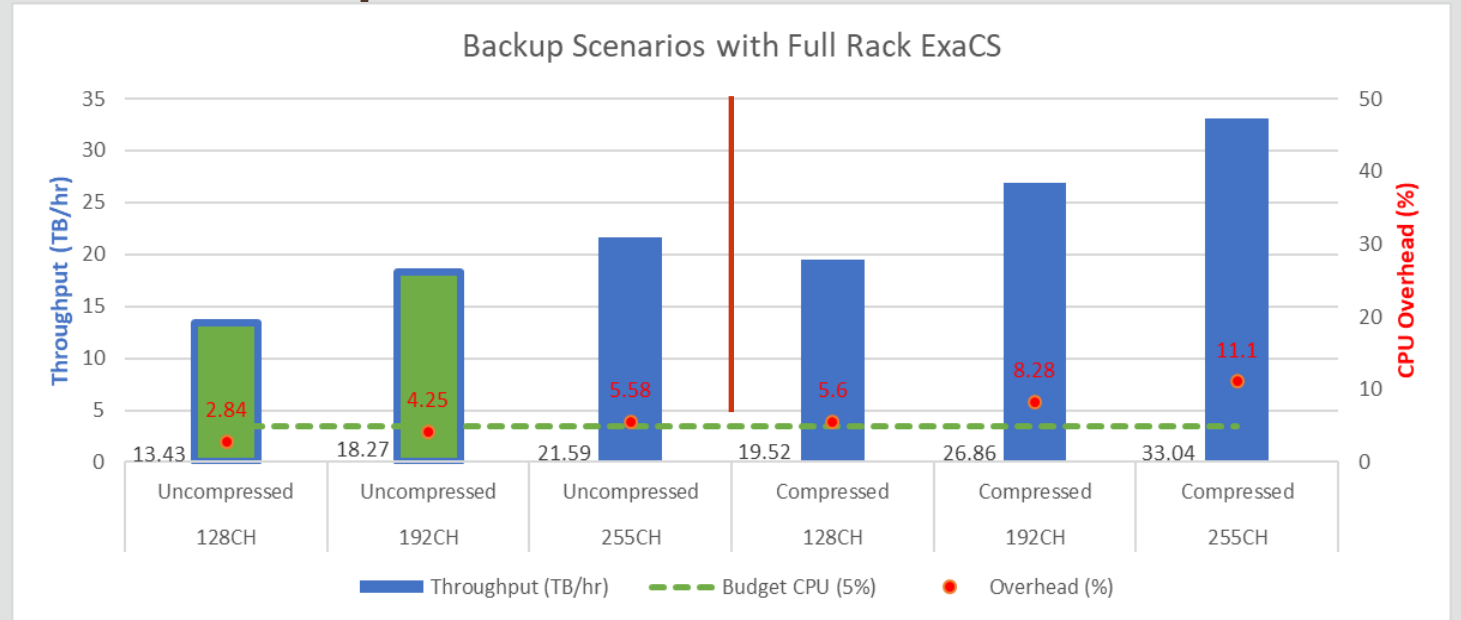
- The red dots represents the CPU overhead averaged per node across the cluster, as percentage
- The green line is the Low Impact 5% targeted CPU
- The blue bar represents the total throughput (TB/hr) across the cluster.
- A green fill within the blue bar represents an acceptable result – CPU overhead is approximately 5% or less
- RMAN compression runs shown to the right of the red bar



ExaCS full rack – 368 OCPUs per cluster

Observations

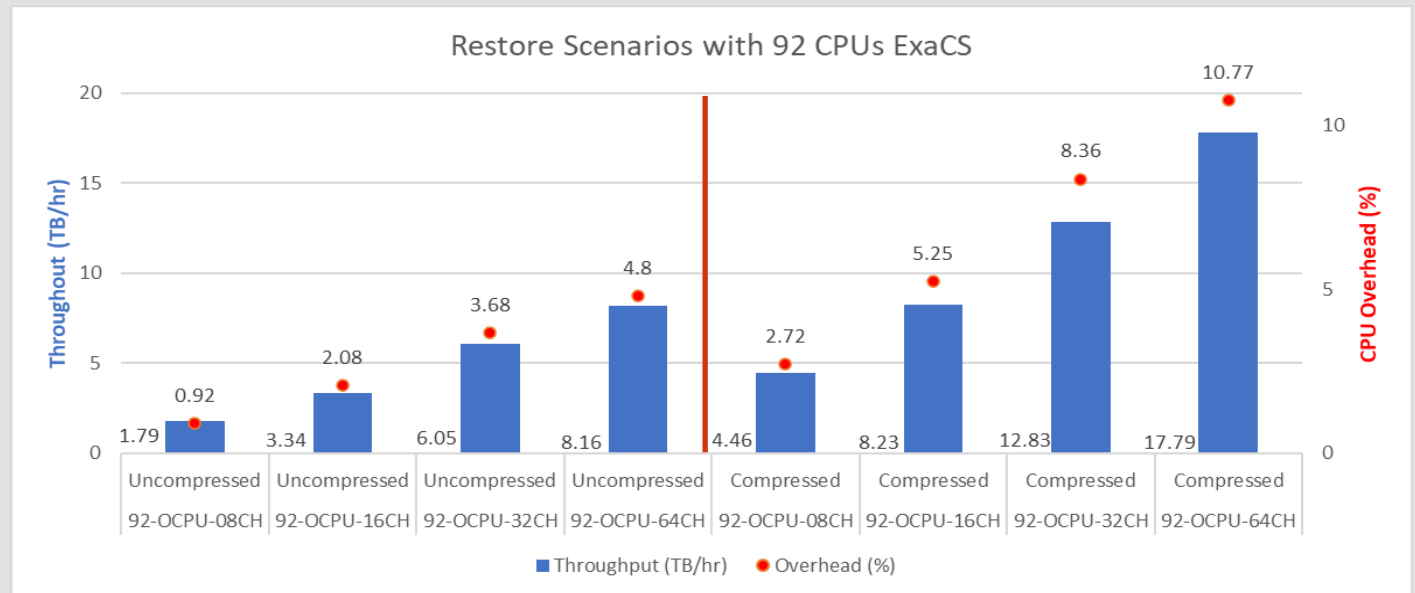
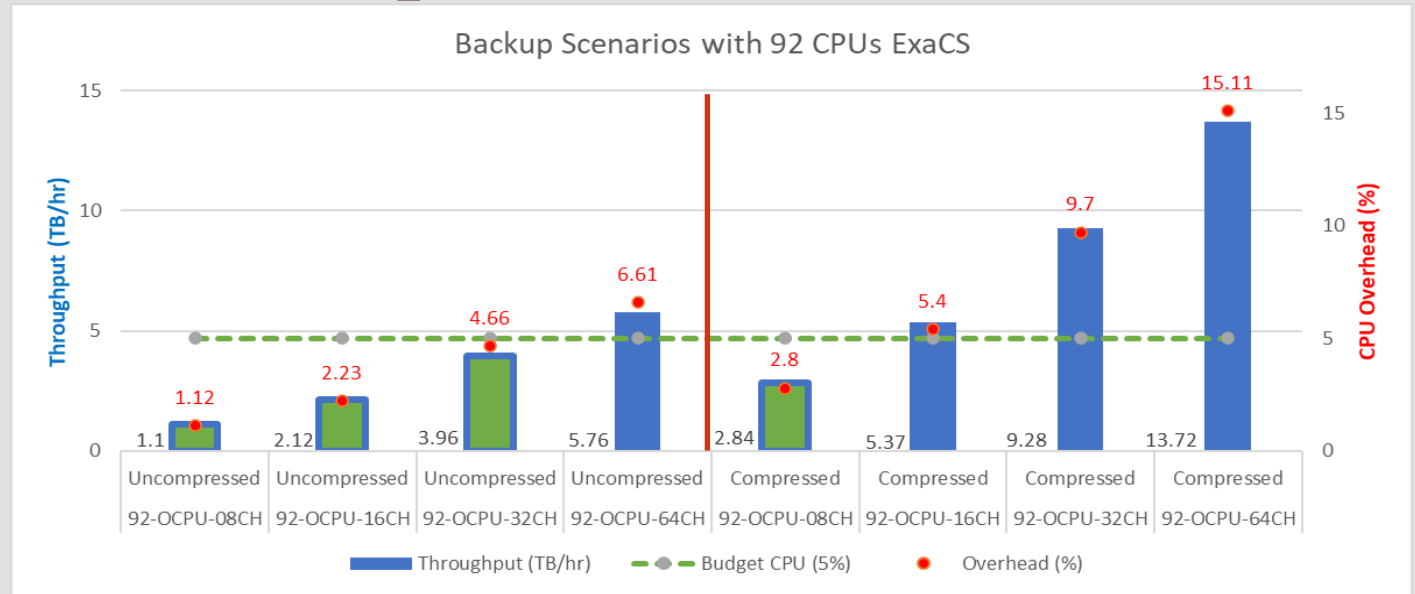
- Backup: 64 channels total + compression
 1. 14 TB/hour
 2. Less than 4% CPU overhead
- Restore: 128 channels total (compressed)
 1. 28 TB/hour
- Effective backup rates for incremental backups can be 2-10X depending on change rate
- To improve backup or restore rates, increase RMAN channel parallelism with trade off of higher CPU and IOPS utilization



ExaCS quarter rack - 92 OCPUs per cluster

Observations

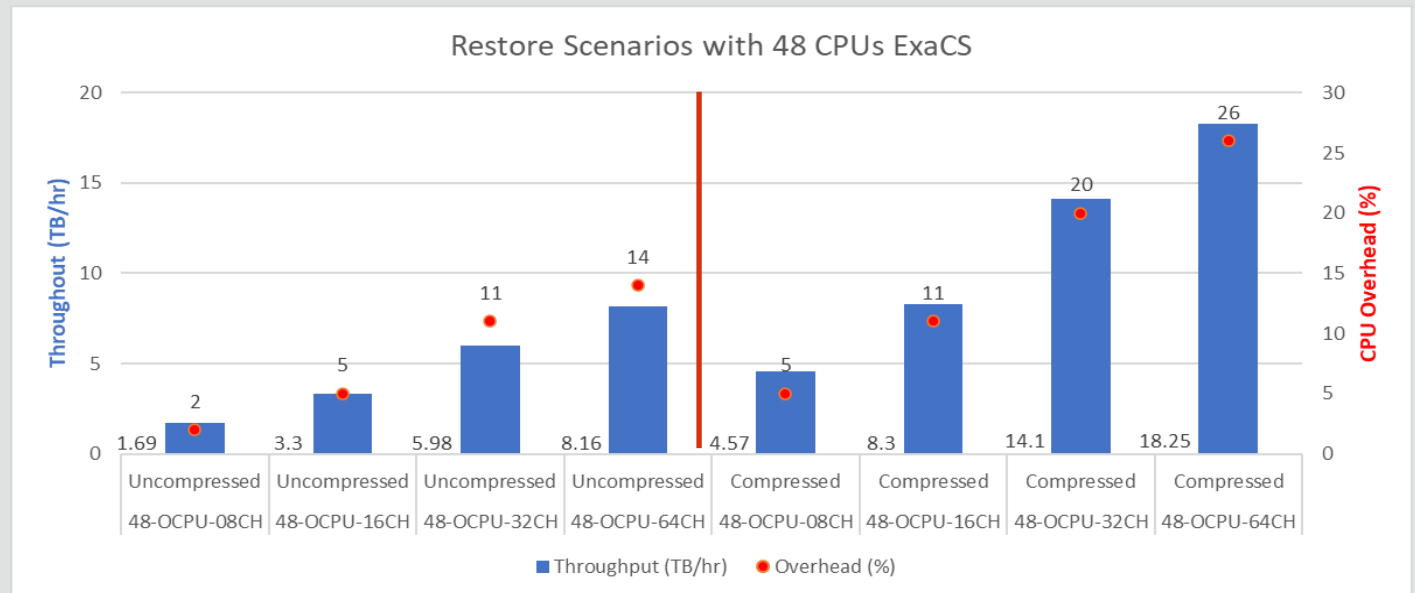
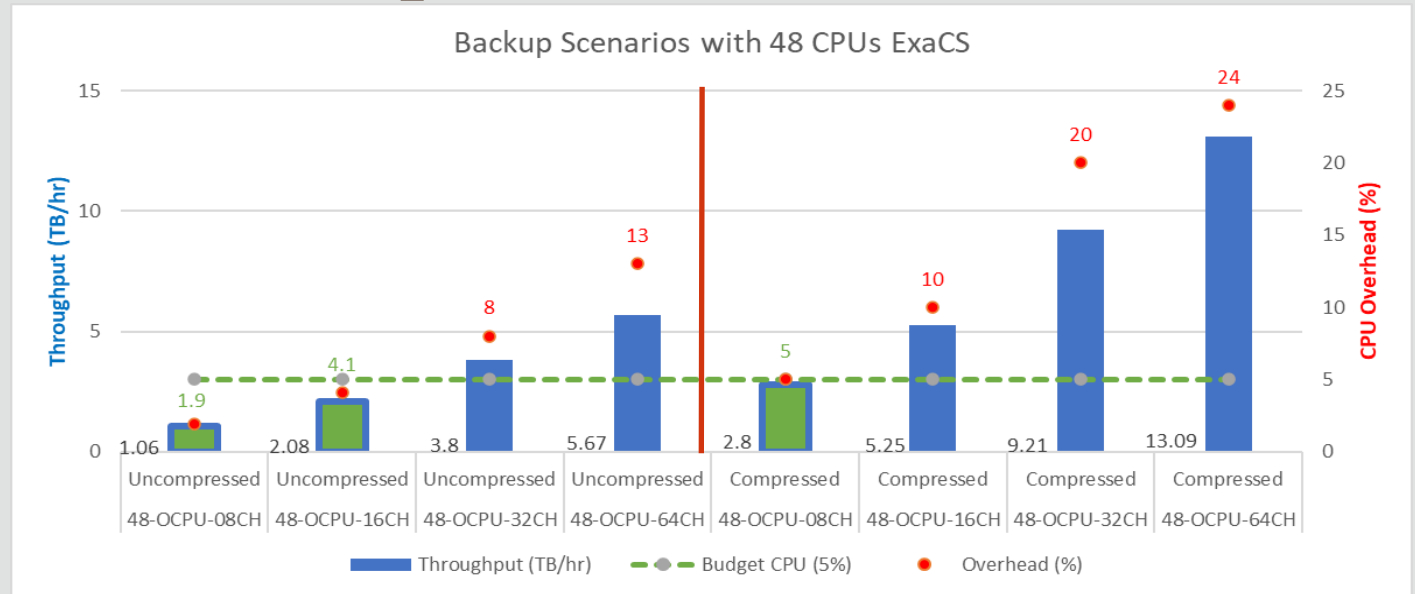
- Backup: 16 channels total + compression
 1. 5.4 TB/hour
 2. Less than 6% CPU overhead
- Restore: 32 channels total (compressed)
 1. 12.8 TB/hour
- Effective backup rates for incremental backups can be 2-10X depending on change rate
- To improve backup or restore rates, increase RMAN channel parallelism with trade off of higher CPU and IOPS utilization



ExaCS quarter rack - 48 OCPUs per cluster

Observations

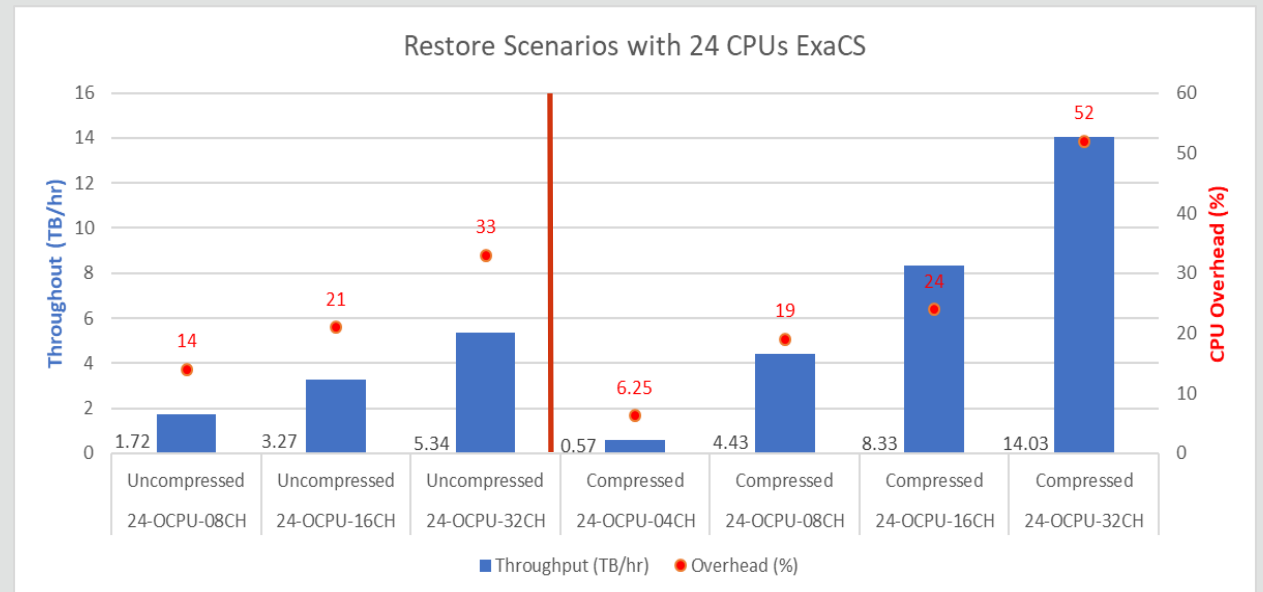
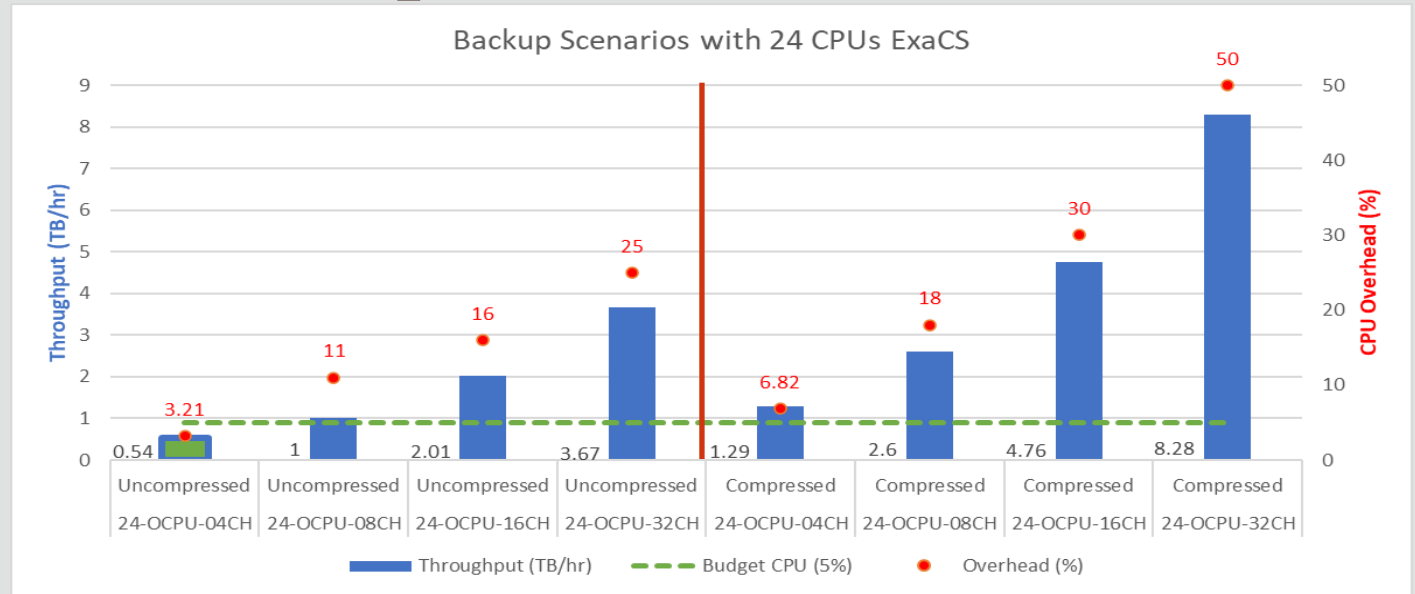
- Backup: 8 channels total + compression
 1. 2.8 TB/hour
 2. Less than 5% CPU overhead
- Restore: 16 channels total (compressed)
 1. 8.3 TB/hour
- Effective backup rates for incremental backups can be 2-10X depending on change rate
- To improve backup or restore rates, increase RMAN channel parallelism with trade off of higher CPU and IOPS utilization



ExaCS quarter rack - 24 OCPUs per cluster

Observations

- Backup: 4 channels total + compression
 1. 1.3 TB/hour
 2. Less than 7% CPU overhead
- Restore: 8 channels total (compressed)
 1. 4.4 TB/hour
- Effective backup rates for incremental backups can be 2-10X depending on change rate
- To improve backup or restore rates, increase RMAN channel parallelism with trade off of higher CPU and IOPS utilization



Appendix A

- Changing RMAN channel parallelism: backups
 - Follow [the documentation](#) to modify `bkup_channels_node`
 - Guidance for setting the value of `bkup_channels_node`

OCPU Criteria per database node	RMAN Channels Low Impact ¹	RMAN Channels Medium Impact	RMAN Channels High Impact
ocpu per node > 24	8	16	32
12 < ocpu per node <= 24	4	8	16
6 < ocpu per node <= 12	2	4	8
ocpu <= 6	1	2	3

¹-Recommended channel settings

Appendix B

- Changing RMAN channel parallelism: restore
 - Allocate channels across all nodes using CONNECT clause
 - Allocate twice the number of channels used by BACKUP.
- Example of an RMAN RESTORE VALIDATE:

```
RUN
{
ALLOCATE CHANNEL SBT_01 DEVICE TYPE SBT parms='SBT_LIBRARY=/path_to/libopc.so, ENV=(OPC_PFILE=/path_to/opcDB.ora)' CONNECT 'sys/****@scan/service as sysdba';
ALLOCATE CHANNEL SBT_02 DEVICE TYPE SBT parms='SBT_LIBRARY=/path_to/libopc.so, ENV=(OPC_PFILE=/path_to/opcDB.ora)' CONNECT 'sys/****@scan/service as sysdba';
.
.
ALLOCATE CHANNEL SBT_16 DEVICE TYPE SBT parms='SBT_LIBRARY=/path_to/libopc.so, ENV=(OPC_PFILE=/path_to/opcDB.ora)' CONNECT 'sys/****@scan/service as sysdba';
ALLOCATE CHANNEL SBT_16 DEVICE TYPE SBT parms='SBT_LIBRARY=/path_to/libopc.so, ENV=(OPC_PFILE=/path_to/opcDB.ora)' CONNECT 'sys/****@scan/service as sysdba';
RESTORE DATABASE VALIDATE;
}
```

- Example of an RMAN RESTORE / RECOVER to latest:

```
RUN
{
ALLOCATE CHANNEL SBT_01 DEVICE TYPE SBT parms='SBT_LIBRARY=/path_to/libopc.so, ENV=(OPC_PFILE=/path_to/opcDB.ora)' CONNECT 'sys/****@scan/inst1 as sysdba';
ALLOCATE CHANNEL SBT_02 DEVICE TYPE SBT parms='SBT_LIBRARY=/path_to/libopc.so, ENV=(OPC_PFILE=/path_to/opcDB.ora)' CONNECT 'sys/****@scan/inst2 as sysdba';
.
.
ALLOCATE CHANNEL SBT_16 DEVICE TYPE SBT parms='SBT_LIBRARY=/path_to/libopc.so, ENV=(OPC_PFILE=/path_to/opcDB.ora)' CONNECT 'sys/****@scan/inst1 as sysdba';
ALLOCATE CHANNEL SBT_16 DEVICE TYPE SBT parms='SBT_LIBRARY=/path_to/libopc.so, ENV=(OPC_PFILE=/path_to/opcDB.ora)' CONNECT 'sys/****@scan/inst2 as sysdba';
RESTORE DATABASE;
RECOVER DATABASE;
}
```