ORACLE

Hybrid Columnar Compression

Kevin Closson, Performance Architect
Systems Technology Group
Oracle

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.

Hybrid Columnar Compression and Performance

- Is compression a cost/space/power/cooling savings feature …
- or a performance feature?
  - Yes

Hybrid Columnar Compression and Performance

- Oracle Database doesn't decompress data for the sheer fun of it. Decompression is a part of query processing
- Query processing starts with I/O and continues on to filtration and column projection
- Compression ratio is also the physical I/O savings ratio
- Decompression is inserted between I/O and SQL processing. Strictly speaking, it is a tax
- Spending CPU cycles to uncompress data takes CPU from query processing
- Decompression needs to produce a lot of output to pay for itself

Hybrid Columnar Compression and Performance

- Compression changes I/O to CPU balance considerations

Hybrid Columnar Compression Performance Characteristics

	GB/s Physical	4GFC HBAs	EHCC Reinflated Data Flow (GB/s)	Required Cores for "Primary Processing"
Generic System *	4	10	40	6.8
Database Machine	21	N/A	210	35.7
	50	N/A	500	85

* Hypothetical as HCC is an Exadata Feature

Hybrid Columnar Compression Performance Characteristics

	GB/s Physical	4GFC HBAs	EHCC Reinflated Data Flow (GB/s)	Required Cores for "Primary Processing"
Generic System *	4	10	40	6.8
Database Machine	21	N/A	210	35.7
	50	N/A	500	85

- Division of work: Sun Oracle Database Machine has 112 processor cores for decompression and filtration (and a lot more)
- A hypothetical equivalent non-Exadata system would need on the order of 10 cluster nodes just for decompression and filtration if provisioned 21 GB/s I/O

Hybrid Columnar Compression and Performance

- Extreme features like Hybrid Columnar Compression are best handled by extreme, balanced hardware!
- Sun Oracle Database Machine is extreme, balanced hardware and HCC is an Exadata feature…an excellent fit!
- Whew, sigh of relief!


ORACLE®