

# Breaking Exadata

*Jonathan Lewis*

*jonathanlewis.wordpress.com*

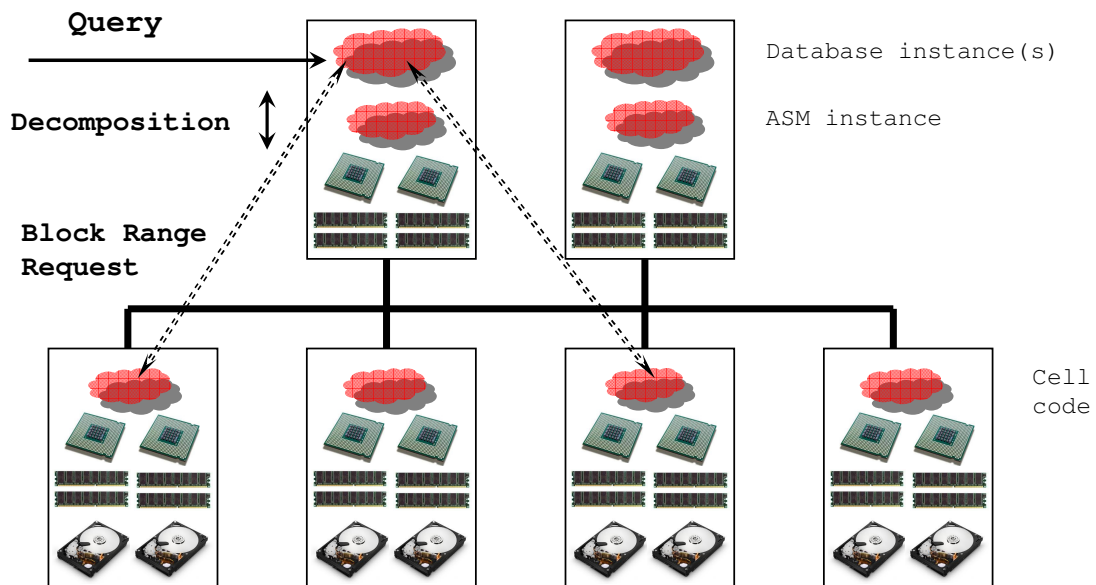
*www.jlcomp.demon.co.uk*

## USPs

---

- Hybrid Columnar Compression
- Storage Indexes
- Smart flash cache
- Smart Scans / Offload

# Implementation



# Interpretation of USPs

- Implicitly parallel scans
- More rows scanned in MB if HCC
- Don't read MB if not needed
- Cells cache their data "intelligently"
- Network traffic is minimal useful data

# ETL

---

- Often not selective on columns
- Often not selective on rows
- Compression failure
  - 1 CU generates more than 1MB => CU forwarded

# Hash Joins

---

- We dump intermediate results to disc
- We read the whole lot back
  
- It's not compressed
- There's no row/column elimination
- It's not in the Smart Flash Cache