Oracle OpenWorld 09

Svetozár Kapusta

17th of November 2009

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

•Outline

- Introduction
- Keynotes
  - *Extreme Innovation*, Scott McNealy, Sun
  - *Lowering Costs with 11G2*, Andy Mendelsohn
- Talks
  - *Bare Bones ASM, What Every DBA Needs to Know*, Jay Caviness
  - *Oracle Advanced Compression: Reduce Storage, Reduce Costs, Increase Performance*, Bill Hodak
  - *Hybrid Columnar Compression: The Next-Generation Compression Technology*, Bill Hodak, Amit Ganesh
  - *Best Practices for Zero Risk, Zero Downtime Database Maintenance*, Joseph Meeks, Michael T. Smith
  - *RMAN 11g New Features R1 and R2*, Michael R. Messina

•Oracle OpenWorld

- Massive Oracle Conference
- ~37 000 Attendees
- 1 Week in San Francisco
- 15-20 degrees Celsius

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

•*Extreme Innovation*
Scott McNealy, Sun Microsystems

- Mentioned all Sun products
- 43rd in the world for budget on R&D
- ORACLE will spend more money developing SPARC, Solaris, MySQL then Sun does now

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Real Application Clusters

Virtualize low cost servers into a shared resource

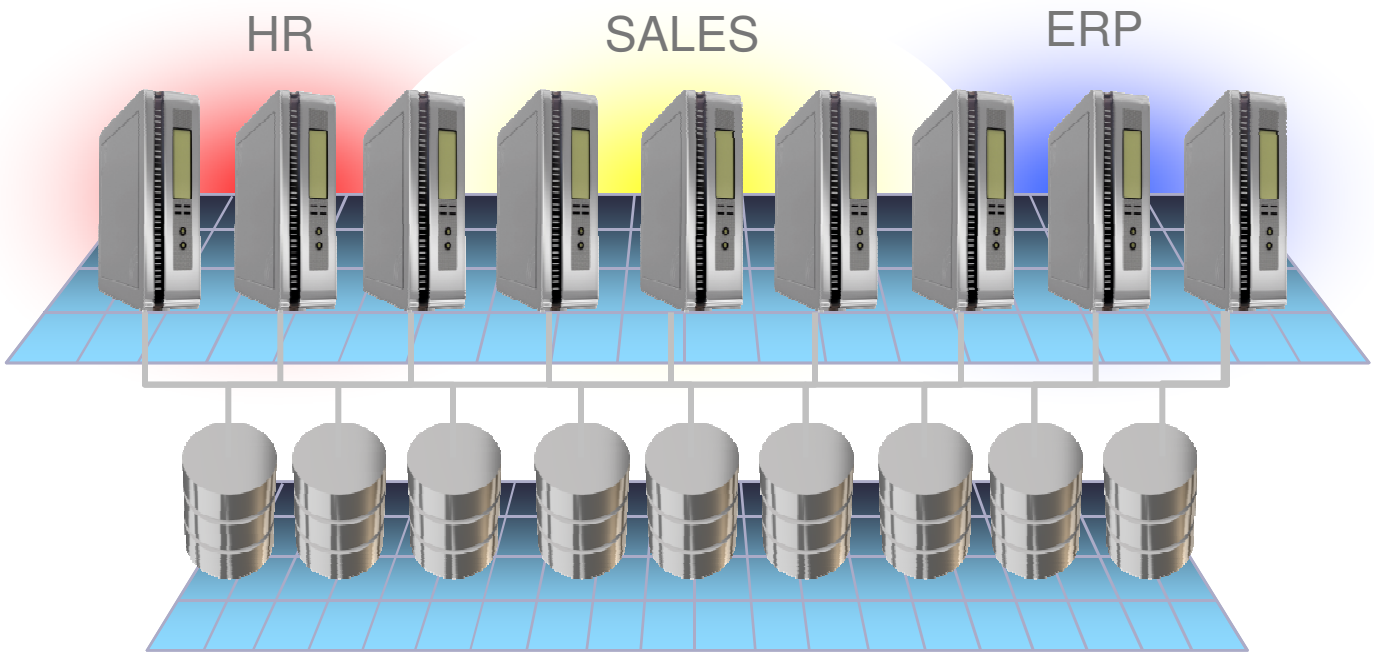

- Run all databases for all applications on shared platform
- Dynamically assigns servers to run groups of related workload
- Allocation is Policy Managed – Min/Max, Relative Importance
- Cluster reconfigures if a pool falls below its minimums

•*Lowering Costs with 11G2*, Andy Mendelsohn (2)

RAC One Node (Option)

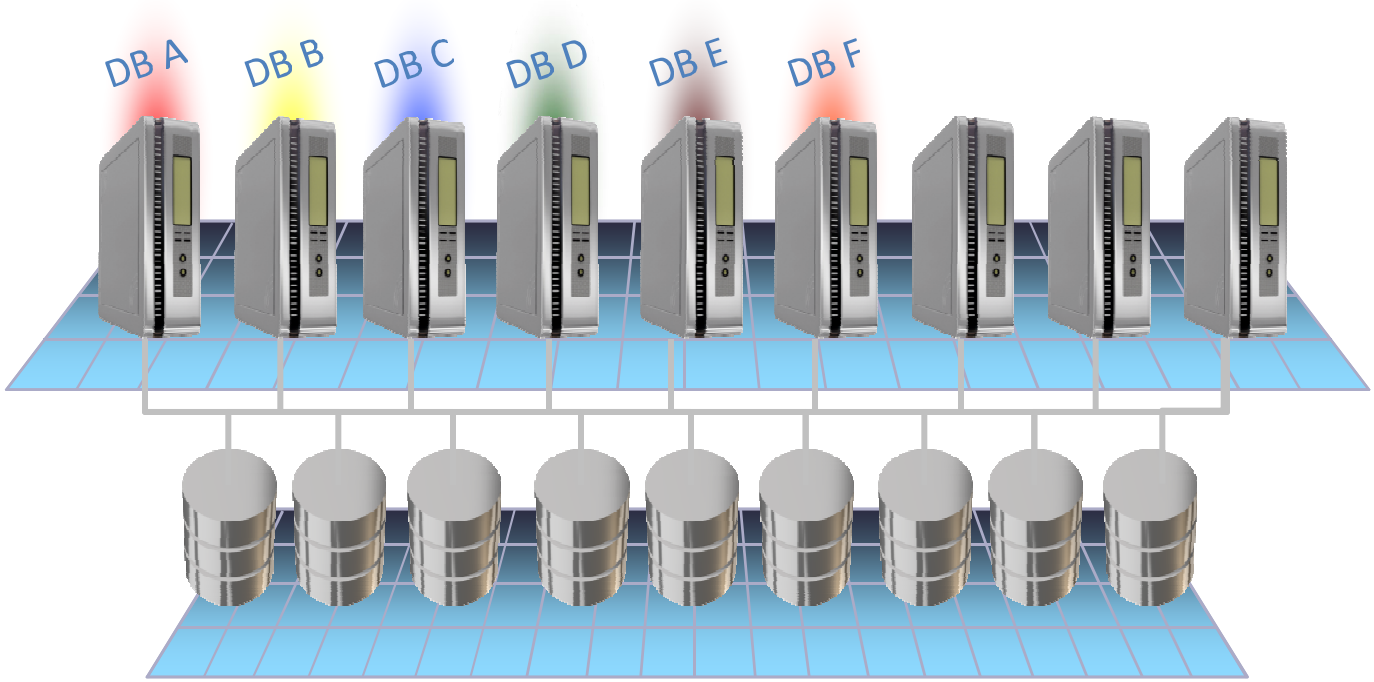

- Low entry cost to grid consolidation
- Automated failover within Grid
- Rolling patches
- Online upgrade to multi-node RAC instances

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

•*Lowering Costs with 11G2*, Andy Mendelsohn (3)

CERN IT Department

Online Application Upgrade

Edition-based redefinition in Oracle Database 11g Release 2

- Prior to 11G2 applications are either
  - unavailable for hours
  - or separate upgrade environments need to be set up and synchronized
- Online application upgrade with uninterrupted application availability
- Pre-upgrade application and post-upgrade application can be used at the same time
- End-user sessions can be rolled over

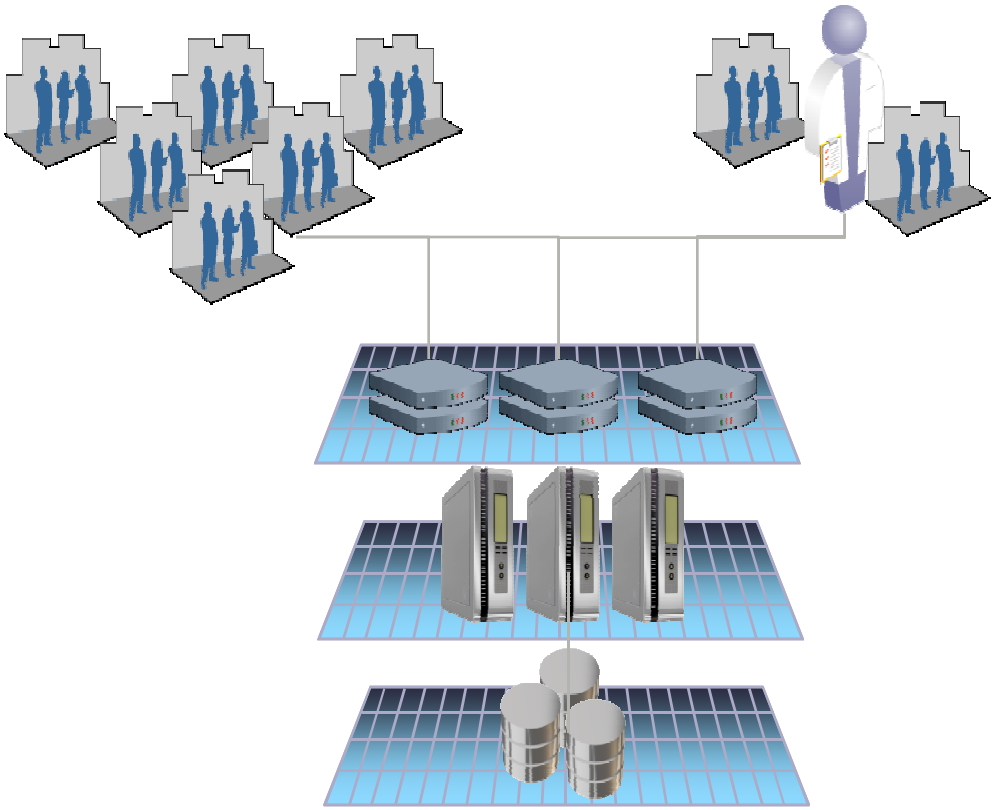

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Sun Oracle Database Machine

2x Version 1 Data Warehousing Performance

World's Fastest Machine for OLTP

Extreme Performance for Random I/O (1M IOPS)

Dramatic new Exadata Software Capabilities

Oracle Database Server Grid

- 8 Database Servers, 64 Cores, 400 GB DRAM

Exadata Storage Server Grid

- 14 Storage Servers, 5TB Smart Flash Cache, 336 TB Disk Storage

Unified Server/Storage Network

- 40 Gb/sec Infiniband Links, 880 Gb/sec Aggregate Throughput

Completely Fault Tolerant

10M$ for anyone not running at least 2X faster on Exadata2 than IBM

•*Lowering Costs with 11G2*, Andy Mendelsohn (5)

Start Small and Grow

Basic System	Quarter Rack	Half Rack	Full Rack
$110,000	$350,000	$650,000	$1.15M

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

•*Bare Bones ASM-What Every DBA Needs to Know*,Jay Caviness

- Introduction to ASM
- ASM slowly getting more popular
- Use 50GB LUNs (didn't hear about 2TB bug)

- New in 11G2
  - ASM part of the clusterware
  - ACFS
  - ASMCA, ASMCMD
  - RAC - OCR/Voting disks in ASM

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

- OLTP table compression
  - Compression factors 2X-8X
  - Also quoted our results

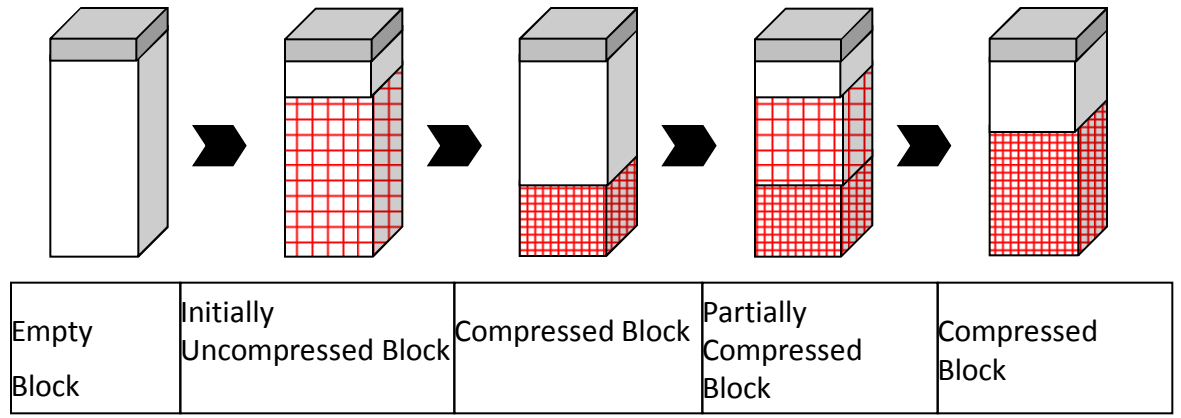

•*Oracle Advanced Compression: Reduce Storage…*, Bill Hodak (2)

- SecureFiles compression and deduplication
  - LOW / [MEDIUM] / HIGH
  - 2-3X compression for typical files
  - LOW: 80% compression of default while 1/3 CPU
- Data Pump compression
  - No need to decompress before import
  - Internal tests reduced dump file size up to 75%

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

•Oracle Advanced Compression: Reduce Storage…, Bill Hodak (3)

- RMAN fast backup compression
  - LOW / MEDIUM / HIGH / [DEFAULT]
  - 10G and 11G medium achieved 6X compression
  - 11G medium is 2.5X faster than 10G

- Data Guard Redo transport compression
  - Lower bandwidth networks (<100Mbps)
    - 15-35% less time required to transmit 1 GB of data
    - Bandwidth consumption reduced up to 35%
  - High bandwidth networks (>100 Mbps)
    - Compression will not reduce transmission time
    - But will reduce bandwidth consumption up to 35%

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

•Hybrid Columnar Compression:… Bill Hodak, Amit Ganesh (1)

- Tables are organized into Compression Units (CUs)
  - Usually 32kB
- Data is organized by column instead of by row
  - Column organization brings similar values close together, enhancing compression

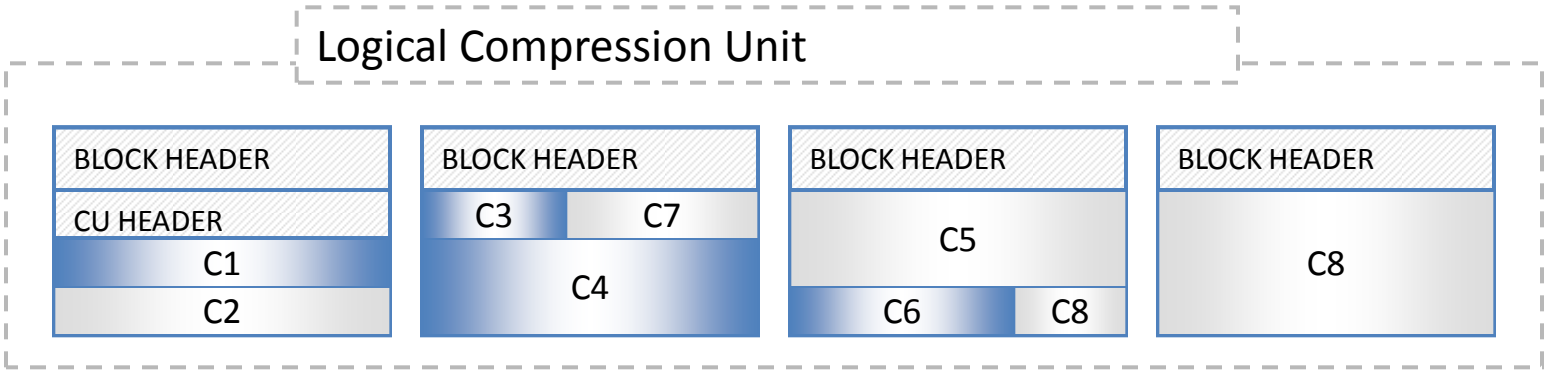

•Hybrid Columnar Compression:… Bill Hodak, Amit Ganesh (2)

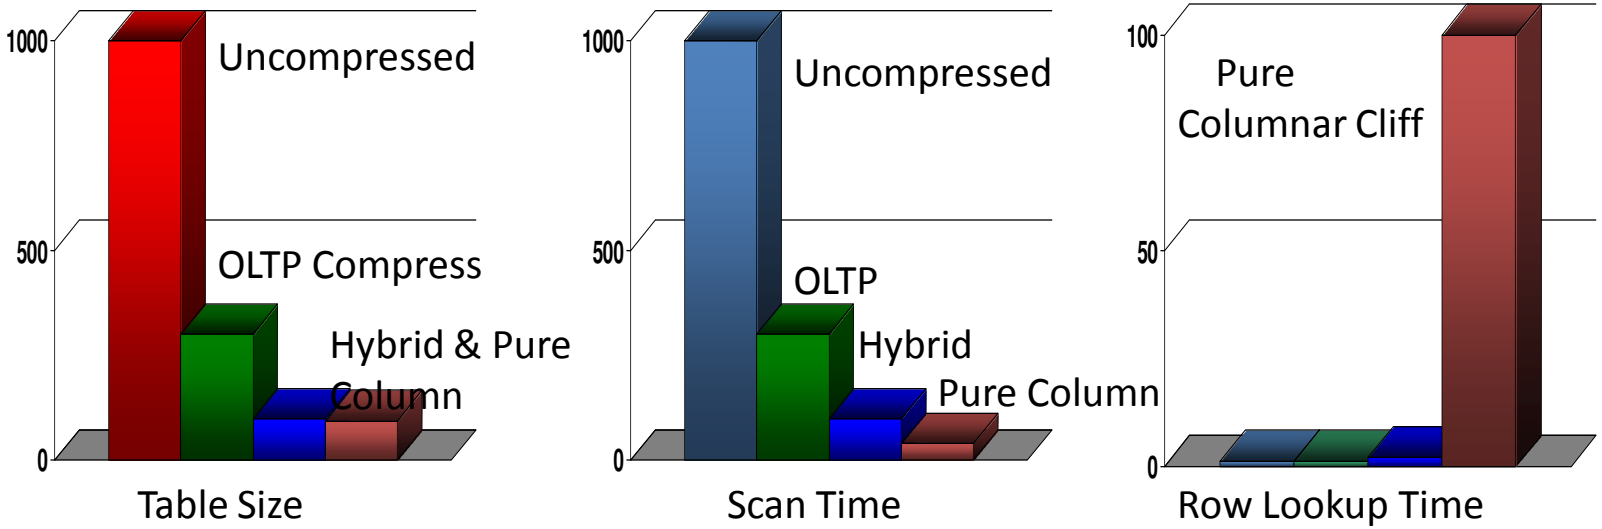

- Hybrid Columnar Compression
  - second generation columnar technology
  - combining the best of row and column formats
  - Best compression – matching full columnar
  - Excellent scan time – 93% as good as full columnar
  - Good single row lookup – no full columnar “cliff”
- Row format remains best for workloads with updates

•*Hybrid Columnar Compression:…* Bill Hodak, Amit Ganesh (3)

- Warehouse Compression
  - 10X average storage savings
  - 10X average scan improvement
- Archive Compression
  - 15X average storage savings (up to 70X)

- Can mix OLTP and hybrid columnar compression by partition for ILM
  - Heavily accessed data
    - Partitions using OLTP Table Compression
  - Cold or historical data
    - Partitions using Online Archival Compression

•*Hybrid Columnar Compression:…* Bill Hodak, Amit Ganesh (4)

- Fully supported with…
  - B-Tree, Bitmap Indexes, Text indexes
  - Materialized Views
  - Exadata Server and Cells including offload
  - Partitioning
  - Parallel Query, PDML, PDDL
  - Schema Evolution support, online, metadata-only add/drop columns
  - Data Guard Physical Standby Support
- Logical Standby and Streams supported in a future release

•Best Practices for Zero Risk, Zero Downtime DB Maintenance, (1)

- 80% of all downtime is planned downtime

- DG can significantly reduce downtime
  - Downtime only during switchover
  - UPS - Database rolling upgrade reduced downtime by 93%
  - Bielefeld University, Germany - Upgraded to Oracle Database from 9i to patched 11g with the full HA stack with downtime less than 2 minutes

•Best Practices for Zero Risk, Zero Downtime DB Maintenance, (2)

- Oracle
  - 32bit to 64bit migration
  - Migrating to ASM
  - Testing new features –e.g. flashback database
  - Migrating to Exadata storage
  - Regular bi-monthly switchovers to test standby
- Zero downtime database upgrade
  - Oracle TimesTen In-Memory (Option) database cache during switchover

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

•RMAN 11g New Features R1 and R2, Michael R. Messina (1)

- RMAN R1
  - Data Recovery Advisor
    - Simplifies diagnosis, analysis and recovery steps for a database failure that will require media recovery
    - CLI and OEM
  - Proactive Health Check
    - Proactively check database for corrupt blocks
    - Database, a tablespace or a specific datafile
  - Block Recovery Enhancement
    - Recover from flashback logs
    - Recover from the physical standby
  - Archived Log Deletion Policy Enhancements
    - APPLIED ON [ALL] STANDBY | BACKED UP integer TIMES TO DEVICE TYPE, …

•*RMAN 11g New Features R1 and R2*, Michael R. Messina (2)

- New Compression Type (ZLIB)
  - Faster then original BZIP2
  - Requires Advanced Compression License
  - RMAN> configure compression algorithm 'ZLIB' ;
- Parallel Backup of Same Datafile
  - Allows large datafile to be broken into “sections”
  - Improves speed in which large datafiles are backed up
  - RMAN> backup section size 50m datafile 4;
- Virtual Private Catalog
  - Enables multiple “virtual” catalogs within the RMAN catalog
  - Improved Security - Catalog owner grants access
  - Separation of databases or groups of databases within the RMAN catalog

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

•*RMAN 11g New Features R1 and R2*, Michael R. Messina (3)

- Better Recovery Catalog Management
  - Move a Catalog to another catalog/database via import catalog
- Active Database Duplication
  - duplicate a database without using or having an existing RMAN backup
- Fast Incremental Backup on Physical Standby
  - Block change tracking on physical standby
- Read-Only Transported Tablespaces Backup
  - No need to set them read-write

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

•*RMAN 11g New Features R1 and R2*, Michael R. Messina (4)

- RMAN R2
  - Automatic Block Repair
    - Corrupt blocks on primary automatically repaired
  - RMAN Web-Services Backup
    - Offers backup to web-based storage services (Amazon)
    - Reduces cost and time to manage in-house backup infrastructure
  - Tablespace Point in Time Recovery Enhancements
    - RMAN automatically determines if the tablespaces in the recovery set are self-contained
  - Duplicate Database attempts to continue where it left off

•Personal Conversations

- Bill Hodak (Compression Product Manager)
- Kevin Jernigan (Compression Principal Product Manager)
- Joseph Meeks (MAA Director, Product Manager)
- Paul Parsons (CTO and Founder of the SERVER LABS)
- Arup Nanda
- Amit Ganesh (Senior Director Data and Systems Technology)

- Many others…

•Conclusions

- Great conference, business oriented
- Unique opportunity to meet and discuss with project managers and other DBAs
- CERN has and continues to have an excellent reputation in using Oracle technology
  - Arup Nanda mentioning Luca's work
  - Joe Meeks mentioning CERN DG 32->64bit migration and Jacek's and Dawid's visit

Thank you for your attention

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it