Exadata High Volume Testing

- Findings from a Customer POC

Martin Roselieb / Gerald Schreyer
Oracle Solution Center for SAP Competency

Customer POC: High Volume Testing on Exadata

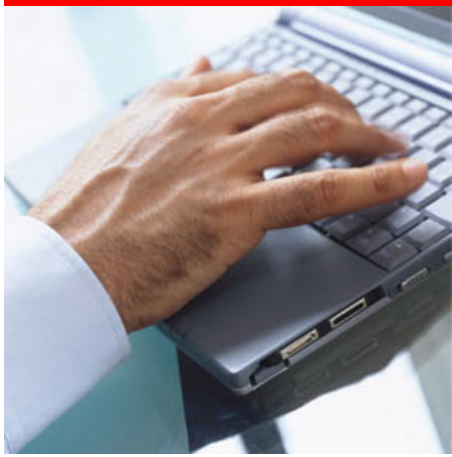

- Customer Case
- Decision For POC
- POC Preparation & Execution
- Results & Achievements
  - Volume tests
  - Demo on High-Availability Features
  - Lessons Learned / Conclusion
- Q&A

Customer Case

Business

- B-2-B retailer for stationeries and office supplies
- operates in 30+ countries
- 30% growth expected within next 4 years

Challenges

- long batch run times for invoicing: no daily processing was possible
- cost and effort to administer large number of SAP databases
- new B-2-B process could not be supported by current infrastructure

Renewal of entire IT infrastructure was considered

- Exadata for high performance and consolidation
- Application servers on x86 running Solaris OS

Decision For POC

Goals

- Expose customer's business-critical „real“ SAP workload to an Exadata environment
- Measure performance increase (throughput & response times)
- Show how Oracle technology can help to minimize the impact of hard-/software failures for the business

Rules/Conditions

- No „lab“ tuning, no DB „tweaking“
- Use standard Exadata database setup without optimization
- → all improvements should be as well achievable by customer's staff for their productive systems

Conducting the POC

Preparation (3 weeks):

- Decision taking on volume test cases:
  - Simulation of online users in a web shop
  - Batch job net for invoicing
- Installation of the entire IT infrastructure for SAP
- Loading production database from tape
- 6.5 TB DB migration from Oracle 10.2 (SPARC) via Oracle 11.2 (x86/Linux) into an Exadata ASM/RAC environment

POC Execution (1 week):

- on-site at OSC Benchmark Center in Linlithgow, U.K.
- load injection (Loadrunner, Control-M) from customer's site

Hardware* Used for Testing

*) without servers for load generation and the application tier for JAVA web shop

Application Servers
5 x Sun Blade 6270 M2
(each 2 x Intel Xeon x5680, 3.3 Ghz)
72 GB RAM
Solaris 10 x86

Exadata Database Machine
Half Rack X2-2
4 database nodes
7 storage cell nodes
High Performance Disks

Database Server
(for DB migration)
1 x Sun Fire x4800
(8 x Intel Xeon x7560, 2.27 Ghz)
128 GB RAM
Oracle Linux

HDS 9990V
(DB for migration)
connected via four 4Gps HBAs
16 x RAID5 LUNs with 8 FCAL disks each

Results Volume Testing - Batch

Invoicing of sales subsidiaries:

- job net starts 900+ batch jobs
- varying amount of jobs running in parallel (1 to 130)

Runtimes:

- 11h 30min on customer's current environment
- 3h 45min on Exadata.
  Two active nodes of a Half-Rack Exadata had only been utilized to 25%.

New environment reduced runtime to one third with a lot of spare capacity

CPU Utilization

Mixed Workload: single & parallel runs; many SAP programs

SAP App

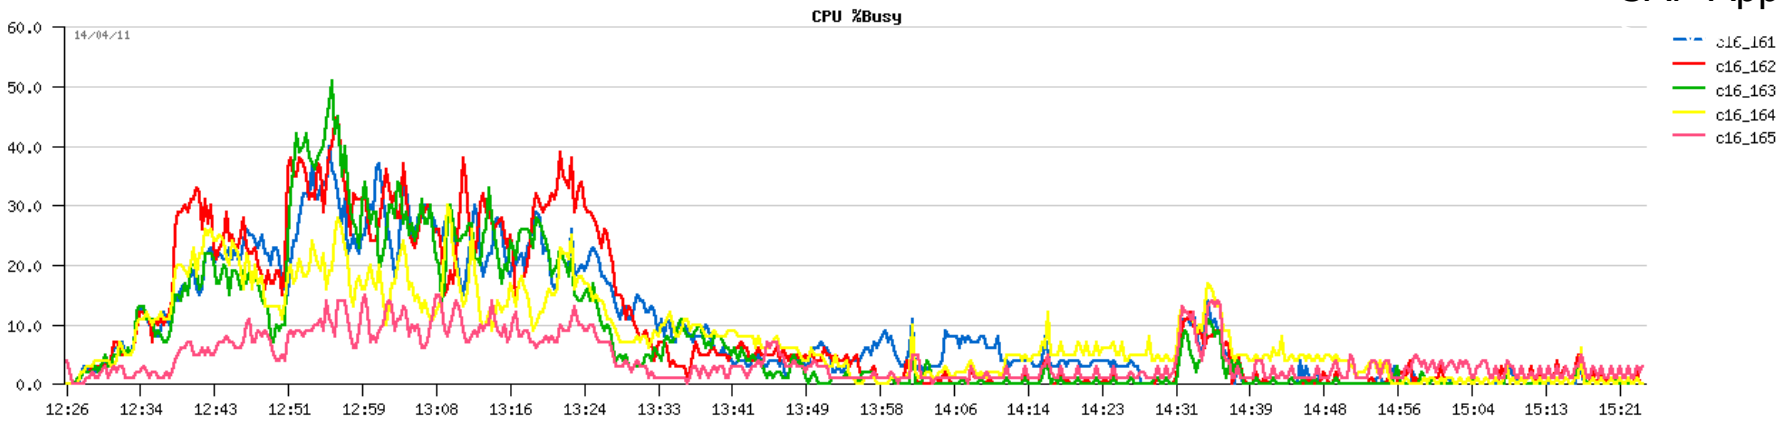

RAC Nodes

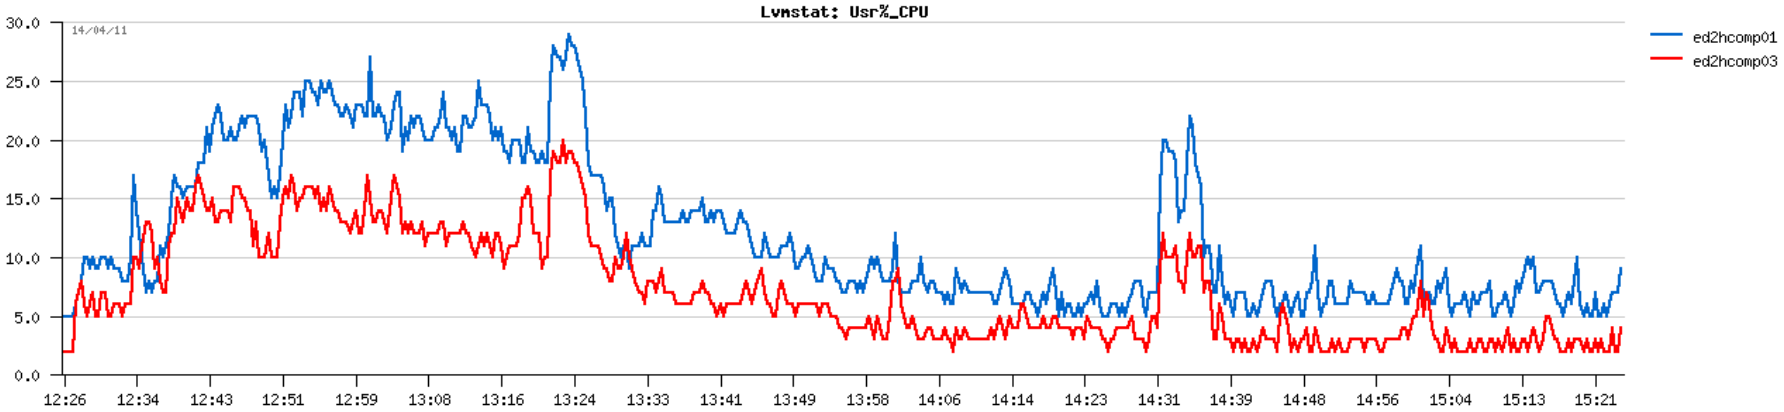

What's Exadata's Impact ?

Response time decrease to 3:45h is a significant improvement, but:

- faster Intel CPU technology had been used compared to customer's current environment
- → Rerun tests using DB staging environment (no Exadata usage)

Runtime by using single-instance DB server was 7:20h

Interpretation:

- 11:30h to 7:20h improvement is due to faster CPUs on app and DB tier
- 7:20h to 3:45h improvement is due to Exadata technology (Flash, Infiniband) and Exadata software (Smart Flash Cache..)

Results Volume Testing - Number of Online Users

Simulation of online orders entered in a web shop:

- Runs with various parameters (#users, thinktime, #RAC nodes)
- After an initial throughput was achieved, scalability bottlenecks within the SAP application had been identified and eliminated

Throughput:

- one node RAC: 7,500 users with 4s thinktime:
  => ~1s response time and 64% CPU utilization (one node)
- four node RAC: 6,000 users with 4s thinktime:
  => ~0.9s response time and 17% CPU utilization (all nodes)

We over-achieved the goal of 2,500 users still having a lot of spare capacity in Exadata

CPU Utilization

SAP App

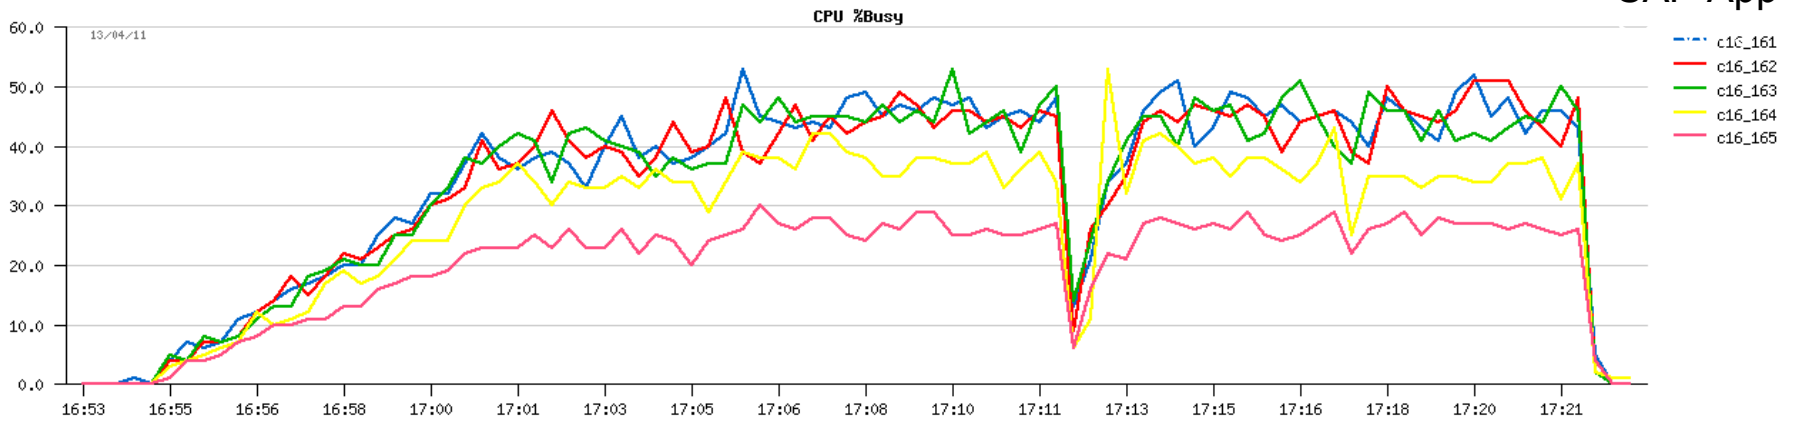

RAC Nodes

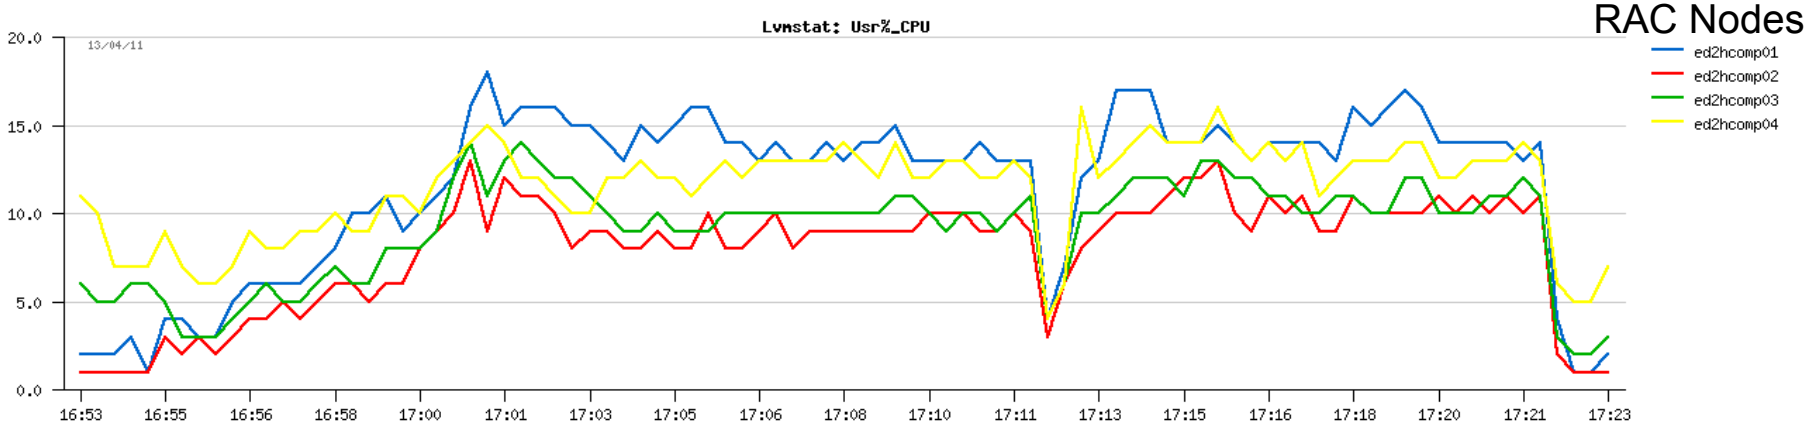

AWR Statistics from Highload Phase

from ST03N, DBACockpit, AWR

Keyfigure / per second	Batch [13% util. of POC hardware]	Online [17% util. of POC hardware]
User Calls	39.383	69.285
Physical Writes	6.538	2.952
Physical Reads	14.727	4.404
Logical Reads	348.300	389.741

Demo of High-Availability Features

Shutting down RAC nodes during SAP operations

- Showcasing the automatic reconnect from SAP application servers to RAC database nodes (TAF)
- Significant reduction of unplanned downtimes compared to a traditional server cluster with the need for a subsequent instance recovery

SAP Replicated Enqueue with Sun Cluster

- Demonstrating a solution which ensures continuous operation even in the case when losing the central SAP lock table held in main memory

SAP Instance Reconnect

Transparent Application Fail-Over (TAF)

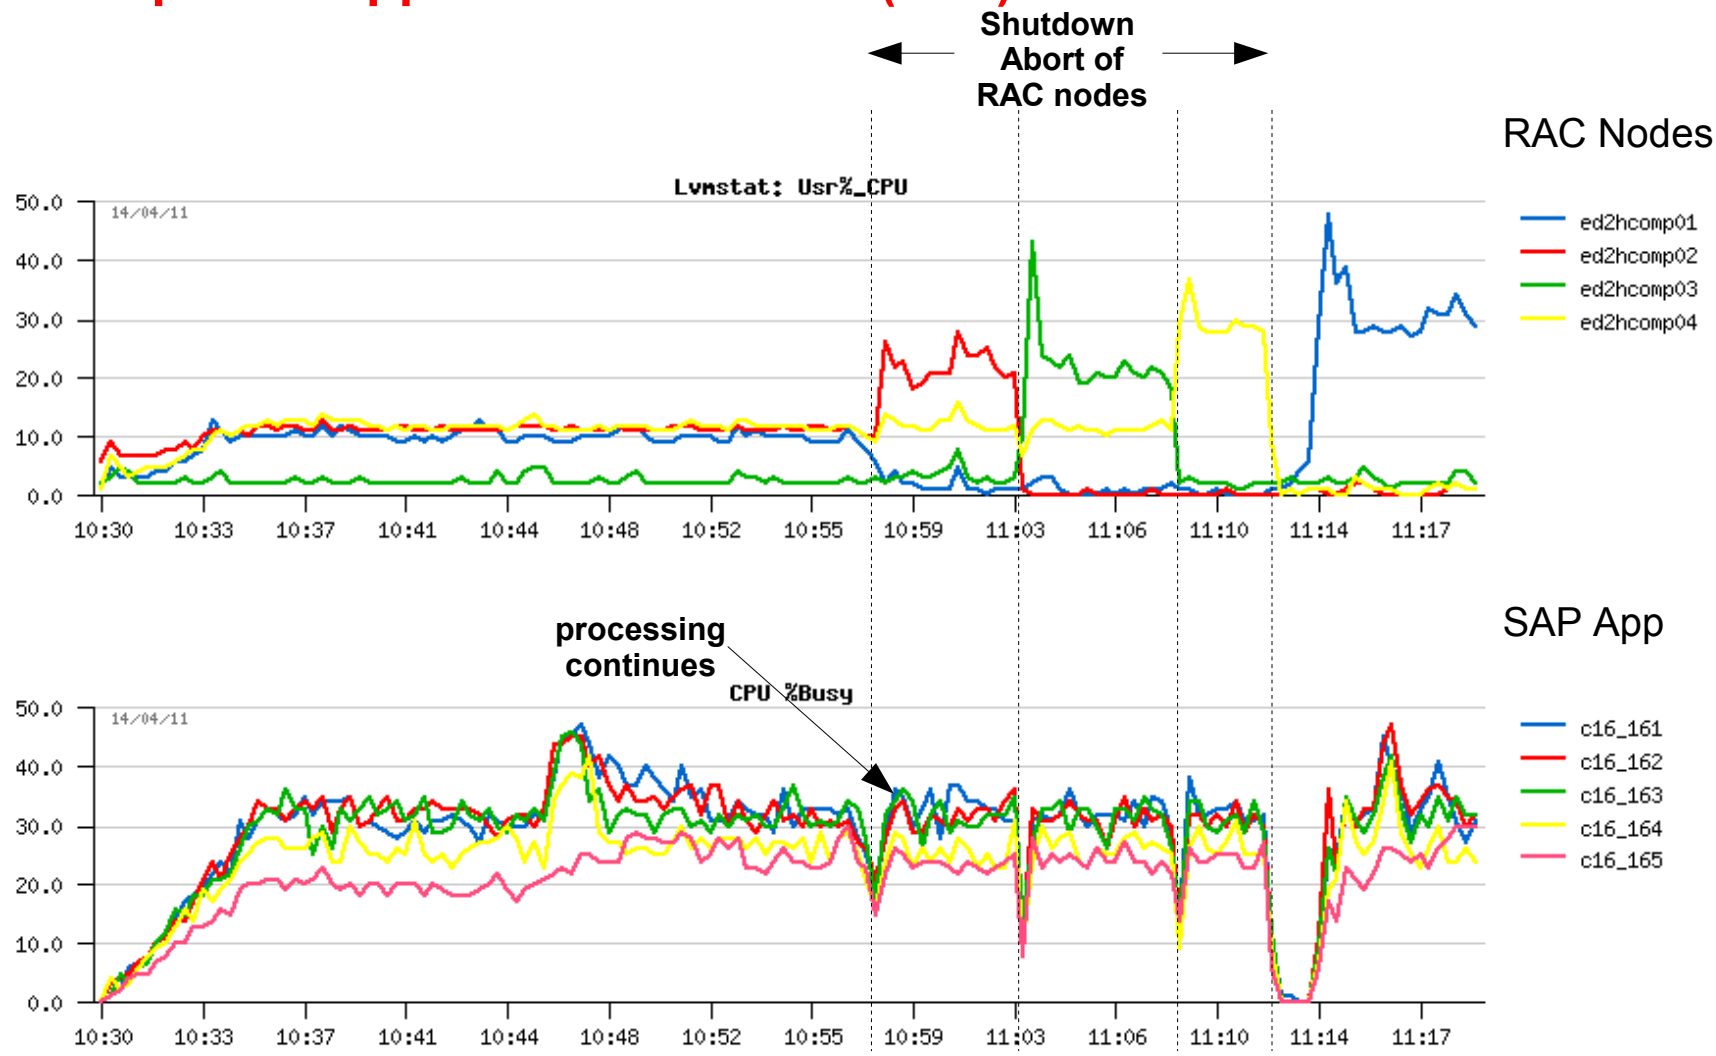

Lessons Learned

Exadata provides better performance „out-of-the-box“

- you see instant performance benefits
- no Exadata- or DB specific tweaking/optimization was needed
- however, effect of new technology should be considered as well when making comparisons (e.g. faster CPUs on app server side)

We were not able to saturate Exadata

- no bottlenecks in the hardware/infrastructure were seen
- throughput was limited by contention within the application:
  - assignment of unique document numbers for new postings (SAP)
  - „hot spot“ inserts at the end of the tables (ascending primary key)

  → **in order to achieve even higher throughput numbers and to take full advantage of Exadata you need to resolve these bottlenecks**

Lessons Learned - cont.

User I/O is no more the major wait event

- „cell single block physical read“ (corresponds to „db file sequential read“ in non-Exadata environments) is very fast (< 1msec)

RAC scales

- almost same throughput numbers in 1-2-4-node configurations
- slightly more CPU consumption in n-node RAC configs

Collecting DB workload & utilization data for sizing

- During test runs AWR reports have been captured along with Exadata utilization data
- Load profile can now be compared with DB workload from other customers for sizing & migration considerations

Lessons Learned - cont.

No Smart Scans in these business-critical workloads

- Well tuned application setup and DB-index design:
  - unique index access
  - selective secondary indices
  - no tablespace / index scans

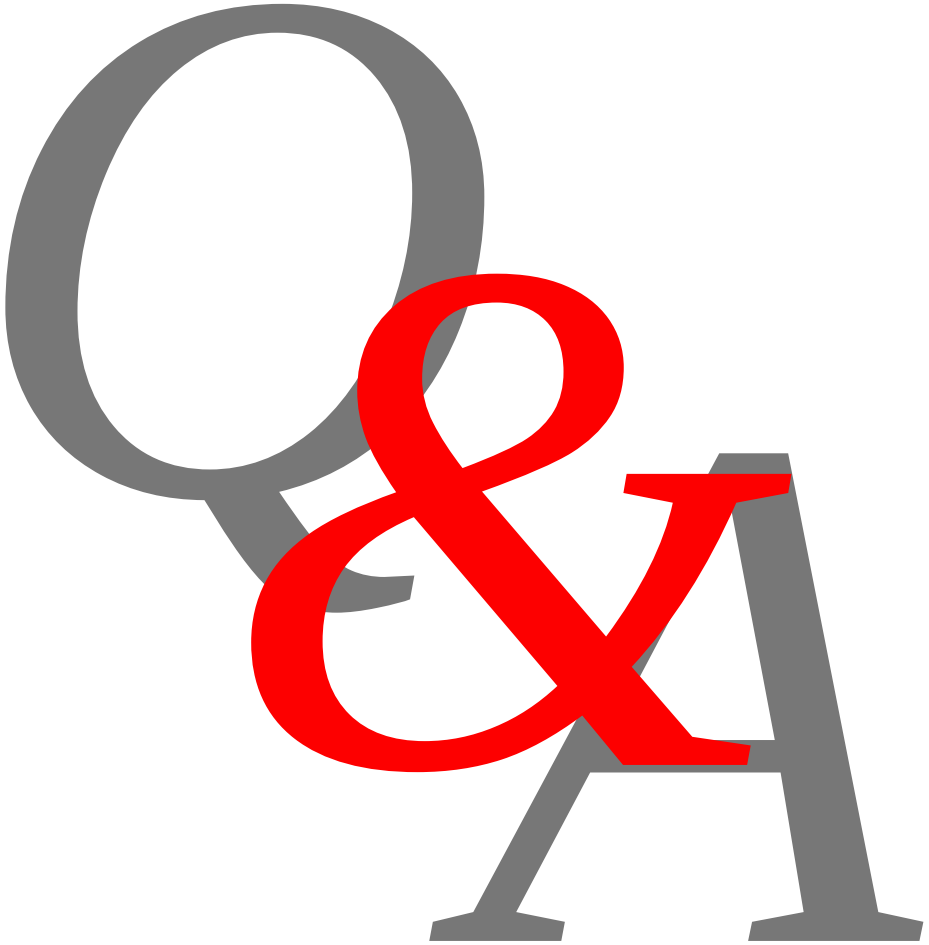
Q&A