



January 16, 2006

**Experiences of
LGT Financial Services AG
with
z/VM 5.2, Linux for zSeries and Oracle in production**

**Renate Gut and Rolf Künzler (LGT)
Dr. Manfred Gnirss and Peter Mühlemann (IBM)**

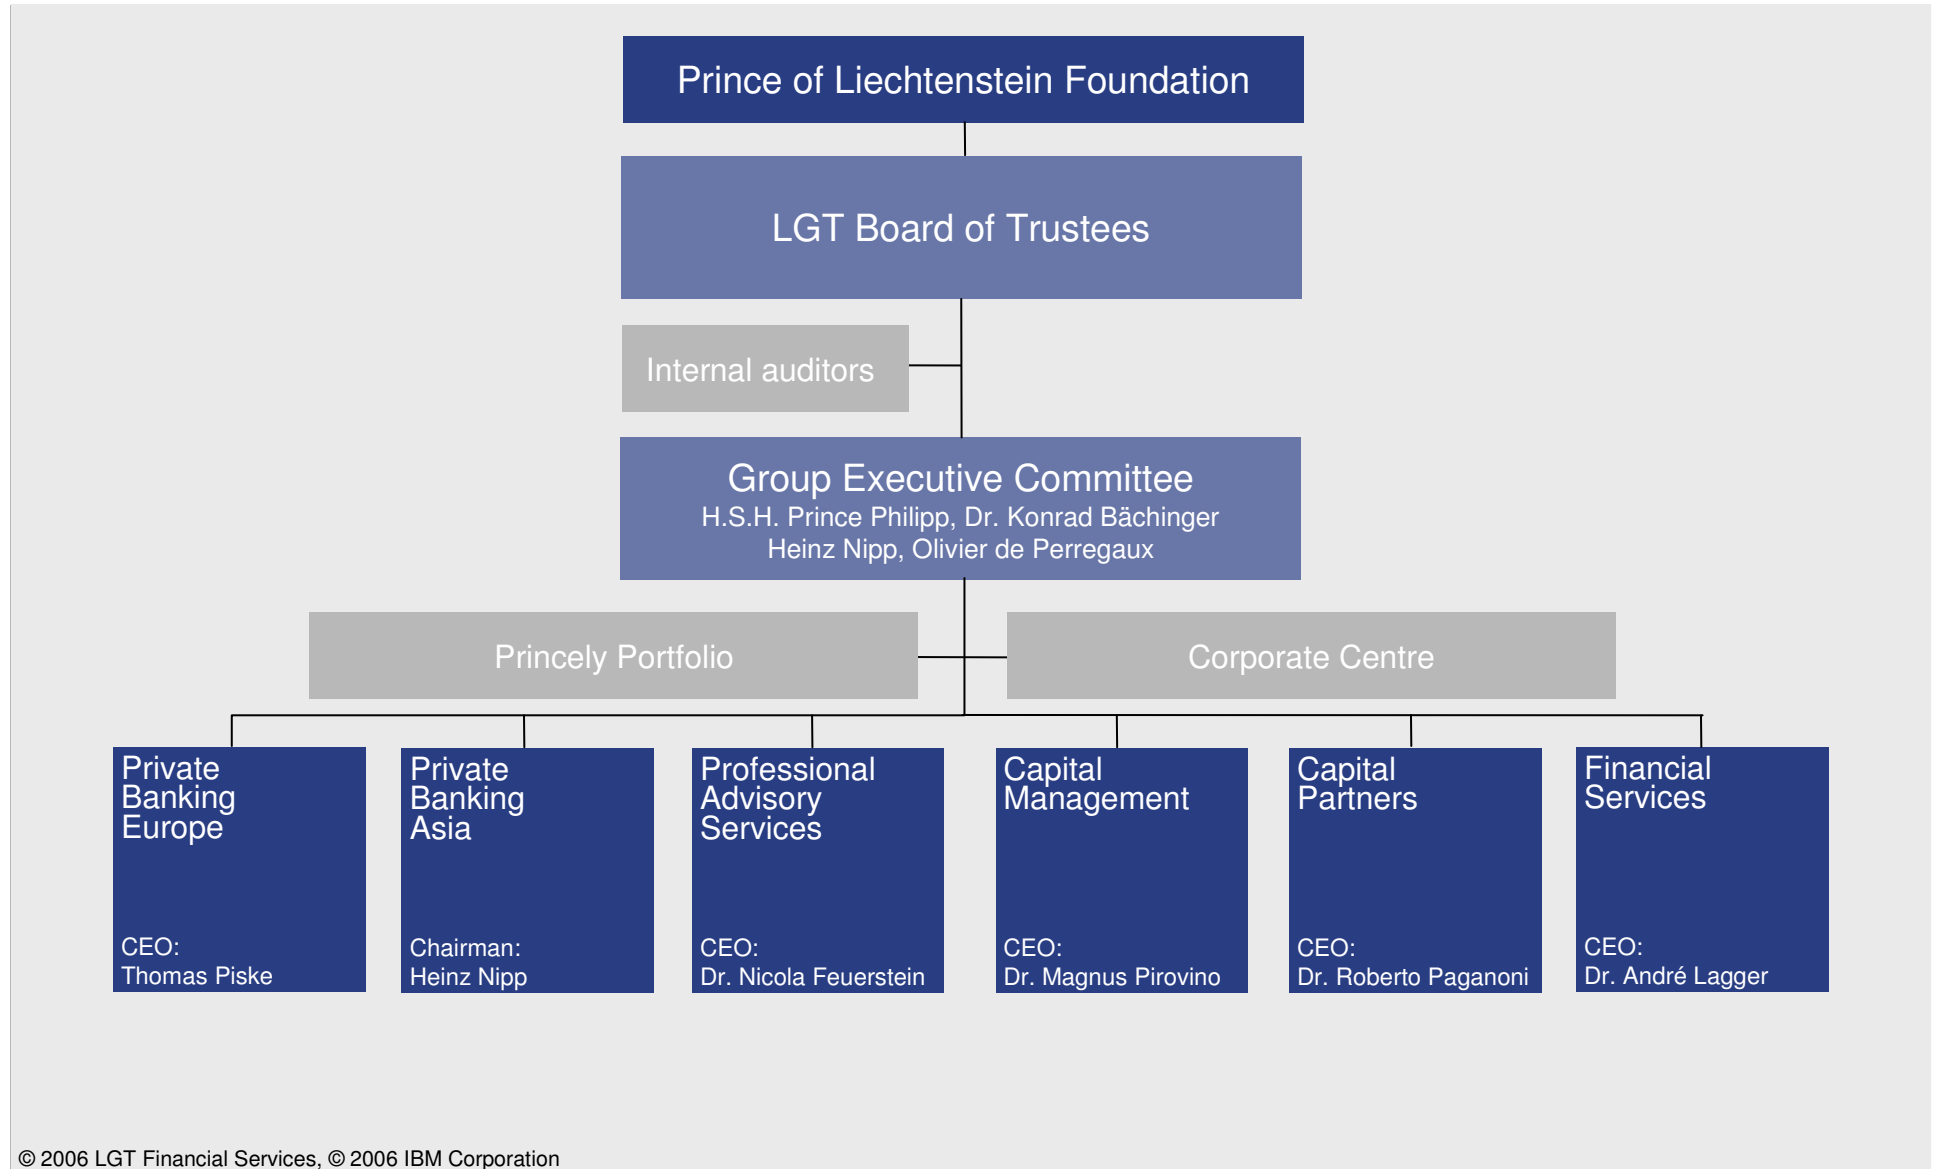
**LGT -
The wealth management experts of the Princely House of Liechtenstein.**





LGT Group

The LGT Organisation





LGT Group International Presence





LGT Group in Figures

Key figures

	30.06.2005 (1.1.05 – 30.6.05)	31.12.2004 (1.1.04 – 31.12.04)
Group profit (in CHF m)	68	110.3
Balance sheet key figures (in CHF bn)		
Balance sheet total	14.9	13.0
Equity capital (prior to appropriation of CHF 100 m to the Prince of Liechtenstein Foundation)	2.2	2.1
Assets under administration (in CHF bn)	70.1	63.2
Client assets under Management	23.5	20.3
Other client assets under administration	44.5	41.1
Princely Portfolio	2.1	1.8
number of staff	1'311	1'278



LGT Financial Services Information Technology

Five Divisions

- Customer Advisory Services
- Architecture and Strategy Development
- IT Support
- IT Application Development
- IT Engineering

Key Information

- Headoffice in Vaduz, decentralized IT support at Banking locations
- 107 Employees and 8 apprentices
- Approximately 1300 Clients worldwide
- Backoffice Application on two zSeries mainframes in different locations
- Frontend Applications (Customer Information System, Output Management System, Financial Reporting, Electronic Archive) decentralized



zSeries and Oracle environment 1/2005

- Two sites, with z890-160 and z890-170 zSeries
- DASD (HDS 7700E)
- z/OS 1.4

- Server Systems with about 92 Oracle 9i instances

The following servers were selected for the PoC

Production Environment:

- MS Server Prod11 with 2 Oracle instances (Lens)
- Unix Server Prod21 with 1 Oracle instance (Valor)
- MS Server Prod31 with 2 Oracle instances (Elas)
- Unix Server Prod41 with 1 Oracle instance (RMAN)

Test Environment:

- MS Server Test13 with 3 Oracle instances (Lens)
- Unix Server Test23 with 1 Oracle instance (Valor)
- MS Server Test33 with 1 Oracle instances (Elas)
- Unix Server Test43 with 1 Oracle instance (Restore-Pool)
- Unix Server Test53 with 1 Oracle instance (RMAN)



PoC for Oracle Consolidation on zSeries

- **Project content:**
Consolidate 6 production Oracle databases and additional test databases on zSeries with z/VM and Linux using IFLs

- **Objective:**
Show z890 with z/VM and Linux is a suitable platform to host Oracle databases for LGT Financial Services AG and therefore an alternative to existing Unix installation by providing
 - Cost effectiveness
 - Better stability
 - Improved availability
 - Good performance
 - Secure and manageable platform
 - Scalability



PoC for Oracle Consolidation on zSeries

- **Roadmap / Idea for PoC:**
 - Use external know-how for z/VM, Linux for zSeries and Oracle on zSeries during PoC – cooperate with Business Partners
 - 1:1 Migration (no upgrade from Oracle SE to Oracle EE during PoC, no consolidation of multiple instances) to keep risk low
 - BP creates documentation about installation and setup of new environment
 - Necessary education of LGT Financial Services AG employee after PoC

- **PoC environment:**
 - 1 IFL on each z890 with z/VM and Linux on try-and-buy basis
 - z/VM 4.4
 - SuSE Linux SLES 8 (as of Oracle prereq) in 31 bit mode
 - Oracle 9i (31 bit)
 - Use RMAN for Backup/Restore of Oracle
 - TSM and TDPO for Linux for zSeries

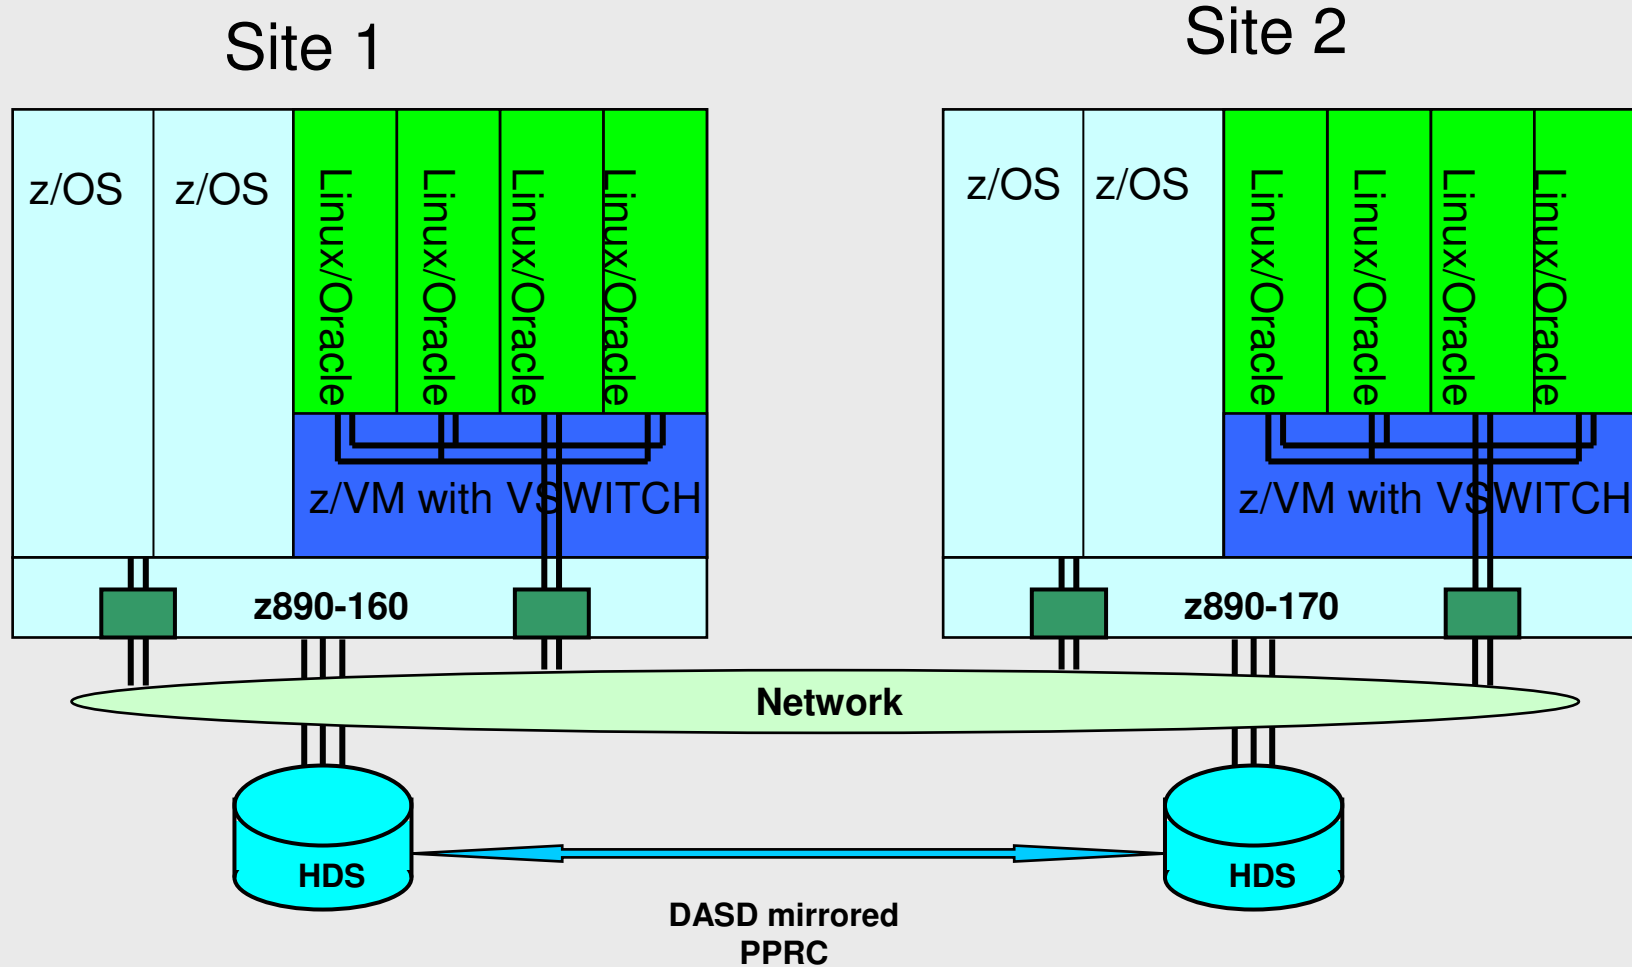


zSeries and Oracle environment 11/2005

- Two sites, with z890-160 and a z890-170 zSeries with 1 IFL in each site
- Per z890 machine:
 - 3 FICON chpids (shared) for DASD attachment and 2 OSA adapters
 - 2 LPAR with z/OS 1.4
 - LPAR with z/VM 5.2 (9 GB, 8 GB Main + 1 GB XSTOR)
 - 4 Linux Guests (each with 2 GB virtual memory) with SuSE Linux SLES 8 SP4 64-bit, with 4 Oracle 9i (31bit) instances, database size is 250 GB
 - Linux guests connected to network via two VSWITCHes (failover)
- DASD: Two HDS TagmaStore USP 100 with 32-38 3390-3 used per Linux for zSeries guest
- 13 Oracle 9i instances on Linux for zSeries
- Unix/MS Servers with about 79 Oracle 9i instances



zSeries environment 11/2005 with LGT





Project history: overview

- 11/18/2004 Workshop with IBM and Business Partner:
Understanding the potential of Linux on the mainframe
- 12/2004 Purchased two z890 with 1 IFL each on try-and-buy basis
- 2/2005 Start of PoC: Install z/VM 4.4, SuSE Linux SLES 8 and Oracle 9i together
with an experienced partner knowing zLinux on zSeries
- 3/2005 PoC in critical phase
Ask for IBM involvement for PoC (IBM Switzerland)
Oracle configuration support by IBM Montpellier/France
- 4/2005 Start involvement of the Linux for zSeries Performance team
Böblingen/Germany for analysis and tuning support
- 6/2005 Tuning of environment improved performance to satisfying level
Decision to participate in z/VM V5.2 ESP (final solution for performance)
- 7/2005 Bought IFLs and memory
- 8/2005-11/2005 Participation in ESP for z/VM 5.2 (remove perf. problems based on 2 GB
limitation)
- 9/2005-11/2005 Optimization and systems management concepts (cloning, etc.) by ITS (CH)
- 10/2005 Start production for Oracle Test and Development environment
- 12/2005 Start production for Oracle Production environment
- 12/2005 Sign for reference
- **Future/Outlook** Additional Oracle database to be consolidated on zSeries – zSeries growth



Analysis of performance situation after start of PoC

Unstable and insufficient performance after initial installation of z/VM 4.4 with Linux and Oracle.

Upgrade to z/VM 5.1 did not solve performance problems

Therefore contacted IBM for help.

Analysis showed that insufficient performance behavior is based on high I/O traffic and I/O limitations due to 2 GB memory constraint for I/O in z/VM 5.1: Too much activity and load based on system behavior when moving pages below 2 GB.



Memory usage in system after IPL

```
FCX113      CPU 2086  SER 39D7E  Interval 10:49:14 - 10:50:14  Perf. Monitor
-----
>
Data > <----- Number of Pages ----->
Spaces >
          <-Resident->
Userid   Owned >  WSS Lockd Resrvd  <2GB  >2GB  XSTOR  DASD  Stor  Nr of
>System< .0 > 89493  48    0   35271  54329  0      0    406M  10
BEV30T10 0  261876  13    0   68565 193327  0      1   1024M
BEV30T20 0  522132  247   0  253236 269340  0      1   2048M
BEV30T30 0  100386  13    0   28872  71530  0      1    700M
DISKACNT 0   1245   0     0    16    1229  0      0    32M
EREP      0   1229   0     0    21    1208  0      0    32M
OPERATOR  0     1    0     0     1     0    0      0    32M
OPERSYMP  0   1259   0     0    11    1265  0      0    32M
PERFSVM   0   1740   0     0     28   1712  0      0    64M
TCP/IP    0   2848  185   0   1092   1941  0      0    64M
VSWCTRL1  0   2215   26    0    870   1736  0      0    32M

Select a user for user details
Command ==> _____
F1=Help  F4=Top  F5=Bot  F7=Bkwd  F8=Fwd  F10=Left  F11=Right  F12=Return
```

After IPL: „Some“ pages used below 2 GB line



Memory usage when running Oracle with RMAN

FCX115 CPU 2086 SER 39D7E Interval 13:38:56 - 13:38:56 Perf. Monitor

Detailed data for user BEV30T20

Total CPU	: 32.6%	Storage def.	: 2048MB	Page fault rate:	: .0/s
Superv. CPU	: 1.0%	Resident <2GB:	86138	Page read rate:	: .0/s
Emulat. CPU	: 31.6%	Resident >2GB:	403004	Page write rate:	: .0/s
VF total	:%	Proj. WSET	: 488273	Pgs moved >2GB>:	1545/s
VF overhead	:%	Reserved pgs	: 0	Main > XSTORE	: .0/s
VF emulation	:%	Locked pages	: 75	XSTORE > main	: .0/s
VF load rate	:/s	XSTORE dedic.	: 0MB	XSTORE > DASD	: .0/s
I/O rate	: 138/s	XSTORE pages	: 0	SPOOL pg reads	: .0/s
DASD IO rate	: 138/s	DASD slots	: 1	SPOOL pg writes	: .0/s
UR I/O rate	: .0/s	IUCV X-fer/s	: .0/s	MDC insert rate	: 5.0/s
Diag. X'98'	: .0/s	Share	: 2000	MDC I/O avoided	: .0/s
*BLOCKIO	: .0/s	Max. share	: ...		

#I/O active	: 4	Active	: 100%	PSW wait	: 0%	I/O act.	: 83%
Stacked blk	: ..	Page wait	: 0%	CF wait	: 0%	Eligible	: 0%
Stat.:	EME, QDS, RNBL	I/O wait	: 0%	Sim. wait	: 0%	Runnable	: 100%

Data Space Name Size Mode PgRd/s PgWr/s XRd/s XWr/s Migr/s Steal/s

Enter 'STOrage Display' for storage details

Command ==> _____

F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F12=Return

High activity to move pages below 2 GB line during run of Oracle with RMAN



Memory usage of system while/after Oracle with RMAN

```
FCX113      CPU 2086  SER 39D7E  Interval 14:16:14 - 14:17:14  Perf. Monitor
-----
Data > <----- Number of Pages ----->
Spaces > <----- Resident----->
Userid  Owned >  WSS Lockd Resrvd  <2GB  >2GB  XSTOR  DASD  Stor  Nr of
>System<  .0 > 80932  20    0  45408  35243  2922  0  381M  11
BEV30T10  0  250819  2    0  83874  186947  11142  0  1024M
BEV30T20  0  516840  2    0  400467  116375  5734  0  2048M
BEV30T30  0  95468  2    0  28580  66890  4933  0  700M
BEV3ST10  0  23183  2    0  6317  14814  2115  1  128M
DISKACNT  0    65  0    0    0    50  1180  0  32M
EREP      0    66  0    0    6    60  1163  0  32M
OPERATOR  0    1  0    0    0    0  1  0  32M
OPERSYMP  0  1259  0    0    0    0  1259  0  32M
PERFSVM   0    618  0    0    13   605  1119  0  64M
TCPIP     0    86  185  0    192   79  2762  0  64M
VSWCTRL1  0  1847  26   0    31  1842  733  0  32M

Select a user for user details
Command ==> _____
F1=Help  F4=Top  F5=Bot  F7=Bkwd  F8=Fwd  F10=Left  F11=Right  F12=Return
```

While/after using Oracle with RMAN: High activity in XTOR and majority of pages are moved below the 2 GB line



First Tuning activities

- Major tuning activities: Try to reduce free memory for Linux to avoid high load based on 2 GB I/O constraint of z/VM 5.1:
 - Adjust memory size of Linux guests
 - Adjust size of XSTOR
 - Oracle tuning: use fixed buffer pool in Linux storage: loc_sga
 - Oracle tuning: Allocate memory during IPL: pre_page_sga
 - CMM (Collaborative Memory Management): After shutdown of Oracle-DB, free memory not given back to Linux memory management.

- Note: Above Oracle tuning methods are now documented. See also:
 - <http://www.vm.ibm.com/perf/tips/2gstorag.html>
 - Problem record 11717,031,724

- Other improvements:
 - Spool, Swap-space
 - Timeout values
 - Minidisk caches
 - Network re-architected
 - Upgrade z/VM 5.1 to latest level

- **Result:**
 - **Sufficient performance now for usual work**
 - Slow DASD performance due to using old DASD system without distribution to multiple ranks and no striping due to organizational reasons.
 - **Still unacceptable performance issues caused by usage of RMAN** (Oracle Recovery manager), as RMAN uses direct I/O which „destroys“ all performed tuning activities.
Note: Replacement of RMAN with another backup software is not applicable, as RMAN highly embedded in application environment



Resolving the remaining performance issues

- All conventional tuning activities could not resolve the performance issues for all situations (2 GB I/O constraint when using RMAN).
- There are only two potential solutions for customer's environment:
 - Linux fixed buffer
 - z/VM 5.2
- Status:
 - Linux fixed buffer was not available in summer 2005 for SuSE SLES 8 and it was an open question, whether it will be retrofitted to SLES 8.
 - Using SLES 9 is not an option for LGT due to prerequisites of Oracle 9i
 - z/VM 5.2 not yet available in summer 2005 (planned GA in 12/2005)
- Solution:
 - Participate in ESP of z/VM 5.2 (solution for long term time frame, as z/VM 5.2 removes the 2 GB I/O constraint)
 - Until z/VM 5.2 is operational, continue PoC and avoid using RMAN (circumvention for backup of DB, Oracle export and save complete files)



Project activities related to z/VM 5.2 ESP

- Test I/O behavior and processor load with z/VM 5.2 (ESP)
- New DASD environment allowed additional optimization of structure for Linux for zSeries (volumes distributed on different ranks, striping, use mod 3 instead of mod 9, ...)

- **Result of tuning, optimization activities and using z/VM 5.2 (ESP):**
 - **All prior performance issues resolved**
 - Memory consumption below 2 GB line now excellent, **no more bottlenecks**
 - Together with new DASD there is **good performance** – incl. RMAN !
(better than in prior Unix environment)



Memory usage when running Oracle with RMAN after tuning and with z/VM 5.2

```
FCX115      CPU 2086  SER 39D7E  Interval 13:32:36 - 13:32:38  Perf. Monitor

Detailed data for user V2T1L023
Total CPU   : 97.6%   Storage def. : 2048MB   Page fault rate: .0/s
Superv. CPU : 7.1%   Resident <2GB: 0       Page read rate : .0/s
Emulat. CPU : 90.5% Resident >2GB: 522243  Page write rate: .0/s
VF total    : ....%   Proj. WSET  : 522243   Pgs moved >2GB>: .0/s
VF overhead : ....%   Reserved pgs : 0       Main > XSTORE : .0/s
VF emulation: ....%   Locked pages : 0       XSTORE > main : .0/s
VF load rate: ....%   XSTORE dedic.: 0MB    XSTORE > DASD : .0/s
I/O rate    : 1941/s  XSTORE pages : 0       SPOOL pg reads : .0/s
DASD IO rate: 1941/s  DASD slots  : 0       SPOOL pg writes: .0/s
UR I/O rate : .0/s    IUCV X-fer/s : .0/s     MDC insert rate: .0/s
Diag. X'98' : .0/s    Share       : 2000   MDC I/O avoided: .0/s
*BLOCKIO    : .0/s    Max. share  : ...

#I/O active : 0      Active      : 79%    PSW wait : 71%    I/O act. : 5%
Stacked blk : ..    Page wait  : 0%    CF wait  : 0%    Eligible : 0%
Stat.: EME,DSC,RNBL  I/O wait   : 4%    Sim. wait: 0%    Runnable : 25%

Data Space Name      Size Mode  PgRd/s PgWr/s XRd/s XWr/s Migr/s Steal/s

Command ==> _____
F1=Help  F4=Top  F5=Bot  F7=Bkwd  F8=Fwd  F12=Return
```

During RMAN run: high I/O activity, but pages remain above 2 GB



Memory usage after running Oracle with RMAN after tuning and with z/VM 5.2

```
FCX113      CPU 2086  SER 39D7E  Interval 13:38:19 - 13:39:19  Perf. Monitor
>
Data >----- Number of Pages ----->
Spaces >
Userid      Owned >Resrvd  R<2GB  R>2GB  L<2GB  L>2GB  XSTOR  DASD  Stor  Nr of
Size       Users
>System<    .0 >    0      1 53785   1    18     0     0    407M  15
BEV30T30    0      0      0 153279  0     13     0     0    800M
DTCVSW1     0      0      8  2764   8     26     0     0    32M
MAINT       0      0      0  1507   0     0      0     0    128M
OPERATOR    0      0      0  1237   0     0      0     0    32M
OPERSYMP    0      0      0  1281   0     0      0     0    32M
PERFSVM     0      0      0  2980   0     0      0     0    64M
TCPIP       0      0      8  4127   8    166     0     0    32M
VMSERVR     0      0      0  1178   0     1      0     0    32M
VMSERVS     0      0      0  1323   0     1      0     0    64M
VMSERVU     0      0      0  1178   0     1      0     0    32M
V2T1L001    0      0      0 15775   0    13     0     1    256M
V2T1L002    0      0      0 15670   0    13     0     1    256M
V2T1L003    0      0      0 61070   0    13     0     1   2048M
V2T1L012    0      0      0 21165   0    13     0     1    256M
V2T1L023    0      0      0 522243  0    13     0     0   2048M
Select a user for user details
Command ==> _____
F1=Help F4=Top F5=Bot F7=Bkwd F8=Fwd F10=Left F11=Right F12=Return
```

After Oracle backup with RMAN: pages of Linux guests still above 2 GB



System management concept

In last phase of PoC a system management concept has been developed and implemented (using RyO and products). This covers:

- Cloning
- Software update
- Change management
- Backup (for DB as well as for file system using DFDSS, TSM and RMAN)
- User management
- Monitoring (Performance Toolkit)
- High-Availability (2 sites)



(Financial) Benefits

- One important step to a central database server on the mainframe
- Oracle and other licenses consolidation
- Rapid cloning for Linux incl. Oracle instances
- Improved maintenance and handling
- Daily backup is two times faster than before
- Synergy effect with z/OS for backup for disaster recovery (via DFDSS)
- z/OS failover capabilities (PPRC) extended now in z/VM-Linux environment
- Stable and flexible environment for Oracle Databases



Summary / Conclusion

- Successful project
- zSeries with Linux is a stable and flexible platform for Oracle databases
- Select workload carefully
- Now signed for Reference
- Outlook: Further growth of zSeries environment expected, as there is more work . . .



Lessons learned

- Management commitment and existing Linux strategy is essential for success.
- To successfully complete a Linux for zSeries project, it is not sufficient to install hardware and software. Good support and project management is essential for success.
- There are several difficult areas for any project, when a new environment is introduced:
 - Organisational aspects (People from z/OS only department, Unix only department, network only department, ... must cooperate)
 - Technical problems (new OS, Open Source questions, security and service questions)
 - Operational questions when introducing new system into existing environment (define new operation concepts and incorporate into existing environment, cooperation of different groups)
- Complete customer environment, including system management environment, has to be considered, as the new solution must be to well integrated into existing infrastructure. Customer's processes must be carefully adapted.
- „Essential for success is often a short and direct line to development laboratory.“



Additional Information

- LGT web site:
<http://www.lgt.com>
- LGT Reference customer information / Case Study:
<http://www-306.ibm.com/software/success/cssdb.nsf/cs/STRD-6QBLH3?OpenDocument&Site=gicss67fss>
ftp://ftp.software.ibm.com/common/ssi/rep_sp/8/GK124138/GK124138.PDF
- Proof of Concept report:
<http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/TD103325>