

# Running z/OS as a Second Level System on z/VM to Clone and Scale 13075

Steve McGarril

[mcgarril@us.ibm.com](mailto:mcgarril@us.ibm.com)

STG WW Mainframe Client Center

February 6<sup>th</sup>, 2013




# Scenario

- New hardware introduction – z10 to z196
- Limited resources at the home site for testing
- Future revenue growth in jeopardy w/o expansion
- How to drag and drop a production system?
- Provide a hands-on, ‘at home’ feel to the POK benchmark environment

# Who are we?




- POK Mainframe Client Center offers a combined 350+ years mainframe experience with the latest IBM Systems and Storage
- Part of an organization that also contains the EBC, HACoC, Design Center and other platform BCs
- We work with clients, business partners and ISVs to perform complex scaling and performance services, PoCs and PoTs
- Cross platform, multi-tiered environment with remote or onsite access
- Fit for purpose to suit client's requirements
- Offer a 'self-serve' zLight program for both z/OS and Linux on z
- [http://www.ibm.com/systems/services/benchmarkcenter/servers/benchmark\\_z.html](http://www.ibm.com/systems/services/benchmarkcenter/servers/benchmark_z.html)



## Why Use the z/BMC?

- Isolated environment ensures consistent and repeatable test results
- Access to SMEs where needed
- Participate in JEMT program and have access to latest mainframes and storage
- Dedicated DASD, network, processors and workload
- Can quickly build , scale, tune, validate, move or break environments

# Objectives



- High level plan:
  - Restore client's z/OS production system and IPL 2<sup>nd</sup> level
  - Run the nightly batch schedule and benchmark the elapsed time
- Billing application scalability test
  - Grow the number of test cases to allow 4x database size
  - Run batch process with the data increase and record performance statistics

## Why Run z/OS 2<sup>nd</sup> Level?




- Familiarity with disaster recovery exercises that are run 2<sup>nd</sup> level
- Resolves licensing issues with non-IBM software
- Allow client to easily set the Virtual Time Of Day clock to when their system was dumped
- Flexibility with address virtualization

## Who Can Run 2<sup>nd</sup> Level?

- z/VM
- Linux
- z/OS
- z/VSE
- TPF
- Refer to z/VM Running Guest Operating Systems

# System Components



- Hardware
  - Single LPAR on a z196 machine(2817-M80)
  - 7 general purpose processors
  - 64GB central storage + 2GB expanded storage
  - 2107 DS8700 DASD 256GB cache
- Software
  - z/VM V6.1 service level 1003 (64-bit)
  - z/OS V1.11
  - DF/DSS
  - Tivoli Workload Scheduler
  - EKM(Encryption Key Management)
- Data
  - 24TB spread over combination of 1856 3390 mod 3, mod 9 and mod27 devices
  - Consisted of both system and client data and allow for significant database growth



## Miscellaneous


- Anonymous FTP site made available to fill in missing data gaps
- VPN for remote access from client site
- Secure and private ipaddresses for z/VM logon and TSO terminal access
- 2 separate engagements as client was not prepared for building test scenarios

# Schedule

- 2-4 days to get client system up and running
- 1 day to prep the files, prep the batch job schedule and setup the system monitors.
- 1 day to run a baseline billing batch run
- 10 days to prep the files for 4M members, prep the batch job schedule and setup the system monitors.
- 1 day to run the billing batch test
- 1 day to prep for the end-of-month or first-of-month batch schedule and setup the system monitors.
- 1 day to run the month-end or first day of month batch schedule
- 1 day to erase all the data

# Three Phase Plan

- Phase I
  - Restore production z/OS system running 2<sup>nd</sup> level under z/VM
  - Setup the system to be identical to the date and time that it was dumped
  - Perform baseline testing of their nightly batch schedule and record results
  - Prepare the environment to handle significant database growth
- Phase II
  - Tests added for cases and individuals to be 4x current size
  - Update batch schedule to handle this increase
  - Run this new nightly billing test and record the results
- Phase III
  - Perform month end processing
  - Measure the impact of the 4x growth
  - Clean up environment and scrub all confidential customer data




SHARE  
Technology • Connections • Results


# Production System Restore

z/VM 6.1

MVSTEST



TCP/IP  
TCAS  
DFRMM




z/OS

EOPS  
RESVOL




3590  
TAPE  
DRIVE

# Challenges



- Initial requirement for 24TB subsystem
- Client needed to change their JCL to allow for output to be directed to DASD vs. VTS
- Encrypted tape drives and secure network required due to sensitive customer data
- Encryption Key Management(EKM) needed to decrypt above data
- Secure network necessary with system IPaddress available to be included in client's stack
  - Needed for 1<sup>st</sup> level z/VM and 2<sup>nd</sup> level z/OS

## Challenges(cont.)



- Aligning client I/O configuration with our own gen.
  - Including only necessary devices
  - When is a 3390 mod 9 device not really a 3390 mod 9?
    - ❖ No real industry standard
    - ❖ Gigabyte versus cylinder number
    - ❖ Adjustment made from a storage standpoint
- Allowing the client to take the wheel via remote access of the MVS console
- Changes in restored system date and timestamp required creation of new couple dataset.

# Gotchas



- EOPS getting disabled wait state due to dependency on specific real address
- Misleading error messages during IPL
  - *HCPGSP2627I Virt. machine placed in CP mode due to a SIGP initial CPU reset from CPU 00*
- Initial 'encryption socket errors' during tape restores
  - IECIOSxx PARMLIB member incorrect
  - – SETIOS EKM, PRIMARY=129.40.107.33:3801
- Different console address for 1 pack system vs. full blown system
- #cp logoff vs. #cp system reset of 2<sup>nd</sup> level system

## 2<sup>nd</sup> Level Steps



- ***'Whatever needs to be visible 2<sup>nd</sup> level has to be defined on the 1<sup>st</sup> level z/VM system'***
- Lotsa virtualization to map our genned addresses to client's gen
- REXX exec used for DASD attaches
- Emergency 1 pack system used to restore limited environment(DR-like exercise) with tools to help restore full system- TCAS, TCP/IP, DFRMM
- SET VTOD to match date and time of client site when system dumped. Issued prior to 2<sup>nd</sup> level IPL
  - New COUPLE dataset created to accommodate this changes
- SET CPUID xxxxx to match client's CEC serial number
- SYSTEM RESET and TERMINAL CONMODE 3270
- **I 2DDC LOADPARM 2DDCEMM**



## 2<sup>nd</sup> Level Steps (cont.)



- TERM BRKKEY PFx – Switch commands to z/VM commands but invalid for SNA/CCS terminals
- #cp vi vmmsg – routes nonpriority MVS commands from z/VM to the MVS guest
- **‘MVS COMMAND PROCESSING AVAILABLE’** - #CP vi vmmsg v cn(\*),activate
- #cp vi vmmsg r xx,warm,noreq - In response to the JES2 startup option prompt where xx is the WTO number
- Alternate option is using 2 separate terminals for console and MVS logons via CONSOLE and SPECIAL(or DEDICATE) statements

## 2<sup>nd</sup> Level Steps (cont.)



- Synchronized use of DIAL and VARY xxx,CONSOLE commands to bring up MVS console
  - From 1 PCOMM screen, clear screen and enter D MVSTEST 702
  - From z/VM console enter #CP vi vmmsg V 702,CONSOLE
  - Can now reply to WTO prompts directly on new MVS console

## 2<sup>nd</sup> Level Steps (cont.)



- Several 3270 terminals defined in the z/VM directory
- Once 2<sup>nd</sup> level system up clients could access their z/OS image directly via console and TSO logons
- Client then able to proceed with full system restore and begin actual benchmark testing
- Authority given to issue CF CPU(04),OFFLINE for right sizing the environment

# z/VM Directory Entry



```
HEALTHP DIRPART A1 F 80 1 Blks 7/25/12 Line 1 of  
==>  
USER MVSTEST 999999 45G 45G BFG  
INCLUDE CMSUSER  
OPTION MAINTCCW LNKEXCLU QUICKDSP TODENABLE  
STDEVOPT DASDSYS DATAMOVER LIBRARY CTL  
XSTORE 1M  
* IPL  
IPL CMS PARM AUTOCR  
* Define CPs for user  
MACHINE ESA 10  
CPU 00 BASE CPUID 0E39FA  
CPU 01 CPUID 0E39FA  
CPU 02 CPUID 0E39FA  
CPU 03 CPUID 0E39FA  
CPU 04 CPUID 0E39FA  
CPU 05 CPUID 0E39FA  
CPU 06 CPUID 0E39FA  
*  
* Console definition  
SPECIAL 701 3270  
* SPECIAL 702 3270  
SPECIAL 440 3270  
SPECIAL 441 3270  
SPECIAL 442 3270  
SPECIAL 443 3270
```

## Results



- Minimal z/VM overhead when running z/OS 2<sup>nd</sup> level
  - From end user standpoint z/OS appeared to be 1<sup>st</sup> level environment
- Replicated client's production workload consisting of nightly batch and record performance statistics
- During peak times the client reached 10K I/Os per second
- Data manipulated to increase data growth to between 300-400%
- Nightly batch cycle completed in 3.5 hours using 4 cps
- 4x heaviest nightly batch cycle completed in 5.75 hours
- **Engagement resulted in \$850K hardware sale!!**

## Summary

- Having adequate capacity key to cloning current client environment and allowing scalability testing
- Number of active policies increased from 1.4M to 5.2M
- Benchmark objectives and new process testing achieved
- All statistical analysis and reporting completed
- Secure Data Overwrite Service of storage drives generated IBM Certificate of completion



## Our Survey Says....

- ***‘System resource allocation more than exceeded our needs’***
- ***‘Everybody seemed very committed to making our experience at the Benchmark Center not only productive but a pleasant one’***
- ***‘Steve McGarril, my counterpart at the Poughkeepsie datacenter ....helped me learn the VM commands I would later use to IPL our z/OS guest after hours and on weekends’***