z/VM Problem Determination and Data Collection Demo

VM Workshop
June 19, 2020

—

Emily Hugenbruch
z/VM Memory Management
Senior Software Developer
ekhugen@us.ibm.com
@ekhugen

Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

CICS*	IBM*	IBM Z*	z15
Db2*	IBM (logo)*	LinuxONE	z/OS*
GDPS*	IBM Cloud Pak	WebSphere*	z/VM*
HiperSockets	ibm.com	z14*	z/VSE*

*** Registered trademarks of IBM Corporation**

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.
IT Infrastructure Library is a Registered Trade Mark of AXELOS Limited.
ITIL is a Registered Trade Mark of AXELOS Limited.
Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.
Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
Kubernetes and Container Initiative™ are registered trademark of The Linux Foundation.
Red Hat and Red Hat OpenShift are registered trademarks of Red Hat, Inc. Open
Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.
UNIX is a registered trademark of The Open Group in the United States and other countries.
VMware, the VMware logo, VMware Cloud Foundation, VMware Cloud Foundation Service, VMware vCenter Server, and VMware vSphere are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.
Other product and service names might be trademarks of IBM or other companies.

Notes:
Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.
IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.
All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.
This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.
All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.
Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.
Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.
This information provides only general descriptions of the types and portions of workloads that are eligible for execution on Specialty Engines (e.g, zIIPs, zAAPs, and IFLs) ("SEs"). IBM authorizes customers to use IBM SE only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the “Authorized Use Table for IBM Machines” provided at www.ibm.com/systems/support/machine_warranties/machine_code/aut.html (“AUT”). No other workload processing is authorized for execution on an SE. IBM offers SE at a lower price than General Processors/Central Processors because customers are authorized to use SEs only to process certain types and/or amounts of workloads as specified by IBM in the AUT.

Thanks to:

- Tim Greer
- Dan Griffith
- Steve Wilkins
- Brian Wade
- Tim Reynolds
- Bill Bitner

What to do when things go wrong

- Determine what is actually wrong
  - What you see might just be a symptom

- Example, can’t reach the guest
  - Is the Linux guest up?
  - Is networking at the Linux guest level okay?
  - Can you log on to the Linux guest, or is it hung?
  - If you can’t log on to Linux, can you log on to the z/VM userid?
  - If you can’t log on to the z/VM userid, can you log on to any userid on that system?

What to do when things go wrong
Networking problem

- The Linux guest is up and I can log on to the z/VM userid
  - Try ip to look at the networking configuration
- If that's okay, look the next level down, are devices attached?
  - Look at the virtual NICs,
  - Investigate the vSwitch
- If that's all okay, ask your networking people
  - Is it on the right subnet
  - Were there hardware issues

What to do when things go wrong
Networking problem

- The Linux guest is up, I think, but I can’t log on
  - Try from the green screen
  - You can’t get a Linux login prompt
  - Guest is likely hung
  - Get a **`SYSTEM RESTART`** or use standalone dump (more on this in a bit)
    - You will need kdump configured in order for SYSTEM RESTART to work

What to do when things go wrong
Restarting guest

- Try to restart your guest with **`FORCE`** and **`XAUTOLOG`**
- **`FORCE`** will send a signal shutdown to Linux
- Linux will have some time to respond, but if it's hung, it probably can't!
- After an amount of time, (**`QUERY SIGNAL SHUTDOWN`** to see how long) the guest will be forced off.
- Depending on the severity of the problem, it might seem like the guest is still logged on!
  - **`QUERY NAMES`** might show it
  - **`QUERY <userid>`** might show "**`LOGOFF/FORCE pending for user <userid>`**"
  - This might be a temporary issue, perhaps a device attached to the guest is misbehaving. Or it might require an IPL to clear. (see next section)

What to do when things go wrong
Guest is not logged on

I can’t log on to the z/VM userid

- I can get a logo screen and log on to another guest
- Log on to a class A guest (**`OPERATOR`**, **`MAINT`**, etc)
- Query the guest
  - Guest was logged off
    - Get the guest’s console log
    - This assumes that you had **`SPOOL CONS * START`** somewhere in the guest’s directory entry/PROFILE EXEC
    - It’s best to use tooling like Operations manager to manage your consoles, rely on spool only as a backup

What to do when things go wrong
Guest is hung

I can’t log on to the z/VM userid

- Guest is still logged on, just hung at the z/VM level
  - Get a **SNAPDUMP**
  - From a class A userid

```
SNAPDUMP PGMBKS ALL FRMTBL YES <- 7.1 only
```

    - This might result in a very large dump, especially on larger systems, but does give us the maximum amount of information

- Check your **SYSTEM CONFIG System_Userids** statement to see where dumps go (usually OPERATNS)

```
 dumpld2
HCPDLD8237I TO operand not specified.  MDMPxxxx files will be created on A disk.
Dump Found :
ORIGINID FILE CLASS RECORDS  CPY HOLD DATE  TIME     NAME      TYPE     DIST
SYSTEM   2496 D SYS 00039081 001 NONE 10/29 11:59:35 CPDUMP    CPDUMP   SNP001
Continue - Y/N ?
 y
A total of 1 file(s) (00039081 records each) will be created.
HCPDLD8210I 1 PRB00001 MDMPxxxx file(s) will be created on disk A
Continue - Y/N ?
 y
HCPDLD8213I Created PRB00001 MDMP0001 A
Ready; T=0.71/2.40 12:00:07
```

What to do when things go wrong
DUMPLD2 failed!!

- If the first attempt to load failed (maybe disk filled up), a second attempt will show no files available
- Likely in USER HOLD state from the attempted read

```
q rdr * all
ORIGINID FILE CLASS RECORDS  CPY HOLD DATE       TIME     NAME     TYPE
LGRLIN21 0010 V DMP 00036313 001 USER 2019-06-27 05:01:46 VMDUMP   FILE
Ready; T=0.01/0.01 05:21:05


change rdr 10 nohold
0000001 FILE  CHANGED
Ready; T=0.01/0.01 05:21:10
```

What to do when things go wrong
SSI is unstable

I can’t log on to any z/VM userid

- If you can get a logo screen, but it hangs at logon, make sure to try an **IDENTITY** virtual machine
  - If the SSI in a bad state and members can’t communicate, only **IDENTITY** virtual machines will be able to log on
  - Examples of **IDENTITY** virtual machines include: **MAINT, OPERATOR, OPERATNS**
  - Or log on to another member of the SSI and do a **QUERY SSI** to see the state
    - If ISFC links are down, may be able to restart them with **ACTIVATE ISLINK**

What to do when things go wrong
System is hung, HMC restart dump

I can’t log on to any z/VM userid

- Time for a PSW restart dump
- Go to the HMC
- Select your LPAR
- Recovery -> PSW Restart

- Note that this, unlike a SNAPDUMP will restart your system!

What to do when things go wrong
System is hung, SE restart dump

I can’t log on to any z/VM userid

- Time for a PSW restart dump
- Go to the SE
- Select your LPAR
- Select a processor
- CP Toolbox -> PSW Restart

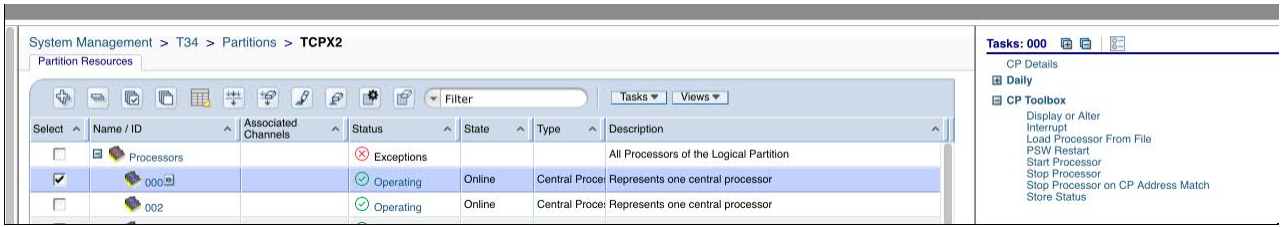

- Note that this, unlike a SNAPDUMP will restart your system!

What to do when things go wrong
Restart dump – 7.1 instructions

- Note that this, unlike a SNAPDUMP, will restart your system!
- PGMBKs and FRMTBL will be dumped or not dumped according to the DUMP settings
- So check DUMP settings now!!!!

```
q dump
01: DASD BE26 dump unit CP IPL pages 337206 PGMBKs DEFAULT FRMTBL
DEFAULT
Ready; T=0.01/0.01 05:00:07
```

What to do when things go wrong
System is hung, Standalone dump

I can’t log on to any z/VM userid, and PSW Restart dump is not working

- Time for a standalone dump!
- You need to have some media prepared for this BEFORE you have a problem
- See Chapter 11 of *CP Planning and Administration* for preparation steps
  - The SDINST EXEC usually resides on MAINT 190
  - Userid you build on must have: at least 256 MB of virtual storage, virtual reader at device number 00C, a virtual punch at 00D, no Class N reader or punch spool files, a read/write 191 minidisk that is accessed as file mode A with at least ten 4K blocks of free space, and access MAINT 400
    - SADU71 IMAGE
    - SSPJ71 IMAGE
    - SSPK71 IMAGE
    - SSPP71 IMAGE
  - Each dump device is an entire 3390 or SCSI LUN.

What to do when things go wrong
System is hung, Standalone dump

1. Go to the HMC/SE
2. **Stop the system via the Stop All function**
3. Load the system, using Load Normal
   - check the Store Status function
   - IPL from the dump device
   - Do not use the CLEAR option and erase any data that is in the Load Parameter field

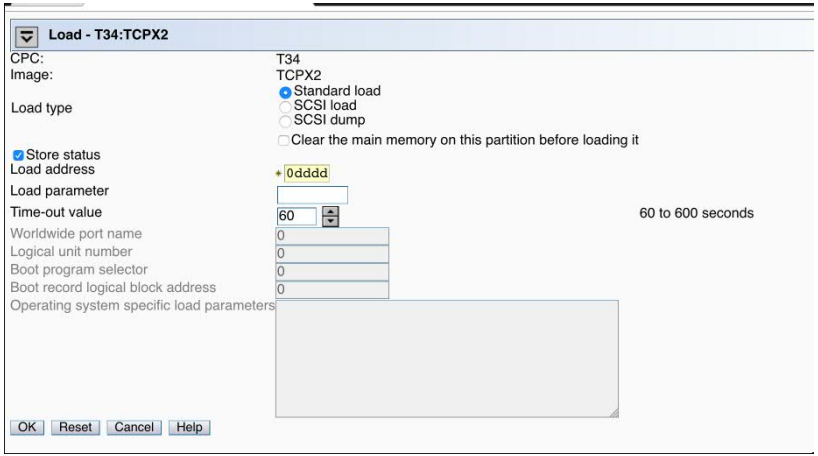

What to do when things go wrong
Collecting the dump

Log on to OPERATNS

```
 att BE06-BE09 *
BE06-BE09 ATTACHED TO OPERATNS
Ready; T=0.01/0.01 12:01:28
 dumpld2 dasd
HCPDLD8237I TO operand not specified.  MDMPxxxx files will be created on A disk.
HCPDLD8278A Enter virtual device number of first dump device.
 be06
DUMP ON BE06 WAS CREATED 2019-10-29 11:44:23
     TOTAL PAGES 59413  DUMPER RC 0

DASD  DUMPLD2         DUMP PAGES DUMP DUMP
TYPE     VDEV VOL-ID   ON DEVICE RDEV   RC
----  ------- ------  ---------- ---- ----
3390     BE06 SAD3A         59413 BE06    0
3390     BE07 SAD3B             0 BE07    0
3390     BE08 SAD3C             0 BE08    0
3390     BE09 SAD3D             0 BE09    0

A total of 1 file(s) (59413 records each) will be created.
HCPDLD8210I 1 PRB00002 MDMPxxxx file(s) will be created on disk A
Continue - Y/N ?
 y
HCPDLD8213I Created PRB00002 MDMP0001 A
Ready; T=0.96/1.85 12:01:52
```

What to do when things go wrong
System is hung, Standalone dump

- Standalone dump may dump everything; if CP memory is sufficiently damaged
  - Even guest memory!
  - Security considerations!
    - Even if it doesn’t dump everything, it will dump the object directory, unlike hard abend or snapdumps
  - Time considerations!
  - Space considerations!

Types of z/VM failures

- CP does not often fail but that does not mean that you should not be ready for an event
  - CP Abends – When CP discovers an unrecoverable error and Dumps
    - Hard - system dumps all of CP's storage and restarts
      - PGMBKs and frame table may or may not be dumped, depending on settings and type of dump
    - Soft - system dumps some pertinent information and continues to run
      - Can **`SET ABEND SNAPDUMP`** to make soft abends take a snapdump
      - Can **`SET ABEND HARD`** to make soft abends into hard ones
        - Not recommended for production!!!
- It is possible to not have enough space for a dump!
  - Use the DUMP option on CP_Owned statement to reserve SPOOL volumes for DUMP space
  - Watch out for warning messages that CP was unable to allocate dump space
  - Use **`SET DUMP`** to not dump PGMBKs or FRMTBL
    - This might not have all the information we need in the dump, then

Things that can happen to you
Wait states

- z/VM will usually restart after a problem, but in rare cases it can't
- You will see the LPAR is down on the HMC/SE:

- Check the hardware messages:

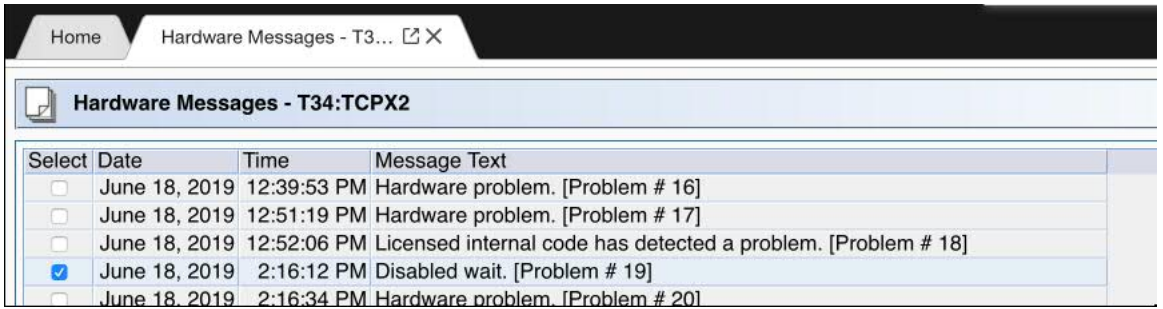

  - Look for a disabled wait (keep in mind these are messages for the whole CEC)
  - Might take a minute or two to refresh the HMC, so be patient!

You will see the disabled wait PSW. We care about the last part of this:

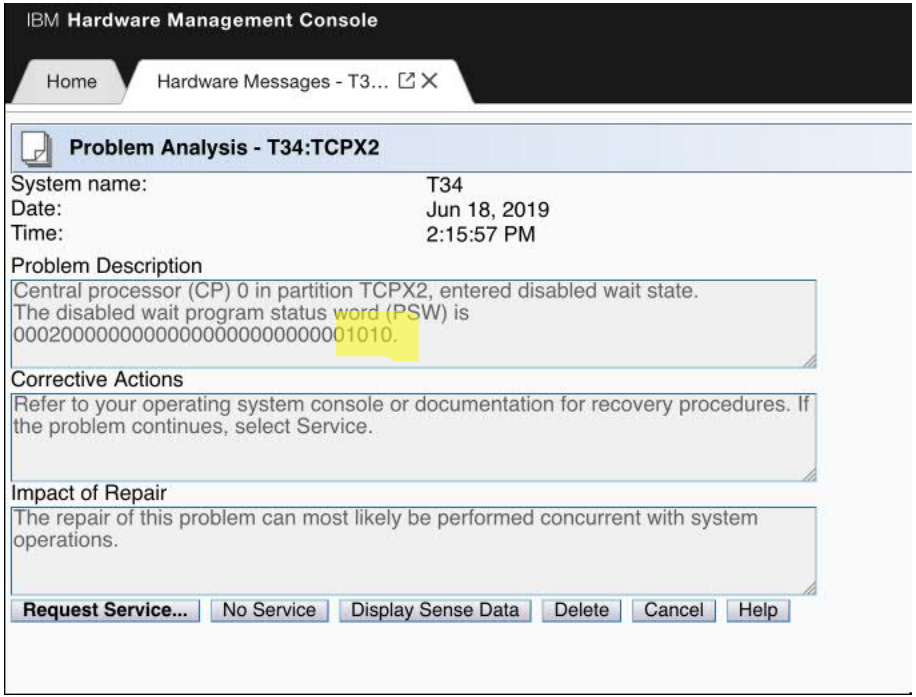

- Look up this number as HCPxxxxW
  - Always look up your wait states before getting a standalone dump!
  - Some are easy to fix problems!

```
MSG HCP1010W                                          All Help Information
(c) Copyright IBM Corporation 1990, 2019

HCP1010W  The operator's console could not be found.

Explanation:  The indicated disabled wait state PSW is loaded when an
operational operator's console cannot be found.

If you were running in virtual mode, CONMODE was not set to 3270.

System Action:  The system enters a disabled wait state (wait state code =
1010). If using the Stand-Alone Program Loader (SAPL), all registers contain
the values held at the time the wait state was entered. Register 2 contains
the console device address.

Operator Response:  There is probably a hardware malfunction. Notify your
system support personnel. Verify the OPERATOR_CONSOLE statement in the system
configuration file properly define the system consoles.

If you were running in virtual mode, enter the TERMinal CONMODE 3270 command,
and re-IPL.
```

This one means your Load Parameter was bad!

Things you need to do
Prepare for z/VM Failures

- Learn how to process a CP dump
  - DUMPLOAD or DUMPLD2 utility
    - DUMPLD2 enables you to create a multi-file dump, which is easier to transfer to IBM
- Collect the OPERATOR's console from the time of failure
- Practice moving files to and from z/VM
- Nearly every problem diagnosis starts with the same questions:
  - Description?
  - Release and service level?
  - What Changed (Workload, Service, HW, …) ???

What could go wrong?
System & Server Performance

- Slow response times
- Applications crashing
- User Hangs
- System hangs
- Just about anything!
  - Depending on what's slowed down, symptoms might be wildly different

What to do when things go wrong
Collecting monitor data

- First, monitor needs a place to put data – the MONDCSS segment
  - Check **`QUERY NSS`** to make sure it's there

```
*NSS      0011 NSS  R  0001 2011-10-07 11:18:12 MONDCSS  DCSS     MAINT620
```

- If not, create this from MAINT:
  - **`CP DEFSEG MONDCSS A000-BFFF SC RSTD`**
  - **`CP SAVESEG MONDCSS`**

What to do when things go wrong
Collecting monitor data

- Make sure you're collecting the correct information:
  - `QUERY MONITOR`
- If not, enable/disable different domains, recommended:
  - `CP MONITOR EVENT ENABLE ALL`
  - `CP MONITOR EVENT DISABLE SEEKS ALL`
  - `CP MONITOR EVENT DISABLE SCHEDULER ALL`
  - `CP MONITOR SAMPLE ENABLE ALL`
- If settings are not what you expect, adjust
  - `CP MONITOR SAMPLE CONFIG SIZE 3072`
  - `CP MONITOR SAMPLE RATE 1 SEC`
  - `CP MONITOR SAMPLE INTERVAL 1 MIN`
- Start monitor generation
  - `CP MONITOR START PARTITION 2048`

What to do when things go wrong
Collecting monitor data

- From userid MONWRITE, start monitor data collection:
  - **`MONWRITE MONDCSS *MONITOR DISK MYDATA MONDATA B`**
- You will need a REALLY large disk for this!
- Stop monitor data collection:
  - **`MONWSTOP`**

- Please don’t send us your monitor data directly!!! Please open a case for this! We <3 GDPR!

- We put some of our execs that we use to view monitor data out on the VM Downloads page: https://www.vm.ibm.com/download/packages/

What you need to do:
Keep track of System & Server Performance

- Other Material
  - Simple monitor instructions
    http://www.vm.ibm.com/devpages/bkw/monsimp.html
  - CP Planning and Administration manual – Estimation
    http://www.vm.ibm.com/library/710pdfs/71627102.pdf
  - Tivoli OMEGAMON XE on z/VM and Linux
    https://www.ibm.com/uk-en/marketplace/omegamon-on-zvm-and-linux
  - Redbook - The Virtualization Cookbook for IBM z/VM 6.3, RHEL 6.4, and SLES 11 SP3 – Section 23.3.1 http://www.redbooks.ibm.com/redbooks/pdfs/sg248147.pdf
  - z/VM: Performance Toolkit Guide Manual
    http://www.vm.ibm.com/library/710pdfs/71630201.pdf

What could go wrong
System Security Policy

- Users might not have access to things they should
  - Example: user not in LOGONBY list
- Users could get access to things they shouldn't
  - Example: User BOB can link MAINT's disk and see the full directory (this could also be catastrophic)
- Users could accidentally crash the system or pieces of it
  - Example: Thought I issued SHUTDOWN second level but I really issued SHUTDOWN first level

What you need to do
System Security Policy

- System Security is very broad and means different things to different enterprises
  - Passwords, Rules, Access Control, Granularity …
  - External Security Manager (e.g. RACF/VM)
    - NOTE: Adding an ESM to an existing SSI cluster is difficult. It is possible to do this after implementation of an SSI, but inconvenient
  - ESMs provide password encryption

- Common Criteria Certification by z/VM
  - A fully defined system
  - It may be too much for you but it gives good ideas

- Other Material
  - z/VM: Secure Configuration Guide Manual
  - VM V6.4 Achieves Common Criteria Certification http://www.vm.ibm.com/security/

What you need to do:
Storage Configuration (FICON DASD and FCP SCSI)

- Have a plan or work toward a plan for your storage configuration
  - Current needs and growth
  - Types of Storage

- Storage Allocation and Maintenance
  - Allocation (Standardization on Size and Device numbers) across LPARS
  - Settings and Error reporting
  - Duplicate VOLSER issues
  - Cylinder zero is special sometimes
    (1-END Minidisk to protect the VOLSER and allocation)
  - Advanced Configurations (e.g. GDPS)

- Other Material
  - CP Planning and Administration Manual **
  - EREP in System Operations Manual **
  - GDPS References and description page – http://www-03.ibm.com/systems/z/advantages/gdps/

What you need to do
A Planned Network Configuration

- Your physical and logical network for z/VM is key to nearly everything
  - Server and Application Connectivity
  - Transaction time and Perceptions
  - Robustness – Built-in failover
  - VSwitch and VSwitch Link aggregation is preferred
    - Lower CPU costs
    - Operates in Ethernet (Layer 2) or IP modes (Layer 3)
    - Supports port isolation
    - Supports link aggregation
  - Involve your network team!! This is really a must

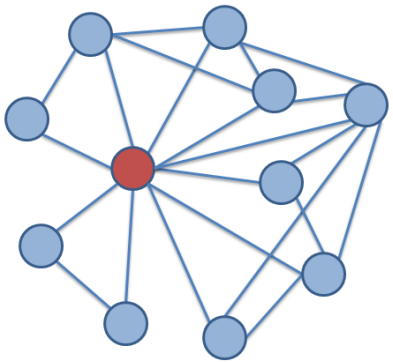

- Other Material
  - z/VM Connectivity Manual **
  - z/VM: Getting Started with Linux on System z **
  - Linux on System z Tuning Hints and Tips for Networking – https://www.ibm.com/support/knowledgecenter/fi/linuxonibm/liaag/tuning/tuning.htm

What you need to do:
BACKUP of z/VM and Server data

- Sometime and for some reason you will need to restore data on your system. Plan on this from the beginning
  - Storage Failures
  - Application failures
  - ...

- How
  - Backups of key data – File level backup (including to email) or Device level
  - Don't backup unnecessary things (paging volumes, redundant SSI data, etc.)
  - Being able to rebuild data
  - Where to backup the data to is your choice
  - Duplicate copies of data (Flash Copy, DDR)
  - Consistent USABLE data. (with I/O Quiesced)
  - TEST YOUR BACKUPs!!!!!

- CP DATA
  - SPXTAPE for Spool and System Data files
  - DDR for CP Volumes (allocation maps etc.)

What you need to do:
BACKUP of z/VM and Server data

- Other Material
  - z/VM: Getting Started with Linux on System z
  - Backup and Restore Manager for z/VM
  - Tape Manager for z/VM
  - SPXTAPE and DDR in the CP Commands Manual
  - DFSMS/VM publications in the VM Library for Tape Handling
  - Tivoli Storage Manager (now: IBM Spectrum Protect)
    https://www-01.ibm.com/support/docview.wss?uid=swg21239546#z%2FVM%20Hypervisor%20Guest

What could go wrong?
Paging

- z/VM generally expects to page, at least a little bit

- Running out of paging will cause a System Outage (PGT004 abend)
  - Messages issued by CP at 90% & 100% also at 90% of spooling space in use as a last effort

- Underpowered paging can cause issues too!
  - Performance is important, backed up paging can back up other tasks too!

What you need to do:
Manage Paging space on the system

- Paging space is not optional
  - Running out of paging will cause a System Outage (PGT004 abend)
  - Messages issued by CP at 90% & 100% also at 90% of spooling space in use as a last effort
    - Often the messages come too late for avoidance actions
  - Monitoring over time will give you a good indication
  - Commands and Tooling to watch and monitor PAGE
    - QUERY ALLOC PAGE

- Consider not only how much page space you have, but also how fast

What you need to do:
Manage Spooling space on the system

- Good – Periodic QUERY ALLOC SPOOL to see where the system is regarding SPOOL usage
  - Allocate Dedicated DUMP space
  - QUERY DUMP

- Better – Queries but also maintaining the history of usage so you can see trends
  - Run tools like SFPURGER & SPOOLPIG to determine more information
  - OMEGAMON will keep spool History

- Best – An automated solution like Operations Manager that will both visually provide the state but will Notify you if some threshold has been exceeded
  - Operations Manager can also run SFPURGER on a schedule or when thresholds have been reached
  - Catch the problem as it is changing in real time.

What you need to do:
Manage Spooling space on the system

- Other Material
  - CP Planning and Administration manual
  - SFPURGER – CMS Commands and Utilities manual
  - Operations Manager for z/VM - https://www.ibm.com/uk-en/marketplace/operations-manager-for-zvm
  - SPOOL PIG and others – z/VM Download Page https://www.vm.ibm.com/download/packages/

What could go wrong
Losing messages

- Could miss a warning message, like running out of disk space
- Could miss the reason a guest went down or restarted
- Could miss an important system warning message, like running out of paging space
- Could miss a system error message, like connectivity was lost in the SSI, or there were errors on a channel

What you need to do
Capture Important Console logs

- Good – Ensure that Spooling of logs is enabled on all servers.
  - Spooling – Set up with COMMAND statement in users Directory Entry
  - Logging in a profile or server start-up
- Better – Monitor Spooling of logs on periodic basis.
  - Close/Purge oldest and open new Console Spool or log keeping newest.
  - EXECs that may use FOR command to remotely do this
- Best – An automated solution like Operations Manager will automatically save and manage server machine consoles and logs, and optionally notify you of critical events
  - Operations Manager VIEWCON tool allows for real time viewing of events that may also make management easier

- Other Material
  - The basics on gathering a Console log in the z/VM Diagnosis guide at: https://www.vm.ibm.com/library/710pdfs/71628001.pdf
  - Operations Manager for z/VM **

What you need to do
Mainframe Social -

- Be Social – This is not a full time task but it really can help
  - Watch what is being done by others
  - Contribute your own thoughts and ideas
  - Ask Questions
  - Walking around – Virtually or Physically
- There are many avenues for material
  - List Servers, Web groups (IBMVM and LINUX-390 LISTSERVs for questions, advice, lessons learned, answers, banter, etc.)
    - Available 24 / 7 / 365
    - Relatively low traffic, low spam, little bad advice
    - Friendly, helpful, potential for lasting contacts
- Other Material
  - VM Community http://www.vm.ibm.com/techinfo/forums.html

What could go wrong
Changing SYSTEM CONFIG

- You could not have a new resource because it was defined incorrectly
- You could lose access to an old resource because its definition got corrupted
- Your system could not IPL

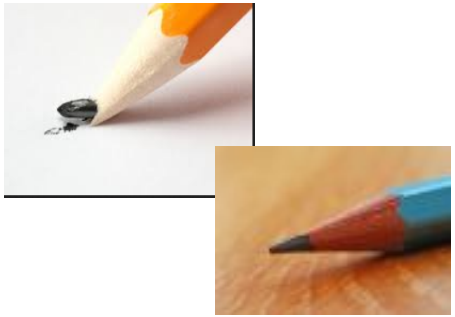

What you need to do
Changing SYSTEM CONFIG

- Develop a process for changes and stick to it. Suggested steps:
  - Make a backup copy before changing anything. This backup can be used in an emergency from the SAPL panel. Save backup in a place you can access in an emergency
  - Have a peer review your changes
  - Without fail, run CPSYNTAX !!!
    - Available on the MAINT 193 minidisk
    - An easy way to avoid embarrassing mistakes at IPL or worse
    - Easy to run – Catches incorrect and unrecognized statements
      - Even Comment Changes
- Changes not effective until next IPL
  (errors may not be discovered for months!)

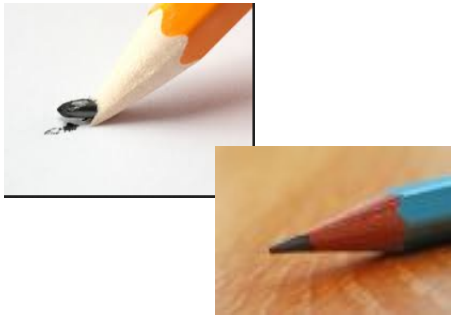

What you need to do
Changing SYSTEM CONFIG – As safe as possible

- Start-up console logs may reveal errors or problems
  - Even if a server or application starts successfully there can be issues.
  - Error messages, Warnings & overrides should be reviewed
- Critical times for reviewing logs.
  - New Releases
  - Maintenance of server or application
  - Common Error messages that could be missed
    - DASD Problems Duplicate VOLID or Offline
    - Spool Problems (e.g. NSS/DCSS …)
    - CONFIG ERRORs
- Other Material
  - CPSYNTAX Described in the CP Commands Manual **
  - CP Planning and Administration manual **
  - CP Messages and Codes Manual **

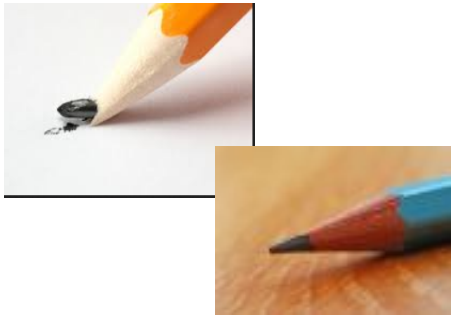

What you need to do

Maintenance is not something that can wait forever

- Apply Recommended Service Upgrade (RSU):
  - Released Periodically (6 months give or take)
  - Contains cumulative service including all pre and co-requisites in a pre-built format
  - Includes service for all integrated components and pre-installed program products
  - Available on 3590 tape, DVD, or electronically (servlink envelope)
  - Includes service required by most customer installations
  - RSUs are proven, tested, and selective
- Monitor Hiper and Red Alert APARs
- Installing:
  - SERVICE
  - PUT2PROD
- Backing out:
  - SAPL – IPL from CPLOLD MODULE
  - VMSES/E - VMFREM

What you need to do

Maintenance is not something that can wait forever

- Other Material
  - RSU Page – as needed. See: http://www.vm.ibm.com/service/rsu/
  - Alert Page -- A great place to watch for the most important items.
    To Subscribe: http://www.vm.ibm.com/service/redalert/
  - News -- http://www.vm.ibm.com/service/news
  - RSU Buckets and other maintenance is still Important
    - http://www14.software.ibm.com/webapp/set2/psearch/search?domain=psp

What you need to do
Change

- Review/Develop your long-term plan
  - What's coming in new hardware/software
  - What do you expect your workload to do over the next 3 months, 6 months, 1 year, 2 years?
  - What's nearing end of service
- Think about organizational changes, budget challenges, etc
- Talk to your IBM rep regularly

What you need to do
Change

- Review/Develop your Disaster Recovery (DR) strategy
  - DR is important in ALL environments
  - DR procedures must be adjusted for SSI members
    - DR site and Home site needs to be the same. A multi-member Home needs multi-member DR or use REPAIR MODE.
    - For desperate circumstances, there is the CLEARPDR IPL parameter on the SAPL panel
  - Some Planning now will help later
- Disaster is not well defined, but I am sure you will know when you experience one
  - TEST Your DR Plans

Thank you!

Emily Hugenbruch

ekhugen@us.ibm.com

@ekhugen