z/VM Startup and Shutdown Best Practices

February 2020

For the most current version of this presentation, please see
http://www.vm.ibm.com/library/presentations/

John Franciscovich
francisj@us.ibm.com

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

Db2*	FlashCopy*	IBM (logo)*	OMEGAMON*	z13*	z/Architecture*	zSeries*
DirMaint	FlashSystem	IBM Z*	PR/SM	z13s	zEnterprise*	z/VM*
DS8000*	GDPS*	LinuxONE*	RACF*	z14	z/OS*	z Systems*
ECKD	ibm.com	LinuxONE Emperor	System z10*	z10 BC	zSecure	
FICON*	IBM eServer	LinuxONE Rockhopper	XIV*	z10EC		

* 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.
Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
OpenStack is a trademark of OpenStack LLC. The OpenStack trademark policy is available on the OpenStack website.
Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.
Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.
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.

Topics

- Describe best practices for starting and shutting down your z/VM system
  - Things to check before IPL
  - Starting (IPLing) your system
  - After IPL
    - Validating system resources
    - Starting servers and guests
  - Shutting down
    - Guests
    - z/VM system

IBM Solutions to Help Manage and Protect Your z/VM System(s)

- Security and User Management
  - RACF Security Server for z/VM
  - Directory Maintenance Facility for z/VM (DirMaint)
- Automation and operational monitoring
  - Operations Manager for z/VM
  - Programmable Operator Facility (PROP)
- Performance monitoring
  - OMEGAMON XE on z/VM and Linux
  - Performance Toolkit for z/VM
- Backup and recovery
  - Backup and Restore Manager for z/VM
  - Tape Manager for z/VM
  - Spectrum Protect (aka Tivoli Storage Manager)
- Interactive provisioning and system resource management
  - IBM Wave for z/VM

Before Starting Your System

Verify Required Disks

- *IPL volume*
  - Where the loader program (SAPL) is
- *"RES" volume*
  - Contains minidisk (CF1 parm disk) where the CP MODULE is
  - Warmstart and checkpoint areas (CP-formatted)
  - Object directory
- *Default IPL/RES volume:*

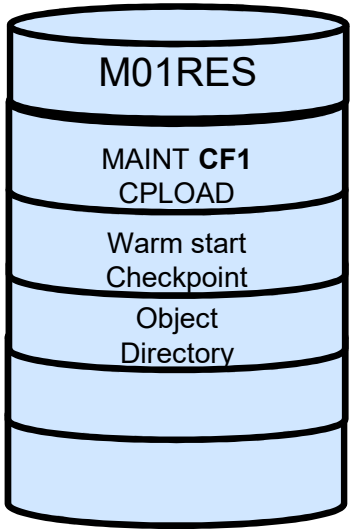

- *CF0 Parm disk*
  - Where the system configuration file is
- *Default CF0 Parm disk:*

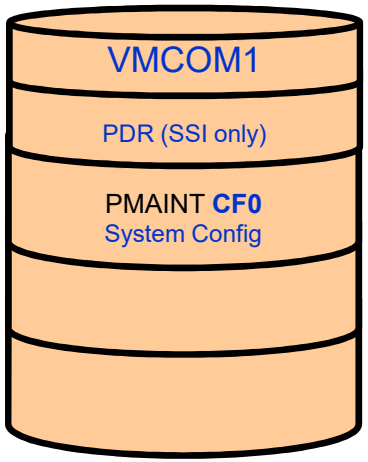

- *Other system and CP-Owned volumes*
  - Page
  - Spool
  - Guest and user data

Verify I/O Configuration

- Helpful tool: **IBM I/O Exerciser for System z (ESAIO)**
- Helps identify cabling or definition errors
  - Validates all paths to every device that the LPAR has access to
- Can be run standalone in LPAR (with SAPL) or on a running z/VM system
- See *z/VM System Operation* (chapter 3) for details

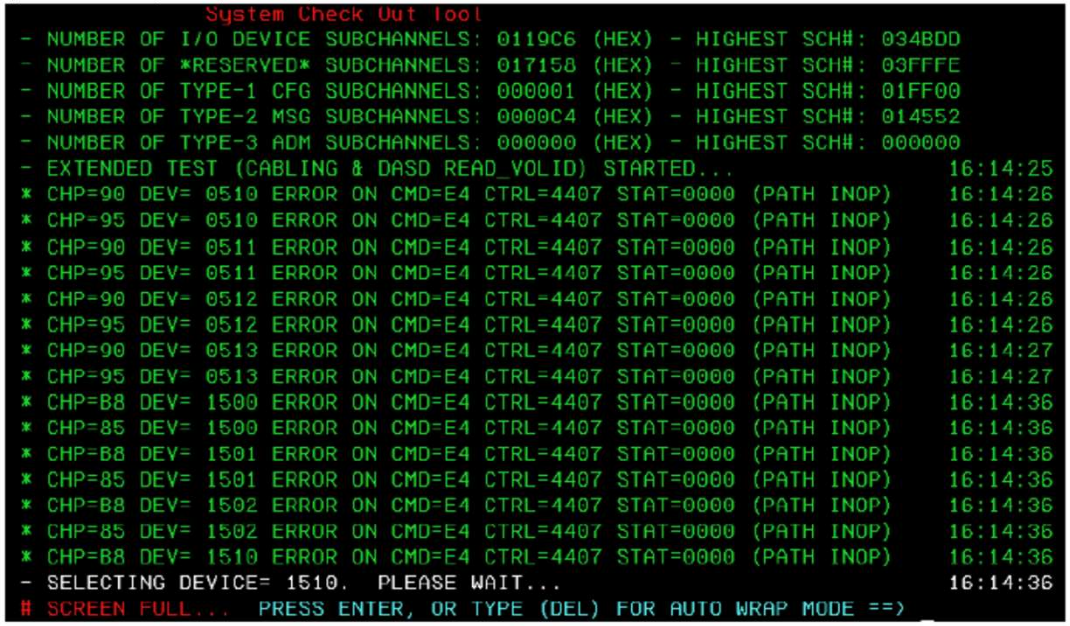

Define IPL Environment and Parameters

- Do you want to set the CP system environment variable (**CP.IPLPARMS.IPLVAR**)?
  - For example:
    - *Production* vs. *Test* vs. *DR*
    - *"Normal" startup vs. System Maintenance/Service*
  - Can be used by automation to determine actions based on system environment/use

- Do you need to display the **SAPL** screen to verify or change default characteristics for this IPL?
  - IPL Parameters
    - Which system configuration file to use
    - Which parm disk to use
    - Which CP module to load
    - System/operator's console
    - Value of the CP system environment variable (CP.IPLPARMS.IPLVAR)
    - etc.

Single System Image (SSI) Activation

- Do you plan to add a new member to your SSI cluster?

  1. Update the I/O configuration to add CTC pairs to the new and existing LPARs.

  2. Update system config file with **ACTIVATE ISLINK** statements for the new member.
     - existing members -> new member
     - new member -> existing members

  3. Activate and IPL the new member.

  4. Issue **ACTIVATE ISLINK** commands on the existing members to connect to the new member

- Do you need to reinitialize the PDR?
  – All members must be down

Specify Location of System/Operator Console

- Operating System Messages Panel of HMC
  - LOADPARM **CONSSYSC**
  - **CONS=SYSC** IPL parameter
  - **SYSTEM_CONSOLE** on OPERATOR_CONSOLES config statement

- Integrated 3270 Console
  - LOADPARM **CONSSYSG**
  - **CONS=SYSG** IPL parameter
  - **SYSTEM_3270** on OPERATOR_CONSOLES config statement

- Device specified by address on OPERATOR_CONSOLES statement

- Selected system console becomes the primary system operator when OPERATOR virtual machine is logged on during IPL

Designate Ownership of CP-owned Volumes?

- Ownership information may be specified for CP-owned volumes
  - Prevents a system from writing CP data on a volume owned by another system
- Ownership definition is mandatory in an SSI cluster, optional otherwise
- Specified with **OWNER** operand of CPFMTXA
  - Owning system name
  - Name of the SSI cluster if system is an SSI member
  - **NOSSI** if not an SSI member

```
q cpowned
 Slot   Vol-ID   Rdev   Type    Status                SSIOwner  SysOwner
    1   ------   ----   -----   Reserved              --------  --------
    2   JFEFE0   EFE0   Own     Online and attached   --------  --------
    .
    8   J1E000   E000   Own     Online and attached   JFSSIA    JFSSIA1
    .
   12   J2E000   E002   Share   Online and attached   JFSSIA    JFSSIA2
    .
   16   J3E000   E003   Share   Online and attached   JFSSIA    JFSSIA3
    .
   20   J4E000   E004   Share   Online and attached   JFSSIA    JFSSIA4
```

Network Definitions

- Define Inter-VSwitch Links (IVLs) in system configuration file
  - Must be defined before any global VSwitches
  - Verify that it is operational and its Uplink port is ready before defining global VSwitches
- Define other VSwitches dynamically after IPL
  - Can use automation to define and verify

Where to put Automation?

- Locations:
  - AUTOLOG1 profile
  - AUTOLOG2 profile
  - Server virtual machine's profile
  - User Directory through COMMAND statements
  - Operations Manager
  - IBM Wave

- Think about prerequisites

- Thoroughly document

System Config File Considerations

- Whenever the system config file is changed, keep several previous versions on the same disk as backup
  - Easier recovery if changes cause a problem
  - Use a naming convention such as -1SYSTEM, -2SYSTEM, etc.

- Run **CPSYNTAX** after every config file change to prevent errors during IPL

- System Config file controls for errors
  - TOLERATE_CONFIG_ERRORs YES (or NO)
  - If NO and there are errors, Operator is prompted to:
    - Stop IPL
    - Continue IPL normally
    - Continue IPL, but do not autolog any virtual machines

Starting Your z/VM System

Stand Alone Program Loader (SAPL)

- "What you really IPL"
- **Loads and starts z/VM**
  - Can also load and start standalone utilities
- Default SAPL is created during z/VM installation
  - Can rebuild and customize using SALIPL utility
- Menu can optionally be displayed when IPLing z/VM
  - Specify/change IPL attributes and IPL parameters
- To display the SAPL menu during IPL
  - Specify a load parameter with the address of a 3270 console
    - Load parameter **SYSG** will cause SAPL to be displayed on integrated 3270 console
- See *z/VM System Operations* (Chapter 2) for more information

Loading and Starting z/VM with SAPL

```
STAND ALONE PROGRAM LOADER: z/VM VERSION 6 RELEASE 4.0
DEVICE NUMBER: 018B MINIDISK OFFSET: 35 EXTENT: -
MODULE NAME: CPLOAD LOAD ORIGIN: 2000
--------------------------------IPL PARAMETERS----------------------
cons=0080 iplvar=PRODUCTION fn=MYPROD
-----------------------------------COMMENTS-------------------------


------------------------------------------------------------------------
9= FILELIST 10= LOAD 11= TOGGLE EXTENT/OFFSET
```

Migrating z/VM from z13 or earlier to an IBM z15 or z14/etc.

- The Stand Alone Program Loader (SAPL) must be rewritten with the z/VM 6.4 or 7.1 SALIPL utility
  - Otherwise you will not be able to IPL
  - Look for current release number in upper right corner of SAPL

- Installation Upgrade does not rewrite SAPL
  - Must be done manually
- See red alert http://www.vm.ibm.com/service/redalert/index.html#SAPLZ14

```
STAND ALONE PROGRAM LOADER: z/VM VERSION 6 RELEASE 4.0
DEVICE NUMBER: 018B MINIDISK OFFSET: 35 EXTENT: -
MODULE NAME: CPLOAD LOAD ORIGIN: 2000
```

- Other stand alone utilities also need to be updated in order to IPL on z14 or later
  - Standalone Dump, DDR, etc.

How Do You IPL?

- Monitoring/responding to messages and prompts
- Automated/not monitored

- Type of IPL
  - Can automate type of IPL on the FEATURES statement in system configuration file …
    - **AUTO_WARM_IPL** (preferred) or AUTO_FORCE_IPL
    - **AUTO_IPL *type***
  - … or reply to prompt

```
Start ((Warm|Force|COLD|CLEAN) (DRain) (DIsable)  (NODIRect)
      (NOAUTOlog)) or (SHUTDOWN)
```

  - Can also disable autolog or terminate IPL (shutdown) in case of errors

IPL Issues

- Watch for errors during IPL
  - Errors in IPL parameter processing
  - Errors in configuration file processing
  - CP-owned volumes that are offline
    - Could result in loss of data (spool files, saved systems, and saved segments)
  - Warmstart or checkpoint errors
    - Might prompt to change to FORCE start
      - Could also result in loss of data

- Decide whether to
  - Stop IPL to fix any errors
  - Continue/complete IPL and attempt to fix them dynamically

- Some errors will result in a wait state

After IPL: Validation, Servers, and Guests

Information About Current IPL

- What caused the system to IPL?
  - System ("normal") IPL
  - SHUTDOWN REIPL
  - Restart from abend or PSW RESTART

- QUERY CPLOAD provides information about the most recent IPL

```
Module CPLOAD was loaded from minidisk on volume IPLDSK at cylinder 0.
Parm disk on volume IPLDSK, cylinders 0 through 79.
Last start was a system IPL.
```

ESM and Service Machines

- Make sure ESM is started before any other service machines
  - Change AUTOLOG1 userid's PROFILE EXEC to autolog **only** RACFVM
  - Move everything else in AUTOLOG1's PROFILE to AUTOLOG2
  - RACF will autolog AUTOLOG2 after it is initialized

- AUTOLOG2: start service machines for system operation and monitoring, including:
  1. Operations Manager, PROP
     - Start these first so they can monitor other service machines (consoles, etc.)
     - Use to verify service machines and manage dependencies among them
  2. DirMaint
  3. MONITOR service, Performance Toolkit, OMEGAMON XE
  4. VSwitch Controllers (DTCVSW1, DTCVSW2, DTCVSW3, DTCVSW4)
  5. z/VM TCP/IP
  6. etc.

Networks and Connections

- Create a "networking script" to set up network definitions and verification in AUTOLOG2

  1. If you have an IVL, verify that it is operational and its Uplink port is ready

     SET VSWITCH IVL IVLPORT PING ALL

  2. Define additional VSwitches and verify that they are operational

     e.g. **DEFINE VSWITCH *name* GLOBAL ETHERNET GROUP *group***

- Virtual networks must be operational before starting the guests that use them
  - Recommend automation to defer guest startup
    - Or add polling logic to AUTOLOG2 that delays dependent guests until Global VSwitch is up.

System Resources

- Use QUERY commands to verify:
  - Available processors
  - How much real memory is available?
    - Make sure all memory is defined as central storage (no expanded storage)
  - Page and spool volumes and available space
  - Dedicated devices for virtual networks are defined and online

- Make sure data is being collected in event of a system failure
  - Console files for Operator and key service virtual machines
    - Retrieve and save regularly to avoid filling spool
  - CP dumps are enabled
    - Make sure enough dump space is allocated on spool volumes

- *Use automation to verify and monitor the above*
  - Define usage thresholds and limits that result in notifications and/or automated actions

Starting Guests

- Start database servers before application servers that are dependent on the database
  - Validate that servers are ready
- "Pace" guest startups to minimize contention for system resources
  - Depends on workloads and guest requirements
    - CPU
    - Memory
    - Network connections and communication
- Verify each group of guests have completed startup before starting next group
- Pacing and verification can be intelligent or simple
  - Automation/verification with system management tools
  - Groups in CMS names file
  - CP SLEEP

Verifying and Changing Guest Resources

- **CP FOR** command can be used to issue commands on a virtual server or guest
- Helpful for querying and/or dynamically changing guest resources
- Response is sent to **CP FOR** issuer

```
cp for lxs00001 cmd cp q v nic

LXS00001 : Adapter 0100.P00 Type: QDIO      Name: UNASSIGNED  Devices: 3
LXS00001 :   MAC: 02-2E-40-00-04-BA         LAN: * None
LXS00001 : HCPFOR069I Command Complete.  CP return code = 0000.
```

Choosing Startup Actions

Why Did We IPL?

- "Normal"
  - Start system servers, applications, networks, etc.
  - Start regular guest workload

- System maintenance
  - Update system resources, configuration, etc.
  - Possibly start system servers and applications
  - Do not start regular guest workload

- Apply service
  - CP (requires system IPL)
  - Other components
  - Do not start system servers and regular guest workload

Choose Startup Actions Based on Reason for IPL

- Define CP environment variables to indicate:
  - Normal IPL and system startup
  - IPL to apply service or do other maintenance
    - Do not start services or regular workload

- Use these environment variables
  - **CP.IPLPARMS.IPLVAR**
    - (CP variable defined as an IPL parameter)
  - **CURRENT.IPLVAR**
    - Define and initialize in system config file

```
SET VARIABLE SYSTEM CURRENT.IPLVAR INITIAL
```

    - Can be changed to allow additional or different actions

Choose Startup Actions Based on Reason for IPL – Example (1 of 3)

- Assumes ESM (RACF) will always be started
- Possible values for **CP.IPLPARMS.IPLVAR**
  - *null* - normal IPL
  - *'NORMAL'* – normal IPL
  - *'NOWORKLOAD'* – service/maintenance IPL
- In AUTOLOG2's PROFILE EXEC:
  - If **CURRENT.IPLVAR** has its initial value, change it according to **CP.IPLPARMS.IPLVAR**

```
/*************************************************************/
/* If CURRENT.IPLVAR contains its initial setting from       */
/* system config, set it based on CP.IPLPARMS.IPLVAR         */
/*************************************************************/

Parse value diag(8,'QUERY VARIABLE SYSTEM NAME CURRENT.IPLVAR') ,
                 with curr_iplvar '15'x

If curr_iplvar = "INITIAL" Then Do
   Parse value diag(8,'QUERY VARIABLE SYSTEM NAME CP.IPLPARMS.IPLVAR') ,
                    with iplvar '15'x
   If iplvar <> '' Then
      'CP SET VARIABLE SYSTEM CURRENT.IPLVAR' iplvar
   Else
      'CP SET VARIABLE SYSTEM CURRENT.IPLVAR NORMAL'
End
```

Choose Startup Actions Based on Reason for IPL – Example (2 of 3)

- In AUTOLOG2's PROFILE EXEC (continued):
  - Check updated **CURRENT.IPLVAR** to determine which actions to take

```
/*************************************************************/
/* Determine actions from new value of CURRENT.IPLVAR        */
/*************************************************************/

Parse value diag(8,'QUERY VARIABLE SYSTEM NAME CURRENT.IPLVAR') ,
                   with curr_iplvar '15'x
```

  - If "**NOWORKLOAD**"
    1. Exit without starting servers or guests

```
/*************************************************************/
/* If IPL'd for maintenance, exit without starting workload  */
/*************************************************************/

If curr_iplvar = 'NOWORKLOAD' Then Do
   Say 'IPLed for service, no servers or workload started'
   Exit
End
```

    2. Do maintenance
    3. Update **CURRENT.IPLVAR** setting

```
SET VARIABLE SYSTEM CURRENT.IPLVAR NORMAL
```

    4. Restart PROFILE EXEC

Choose Startup Actions Based on Reason for IPL – Example (3 of 3)

- In AUTOLOG2's PROFILE EXEC (continued):
  - If **"NORMAL",** start servers, applications, and guest workload

```
/*****************************************************************/
/* If normal IPL, start regular workload                         */
/*****************************************************************/
If curr_iplvar = 'NORMAL' Then
   Say 'All is well - starting normal workload'

/***** Start servers, applications, and guest workload *****/
```

Shutting Down

Shutting Down Guest(s)

- Define order of guest shutdowns
  - Application servers before database servers
  - Shut down monitoring service machines last
    - Can be used to automate guest shutdowns

- Orderly shutdown of guests
  - Maintain data integrity
  - Avoid dirty file system on Linux guests

Shutting Down Guest(s)

- Enable guests to receive shutdown signals

- Guests that are enabled will receive a "termination" signal when
  - System is shutting down
  - **FORCE** command is issued to log the guest off
  - **SIGNAL SHUTDOWN** command is sent to the guest

- Guests that receive signals will have a designated amount of time to shutdown and logoff
  - SHUTDOWNTIME interval for the system
  - Interval or time of day designated when the signal is sent

- SHUTTRAP tool in VM Download Library
  - Detect a SHUTDOWN signal from CP and optionally execute a CMS command when such a signal is received.

Shutting Down Your z/VM System

- Define the amount of time reserved for shutdown
  - For signaled guests
  - For CP
    - System configuration file

```
Set ,
 Signal ShutdownTime 500,
 ShutdownTime 30 /* amount of time reserved for z/VM shutdown
```

    - **SET SIGNAL** and **SET SHUTDOWNTIME** commands

  - **ShutdownTime** is used when CP SHUTDOWN command is issued or the LPAR is deactivated

Summary

- "Best practices" will help startup and shutdown of your z/VM systems and workload go smoothly

- Specific steps will be different for every system
  - Company policies
  - System capacity and workloads

- Using documented steps and procedures makes it easier when
  - Applying service or upgrading your system
  - Transferring processes to someone else

- Using tools to monitor and automate startup and shutdown processes is recommended
  - Detect completion of one process and start the next
  - Problem detection and notification or remedy