Using New CP Features in z/VM 5.3

Session R16 – WAVV – Chattanooga, TN
April, 2008

Romney White

romneyw@us.ibm.com
IBM System z Software – Strategy and Design

IBM Systems

Trademarks

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

DB2*
DB2 Connect
DB2 Universal Database
e-business logo
GDPS*
Geographically Dispersed Parallel Sysplex
HyperSwap
IBM*
IBM eServer
IBM logo*
Parallel Sysplex*
System z
Tivoli*
VM/ESA*
WebSphere*
z/OS*
z/VM*
zSeries*

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Intel is a registered trademark of the Intel Corporation in the United States, other countries or both.
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.
SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC.

* All other products may be trademarks or registered trademarks of their respective 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.

Agenda

- **Product Evolution**
- **z/VM System z10 Compatibility Support**
- **z/VM 5.3 CP Enhancements Overview**
- **z/VM 5.3 CP Enhancements**
  - **Scalability and Constraint Relief**
  - **Virtualization**
  - **Security**
  - **Virtual Networking**
  - **Systems Management**
  - **Miscellaneous**
- **Statements of Direction**

z/VM Evolution

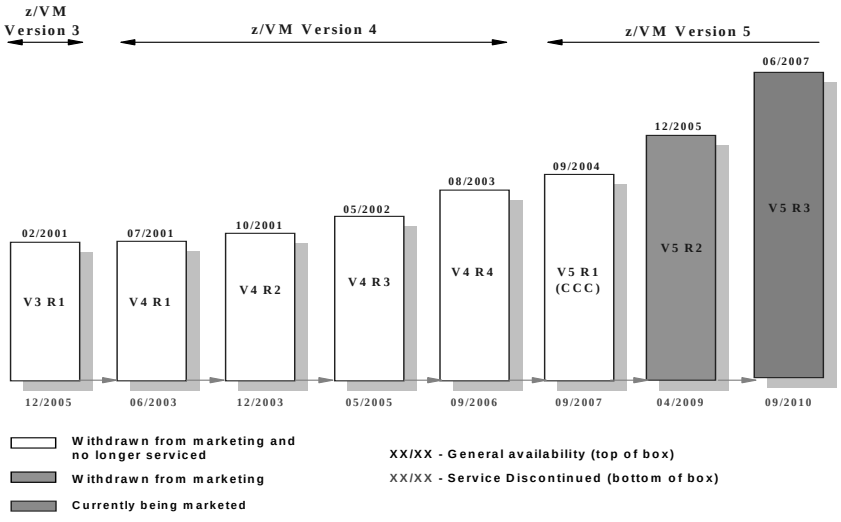

z/VM System z10 Compatibility Support

- **Processor Compatibility (VM64180, UM32277/32278)**
  - **Guest support for Execute-Extensions Facility**
- **IOP Subchannel Recovery (VM64242, UM32285/32286)**
- **Memory Management Improvement (VM64349, UM32250/32251)**
- **Encryption Re-Key Support – z/VM 5.3 only (VM64260, UM32221)**
  - **SET TAPE REKEY**
- **Virtual Switch Port Isolation (VM64281, UM32269/32270)**
  - **Isolate guests on VLAN-unaware virtual switch (DROP or FORWARD inter-guest traffic)**
- **PerfKit OMEGAMON Enhancements (VM64337, UM32233/32234)**
- **PerfKit Enhancements (VM64369, UM32257/32258)**
- **IOCP Support (VM64302, UM32279/32280)**

z/VM 5.3 CP Enhancements Overview

- **Scalability and Constraint Relief**
  - **Support 256 GB of real storage and 8 TB of virtual storage**
  - **Allow 32 real processors in a single z/VM image**
  - **Collaborative memory management**
  - **HyperPAV support for IBM System Storage DS8000**
  - **Enhanced FlashCopy support**
  - **SAN Volume Controller support**
- **Security**
  - **Enhanced system security with longer passwords**
  - **Tape data protection with support for encryption**
- **Systems Management**
  - **Enhanced guest configuration**
  - **Asynchronous CP command responses**
  - **VM event notification**
  - **z/VM integrated systems management**
- **Virtualization**
  - **Guest support for IBM System z specialty engines**
  - **Enhanced Virtual Switch and Guest LAN usability**
  - **Guest support for Modified Indirect Data Address Words (MIDAWs)**
  - **Guest ASCII console support**
  - **Enhanced SCSI support**
- **Virtual Networking**
  - **Improved virtual network management**
  - **Link aggregation**
- **Miscellaneous**
  - **Shutdown message time stamps**
  - **SYSEVENT Query Virtual Server**
  - **TRSOURCE for LDEVs**
  - **QUERY IUCV**
  - **DS8000 Dynamic Volume Expansion toleration**

Scalability and Constraint Relief

- **Support 256 GB of real storage and 8 TB of virtual storage**
- **Allow 32 real processors in a single z/VM image**
- **Collaborative memory management**
- **HyperPAV support for IBM System Storage DS8000**
- **Enhanced FlashCopy support**
- **SAN Volume Controller support**

Scalability and Constraint Relief …

- **Support 256 GB of Central Storage and 8 TB of virtual storage**
  - **PGMBKs (Page Management Blocks) moved above 2G**
    - **Relieves constraint on total amount of virtual storage**
  - **Contiguous frame management improvements**
- **Expanded storage limit remains at 128 GB**

Scalability and Constraint Relief …

- **Support 32 Processors**
  - **Actually extended to 64 processors**
    - **CPU bit masks expanded**
    - **Scheduler lock redesigned**
      - **Shared/exclusive spin lock**
      - **Allows more concurrent scheduling activity**
  - **Performance degradation observed between 36-42 processors**
    - **Dependent on work load characteristics**
  - **New spin locks monitor record**
    - **MRSYTLCK – Formal Spin Lock Data (Sample)**
      - **Data for 26 spin locks**
      - **Extensible format**

Scalability and Constraint Relief …

- **Collaborative Memory Management (CMM)**
  - **Coordinates memory state and page management between Linux and z/VM at the level of individual pages**
  - **Exploits Collaborative Memory Management Assist (CMMA) on System z9 EC and BC and System z10**
    - **New Extract and Set Storage Attributes (ESSA) instruction**
    - **Exploits Host Page-Management Assist (HPMA)**
  - **Linux exploitation under discussion with Open Source community**
    - **Some progress at July 2007 Ottawa Linux Summit**

Collaborative Memory Management (CMM)

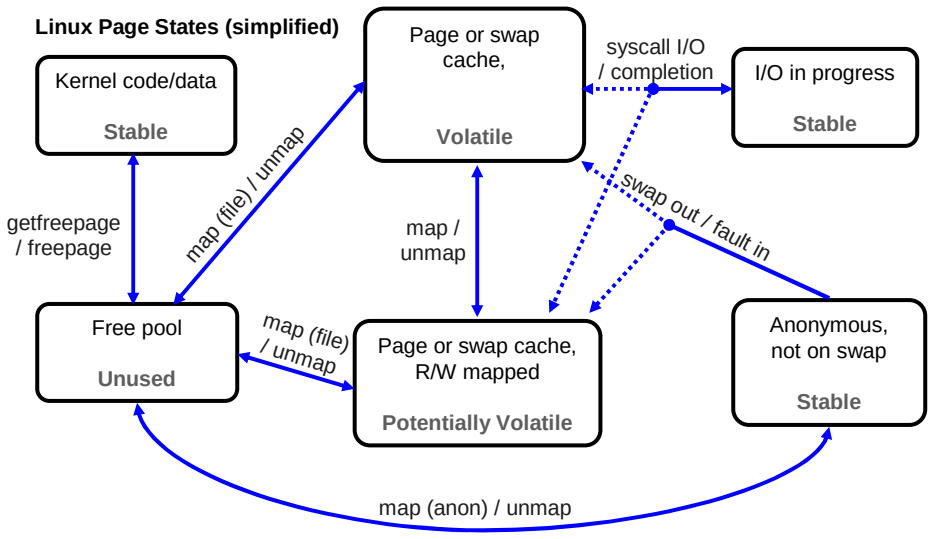

Scalability and Constraint Relief …

- **DS8000 HyperPAV Support**
  - **Dynamic alias assignment**
    - **Only for duration of I/O operation**
    - **Reduces number of real alias device addresses required**
    - **Allows better reaction to shifting bandwidth requirements**
  - **Guest support for HyperPAV-aware guests (z/OS)**
  - **Host support for HyperPAV-unaware guests (everyone else)**
  - **Base and alias devices belong to one of up to 16,000 pools**
  - **Number of virtual aliases limited to**
    `min(254,aliases in base device pool)`

HyperPAV Support

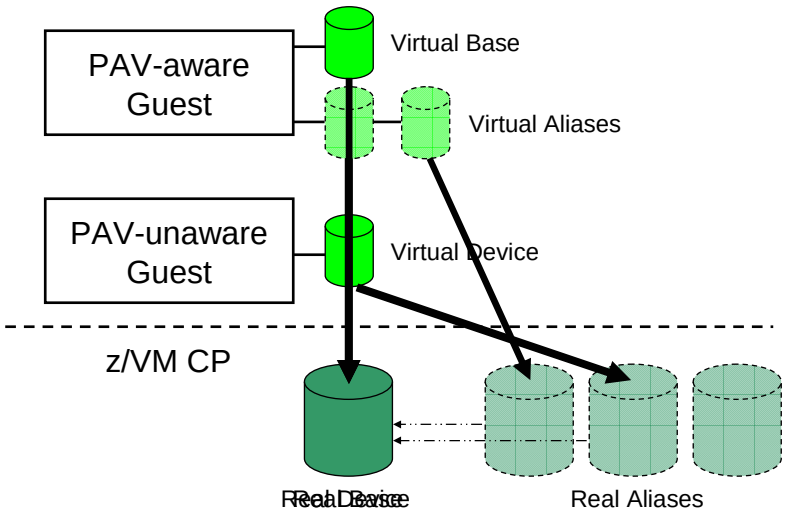

HyperPAV Support ...

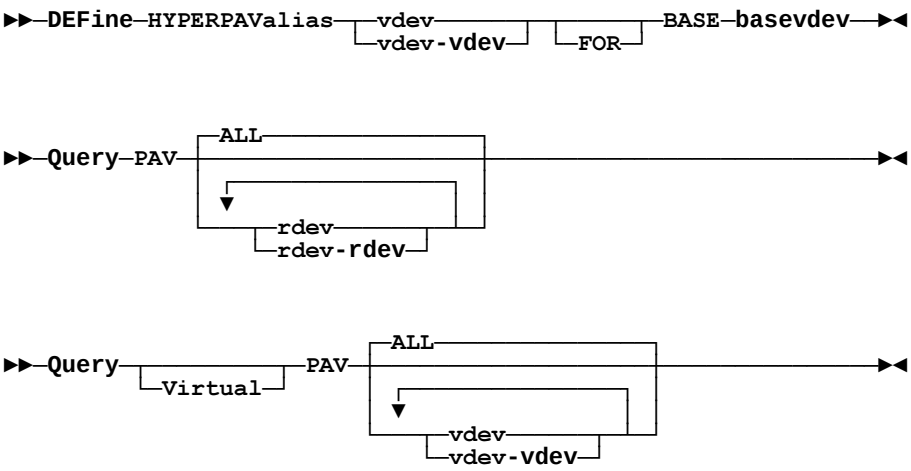

HyperPAV Support ...

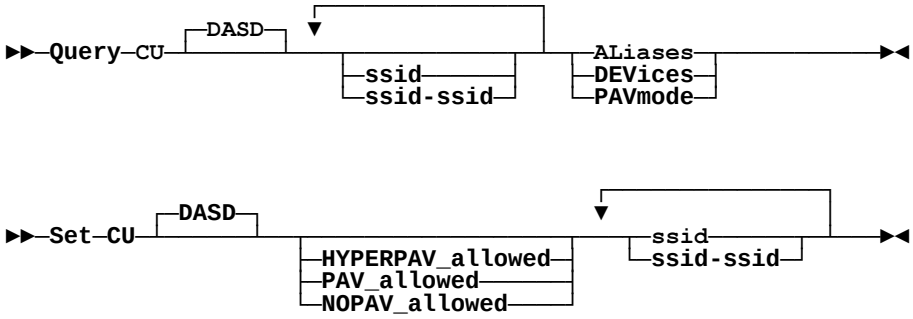

HyperPAV Support ...

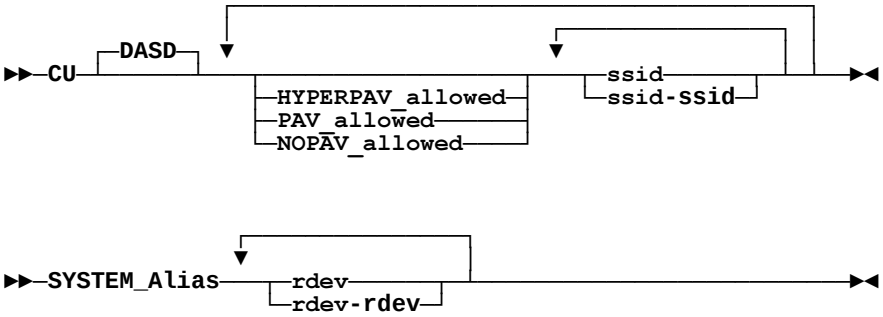

Scalability and Constraint Relief …

- **Enhanced FlashCopy Support**
  - **Multiple FLASHCOPY targets (up to 12)**
    - **Requires FlashCopy Version 2**
  - **New SYNCHRONOUS option**
    - **Automatic for FlashCopy Version 2**
    - **Simulated for FlashCopy Version 1**
  - **New QUERY VIRTUAL FLASHCOPY command**
  - **Improved error handling**
  - **CPHX can be used to terminate FLASHCOPY retry wait loop**

FLASHCOPY Enhancements

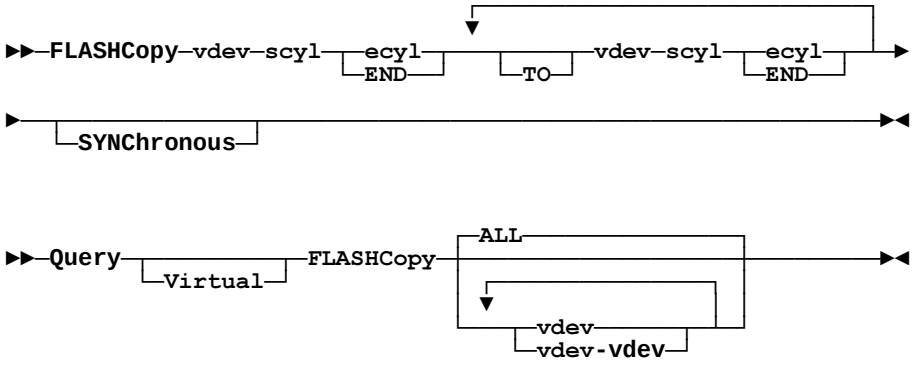

Scalability and Constraint Relief …

- **SAN Volume Controller (SVC) Support**
  - **IBM TotalStorage SAN Volume Controller**
    - **Provides homogeneous support for a wide variety of SCSI storage controllers**
    - **Image or virtualizing mode**
    - **Provides other facilities that underlying hardware might not support**
      - **E.g., copy services**
  - **New emulated device type (2145)**
  - **QUERY EDEVICE reports "2145"**
  - **Monitor record provides 2145 attribute**
    - **MRMTRDEV – Device Configuration Data (Sample Configuration)**
  - **Support also provided for z/VM 5.2 via SPE APAR VM64128**

SAN Volume Controller Support

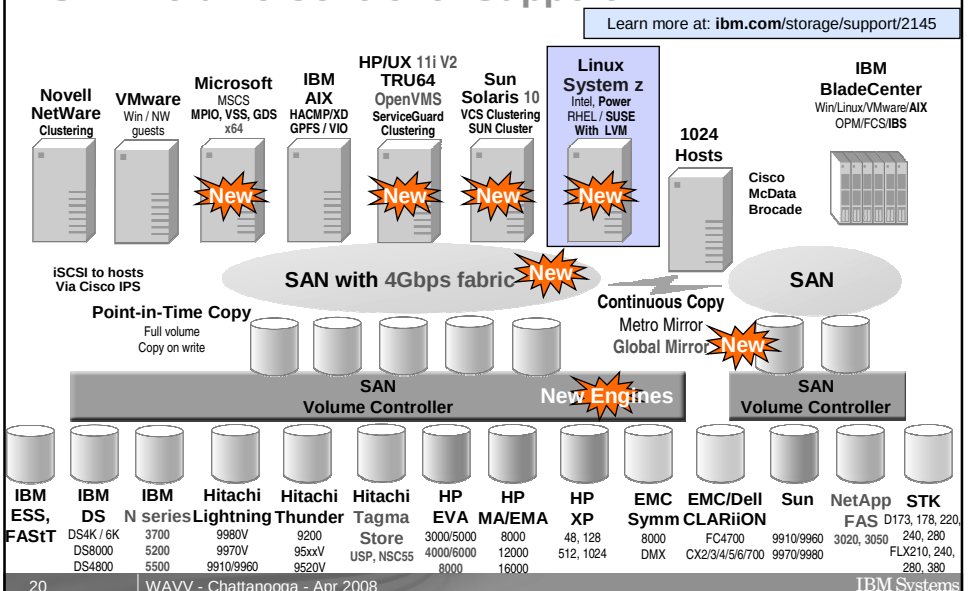

SAN Volume Controller Support …

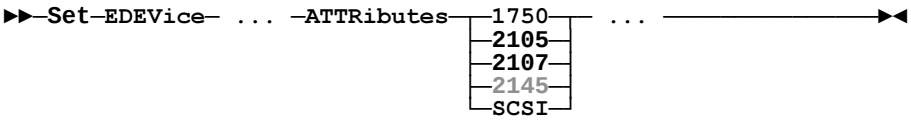

Note:
1750 = DS6000
2105 = TotalStorage ESS
2107 = DS8000
2145 = SVC
SCSI = generic

Virtualization

- **Guest support for IBM System z specialty engines**
- **Enhanced VSWITCH and Guest LAN usability**
- **Guest support for Modified Indirect Data Address Words (MIDAWs)**
- **Guest ASCII console support**
- **Enhanced SCSI support**

Virtualization …

- **Guest Specialty Engine Support**
  - **In ESA/390 logical partitions, virtual machines may have virtual IFLs or virtual CPs**
    - **Virtual machine with virtual CP(s) can have virtual zIIP(s) and zAAP(s)**
    - **zAAP support requires z890, z990, System z9, or System z10**
    - **zIIP support requires System z9 or System z10**
  - **In Linux-only logical partitions with CPs, virtual machines may have virtual IFLs or virtual CPs**
  - **In Linux-only logical partitions with IFLs, virtual machines have virtual IFLs**
  - **Simulation support dispatches virtual IFLs, zIIPs, and zAAPs on real CPs**
  - **Virtualization support can dispatch them on the corresponding real engine type, if available**
    - **CPU affinity can be controlled**
  - **Primary real processor type is type of IPL processor (CP or IFL)**
  - **Described in new section of *Running Guest Operating Systems***

Guest Specialty Engine Support

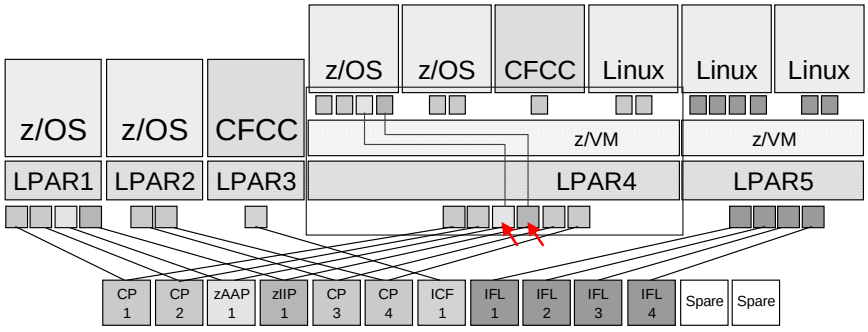

Guest Specialty Engine Support …

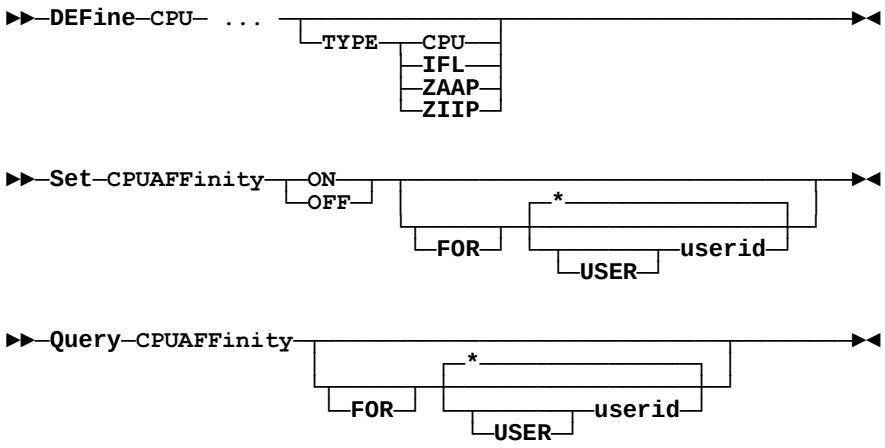

Guest Specialty Engine Support …

- **INDICATE LOAD shows processor type**
  - **PROC 0000-006% CP PROC 0001-003% CP**
    PROC 0002-003% ZAAP
- **INDICATE USER EXPANDED shows virtual processor type and affinity and resources used on real primary and real secondary processors**
  - **CPU 00: Ctime=15 01:49:23 Vtime=0 00:06:58 Ttime=0 00:07:33**
    Rdr=22514 Prt=465884 Pch=7088 IO=353978
    Type=CP CPUAffinity=ON
    VtimePrimary=0 00:06:58 TtimePrimary=0 00:07:33
    VtimeSecondary=0 00:00:00 TtimeSecondary=0 00:00:00
    CPU 01: Ctime=0 00:00:07 Vtime=0 00:00:00 Ttime=0 00:00:00
    Rdr=0 Prt=0 Pch=0 IO=0
    Type=ZIIP CPUAffinity=SUPP
    VtimePrimary=0 00:00:00 TtimePrimary=0 00:00:00
    VtimeSecondary=0 00:00:00 TtimeSecondary=0 00:00:00

Guest Specialty Engine Support …

- **INDICATE USER shows virtual processor type and affinity**
  - **USERID=ROMNEY MACH=ESA STOR=128M VIRT=V XSTORE=NONE**
    IPLSYS=CMS DEVNUM=00022
    PAGES: RES=00001646 WS=00001626 LOCKEDREAL=00000000 RESVD=00000000
    NPREF=00000020 PREF=00000000 READS=00008061 WRITES=00011948
    XSTORE=000284 READS=001585 WRITES=003445 MIGRATES=001249
    CPU 00: CTIME=68:17 VTIME=002:16 TTIME=003:11 IO=001498
    RDR=001474 PRT=000000 PCH=000000 TYPE=CP CPUAFFIN=ON
    CPU 01: CTIME=00:00 VTIME=000:00 TTIME=000:00 IO=000000
    RDR=000000 PRT=000000 PCH=000000 TYPE=ZIIP CPUAFFIN=SUPP

Guest Specialty Engine Support …

- **QUERY PROCESSORS shows real processor type**
  - **PROCESSOR 00 MASTER CP**
    PROCESSOR 01 ALTERNATE CP
    PROCESSOR 02 ALTERNATE ZAAP
- **QUERY VIRTUAL CPUS shows virtual processor type and affinity**
  - **CPU 00 ID FF319B9E20948000 (BASE) CP CPUAFF ON**
    CPU 01 ID FF319B9E20948000 STOPPED ZIIP CPUAFF SUPP
- **DEDICATE and UNDEDICATE only allowed for virtual processor that matches primary real type**
- **Cannot VARY OFF last primary real processor**
- **DEFINE CRYPTO only allowed for virtual CP**

Guest Specialty Engine Support …

- **Accounting records**
  - **Type 01 record includes virtual and real processor type code**
    - **May be multiple records per guest virtual CPU**
      - **CPU time used on primary real processor**
      - **CPU time used on secondary real processor**
    - **Accounting performed when DEFINE CPU command changes type**
  - **Type 0D record includes secondary CPU capability value**

Guest Specialty Engine Support …

- **Monitor records**
  - **Real processor type added to Monitor, System, Scheduler, Storage, and Processor domain processor-specific records**
  - **Virtual processor type and secondary processor CPU times added to User domain Logoff, Activity, Interaction, and Transaction End records**
  - **Virtual processor type added to User domain DEFINE CPU and DETACH CPU records**
  - **New records**
    - **MRSYTSPT – Scheduler Activity by Processor Type (Sample)**
    - **MRSCLSCA – SET CPUAFFINITY Changes (Event)**

Virtualization ...

- **Guest LAN and Virtual Switch Ease-of-Use**
  - **Authorized VLAN identifier set and promiscuous mode authorizations changed immediately**
  - **Native VLAN identifier may be configured for untagged traffic**
  - **IP address timeout support added for Guest LAN and IP virtual switch**
  - **New Virtual Network monitor domain**
    - **MRVNDSES - Virtual NIC Session Activity (Sample)**
    - **MRVNDLSU - Virtual Network Guest Link State - Link Up (Event)**
    - **MRVNDLSD - Virtual Network Guest Link State - Link Down (Event)**

Guest LAN/VSWITCH Ease-of-Use

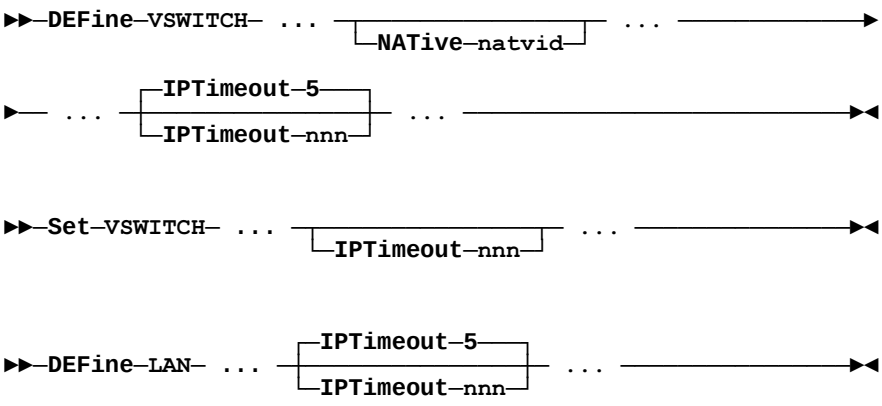

Guest LAN/VSWITCH Ease-of-Use ...

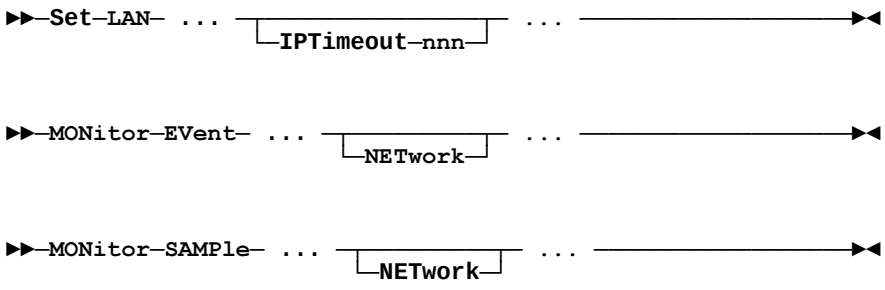

Virtualization ...

- **Guest MIDAW Support**
  - **Modified IDAW – Indirect Data Address Word**
    - **Can designate multiple non-contiguous data areas in storage**
    - **Each MIDAW includes data address, count field, and flags**
  - **Allows guest to exploit System z9 and later hardware capability**
    - **More flexible and performance-efficient than CCW data chaining**
      - **Each IDAW-referenced data area must end at 2K or 4K boundary**
      - **(M)IDAW fetching more efficient than CCW fetching**
  - **Diagnose X'210' (Retrieve Device Information) indicates if device is MIDAW-capable**

Virtualization …

- **Guest ASCII Console Support**
  - **HMC supports integrated ASCII console**
    - **Behaves like VT220**
      - **Enables cursor addressing**
      - **Provides familiar look and feel for Linux full-screen applications (e.g., vi, emacs)**
  - **Can be dedicated to a (Linux or z/VM) guest**
    - **Because Linux ignores errors (e.g., if console DETACHed), can switch from guest to guest at will**
  - **Provides recovery mechanism when normal network access not available**

Guest ASCII Console Support

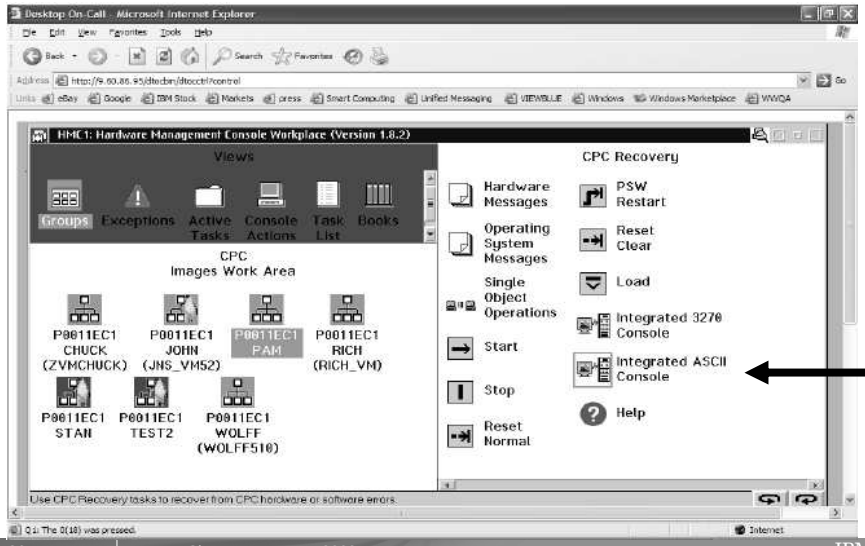

Guest ASCII Console Support …

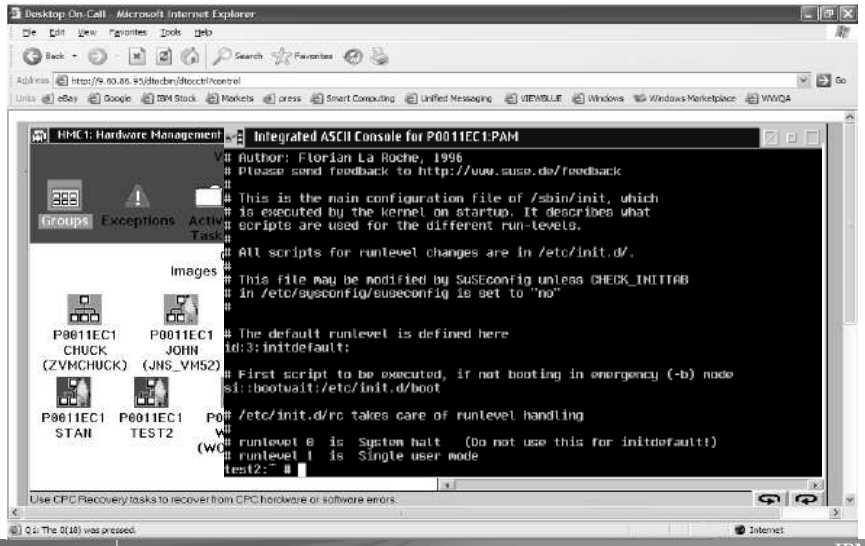

Guest ASCII Console Support …

- **New system object (not a conventional I/O device) – SYSASCII**
  - **Reported “free” or “attached to” guest**
  - **Reported “active” if ASCII console session open on HMC**
  - **Otherwise “inactive”**
- **Operated on by**
  - **ATTACH**
  - **DETACH**
  - **QUERY**
  - **QUERY VIRTUAL**

Guest ASCII Console Support …

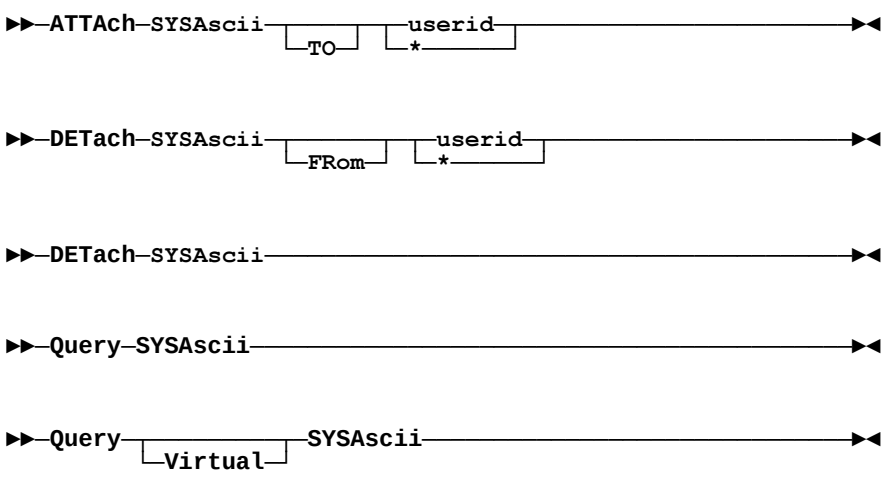

Guest ASCII Console Support …

- **A Linux guest must be configured to use the ASCII console, as described in "Linux on zSeries Device Drivers, Features, and Commands“**
  - **Device `ttyS1` is the full-screen mode ASCII console device driver**
- **Add a line to /etc/inittab**
  - **`<id>: 2345:respawn:/sbin/agetty -L 9600 ttyS1 linux`**
- **Guest must be booted with kernel parameter**
  `console=ttyS1`
  or `console=ttyS0 console=ttyS1`
  to activate the full-screen console device driver

Virtualization …

- **Enhanced SCSI Support**
  - **Dynamic preferred path discovery**
    - **DS6000**
    - **PREFERRED option no longer supported on SET EDEVICE**
  - **Fast format**
    - **ESS and DS8000**
    - **New Diagnose X'A4' (Synchronous Block I/O) Format function**
      - **Support reported by Diagnose X'210' (Retrieve Device Information)**
  - **Duplicate LUN checking**
  - **Point-to-point support**
    - **Eliminates need for switched Fibre Channel fabric**
  - **Additional device information in QUERY EDEVICE DETAILS response**
    - **Device vendor, product identification, revision level, cache status (if applicable), connection type (switched or Point-to-point) for each path**

Security

- **Passphrase support**
- **Tape encryption support**

Security …

- **Passphrase Support**
  - **Allow ESM to support longer and more complex passwords (password phrases or passphrases)**
    - **May be up to 200 characters long**
    - **May include any hexadecimal character, including blank**
    - **May require enclosure in single quotation marks**
  - **RACF/VM Feature supports up to 100-character passphrases**
  - **z/VM logo extends password field to end of line (width – 15)**
  - **No line-editing performed on user identifier (entered at system logon screen) or password**
  - **z/VM User Directory does not support passphrases**
  - **AUTOLOG, XAUTOLOG, LINK, APPC do not support passphrases**

PassPhrase Support

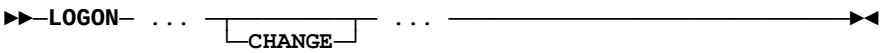

PassPhrase Support …

- New Diagnose X'88' subcode 8
  - Similar to subcode 0 but
    - Supports passphrases
    - Handles password case
    - Invokes ESM if present
    - Optionally asks ESM about agent's LOGON BY authority for target
    - Validates uppercase password against User Directory if no ESM
- New Diagnose X'88' subcode X'FF'
  - Determine ability to use other Diagnose X'88' functions

Security …

- Tape Encryption Support
  - IBM System Storage TS1120 Tape Drive (3592 E05)
  - Support KEY option for various commands and utilities
    - ATTACH command
    - SET RDEVICE command
    - DASD Dump/Restore (DDR) utility
  - Enable encryption automatically for encryption-unaware guests using the default key or a designated key label
  - Enhance related z/VM tape support facilities
    - SPXTAPE command
    - QUERY TAPES DETAILS
    - QUERY VIRTUAL TAPES

Tape Encryption Support

- Add key selection capability
  - ATTACH command
  - SET RDEVICE command
  - SET KEYALIAS command
  - QUERY KEYALIAS command
- Allow encryption-aware guests to exploit hardware facilities
  - Can use in-band key manager
- z/VM support requires an out-of-band key manager
  - IP attachment
- Delivered via PTFs for APAR VM64063 on z/VM 5.1 and 5.2
- ATL support delivered via PTF for DFSMS/VM FL221 APAR VM64062
  - Tape encryption support for z/VSE guests
- Diagnose X'210' (Retrieve Device Information) identifies 3592 E05 devices

Tape Encryption Support …

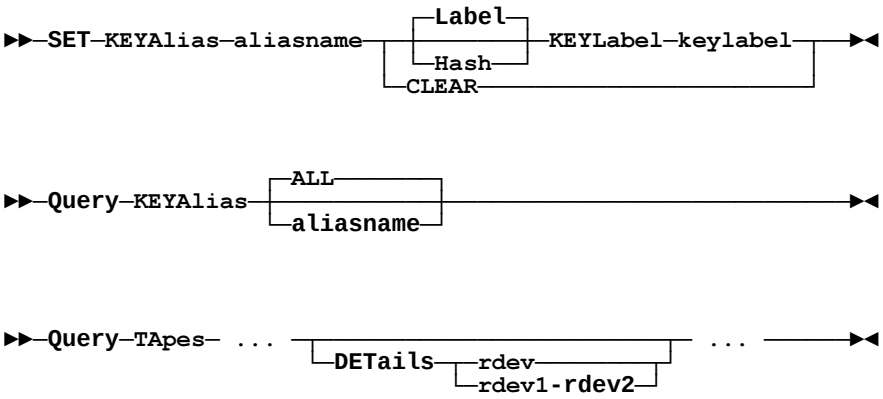

Tape Encryption Support ...

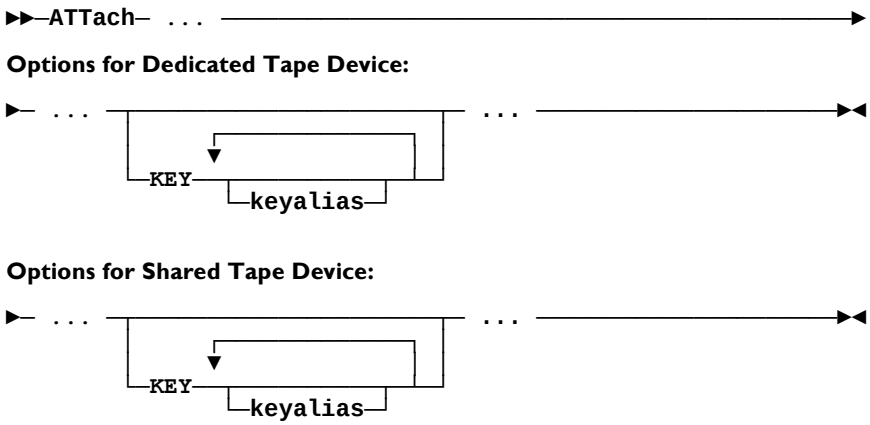

- Multi-user ATTACH allowed if KEY operands are identical

Tape Encryption Support ...

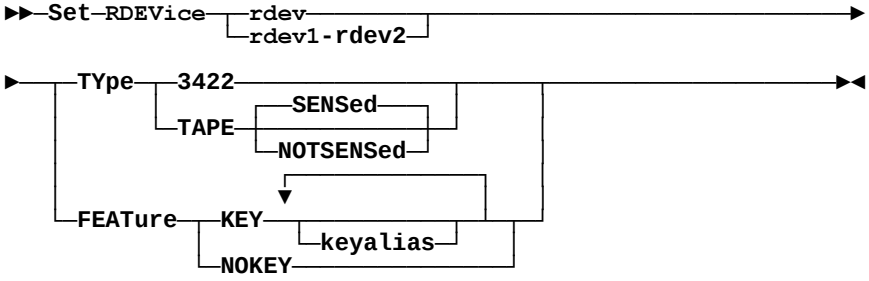

Tape Encryption Support …

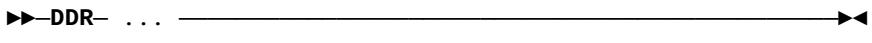

I/O Definition Control Statements:

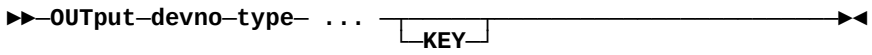

LABEL/HASH Control Statements for Encryption Key Labels:

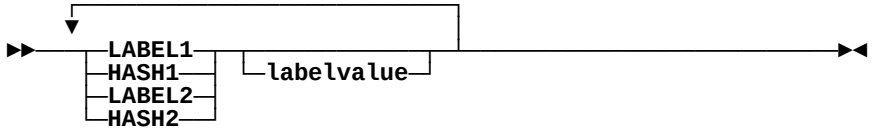

Tape Encryption Support …

- SPXTAPE DUMP honors ATTACH or SET RDEVICE KEY settings and enables for encryption as required
- QUERY TAPES DETAILS reports encryption-capable drives and displays active (set by ATTACH) and inactive (set by SET RDEVICE) key labels
- QUERY VIRTUAL TAPES reports encryption-capable devices

Virtual Networking

- **Enhanced virtual network management**
- **Link aggregation and failover**

Virtual Networking …

- **Enhanced Virtual Network Management**
  - **Support for SNMP agent virtual machine**
    - **MACID on SET VSWITCH**
      - **Concatenated to system MACPREFIX to define virtual switch MAC address**
      - **Can be system-assigned**
    - **Enhanced QUERY VSWITCH response shows**
      - **MACID associated with VSWITCH**
      - **Management ID (TCP/IP stack userid) and IP address associated with SNMP agent**
      - **Port numbers and associated interface indices**
    - **Enhanced QUERY VIRTUAL NIC response**
      - **Shows port number and associated interface index**

Enhanced Virtual Network Management

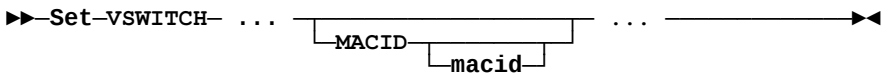

Enhanced Virtual Network Management …

- **New Diagnose X'26C' (Access Certain System Information) subcodes**
  - **X'00000008' - Return virtual LAN system information**
  - **X'0000000C' - Return controller list**
  - **X'00000010' - Return controller information**
  - **X'00000014' - Return guest LAN list**
  - **X'00000018' - Return guest LAN information**
  - **X'0000001C' - Return virtual switch list**
  - **X'00000020' - Return virtual switch information**
  - **X'00000024' - Return virtual port or NIC information**

Enhanced Virtual Network Management …

- **Changed Monitor records**
  - **Management user ID and switch IP and MAC addresses added to**
    - **MRIODVSW - Virtual Switch Activity (Sample)**
    - **MRIODVSF - Virtual Switch Failure (Event)**
    - **MRIODVSR - Virtual Switch Recovery (event)**

Virtual Networking …

- **Link aggregation and failover**
  - **Requires System z9 or later OSA Express-2**
  - **Allow up to eight OSAs to be associated with a virtual switch**
  - **Up to 128 named link aggregation groups supported**
  - **Links aggregated for data transmission**
    - **Optional time-driven balancing of conversations across links in group (30-second granularity)**
    - **Increased bandwidth**
    - **Improved recoverability**
  - **DEFINE VSWITCH and SET VSWITCH allow link aggregation group specification**
  - **QUERY CONTROLLER reports "LINKAGG"**
  - **QUERY PORT displays link aggregation group and device information**
  - **QUERY VSWITCH displays link aggregation group name**
  - **SET PORT defines link aggregation group**

Link Aggregation

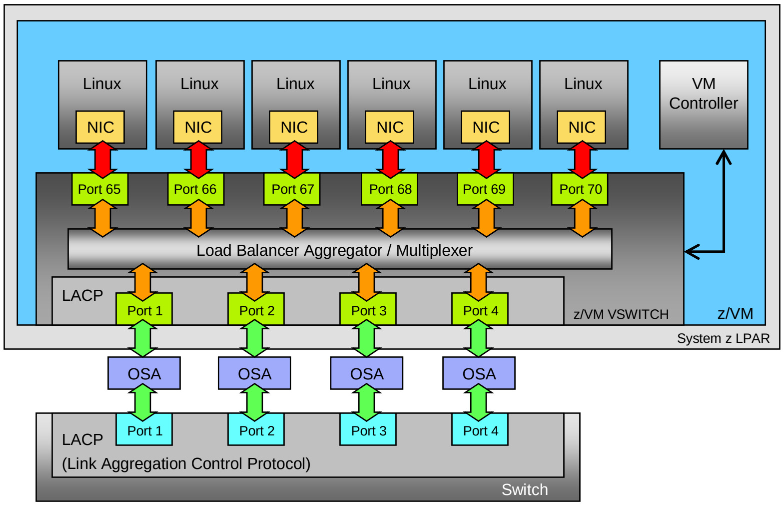

Link Aggregation …

```
►►─DEFine─VSWITCH─switchname─...──────────────────────────────────►◄
```

Ethernet Options:

```
   ┌─NOGroup──────────┐
├──┤                  ├──────────────────────────────────────────────┤
   └─GROup─groupname──┘

                           ┌─ALL─ACTive──────────────┐
►►─Query─PORT──┬─GROup─────┼─────────────────────────┼─┬─┬──────────┬─►◄
               │           │      ┌─ACTive───┐       │ │ └─DETails──┘
               │           ├─ALL──┤          ├───────┤ │
               │           │      └─INActive─┘       │ │
               │           └─groupname───────────────┘ │
               └─RDEV─rdev─────────────────────────────┘
```

Link Aggregation ...

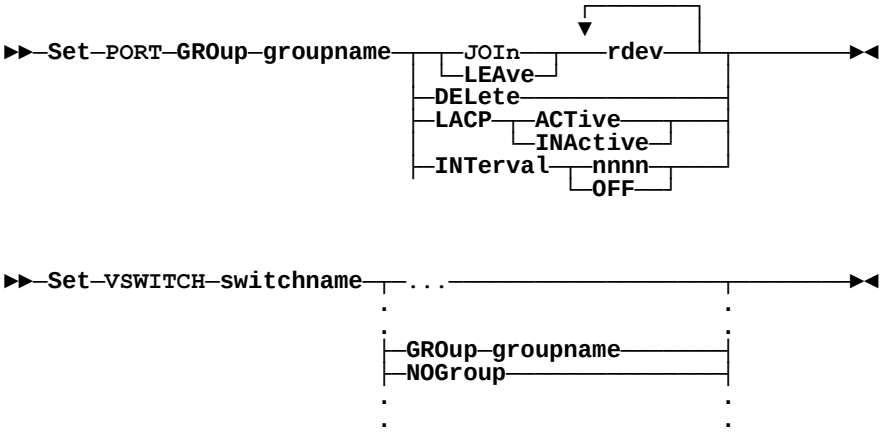

Systems Management

- **Enhanced guest configuration**
- **Asynchronous CP command responses**

Systems Management …

- **Enhanced Guest Configuration**
  - **User Directory COMMAND statement**
  - **Allows (almost) <u>any</u> CP command to be executed for a guest**
    - **After LOGON complete, immediately before IPL**
    - **Any privilege class**
    - **Subject to ESM auditing**
    - **Multiple statements allowed**
  - **Reduces need for future directory control statement enhancement**
    - **E.g., Specialty engines defined via COMMAND DEFINE CPU rather than by extensions to CPU statement**

Enhanced Guest Configuration

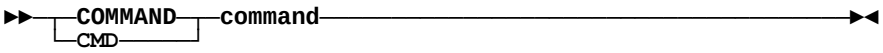


```
USER U1    U1PW 32M 32G G
IPL 190  PARM AUTOCR
COMMAND VARY ON 1234
COMMAND ATTACH 1234 TO &USERID AS 4567
CONSOLE 009 3215 T MAINT
SPOOL 00C 2540 READER A
SPOOL 00D 2540 PUNCH A
SPOOL 00E 1403 A
       .
       .
```

Systems Management ...

- **Asynchronous CP Command Responses**
  - **New FOR command allows one user to execute CP commands on behalf of another**
    - **Issuer must have SECUSER authority or Class C privileges**
    - **Target must have appropriate authorization for subject command**
  - **Responses come to issuer**
    - **`userid : [token : ]text`**
  - **No indication to target (except for side-effects)**
    - **E.g., FOR OPERATOR LOGOFF**
  - **Not recommended for use with asynchronous commands (e.g., SPXTAPE)**
  - **New *ASYNCMD IUCV System Service allows programmatic use**

Asynchronous CP Command Responses

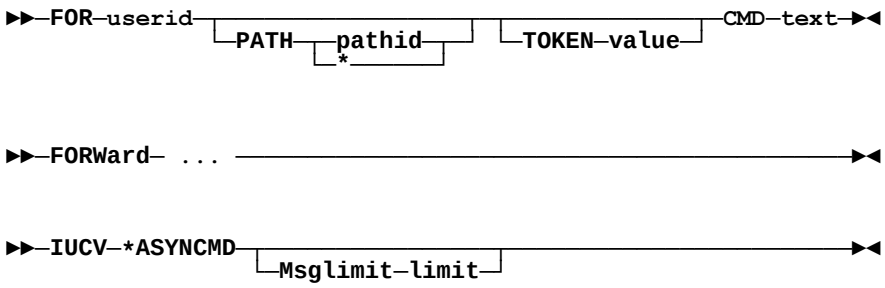

Asynchronous CP Command Responses …

***ASYNCMD Command Response Record:**

1-8	9-24	25	26-28	29-32	33	34-n
User ID	Token	0	Component ID	Message Number	Severity	Text

***ASYNCMD End-of-Command Record:**

1-8	9-24	25	26-29	30-33
User ID	Token	1	Return Code	Messages Discarded

Miscellaneous

- **Shutdown message time stamps**
- **SYSEVENT Query Virtual Server**
- **TRSOURCE for LDEVs**
- **QUERY IUCV**
- **DS8000 Dynamic Volume Expansion toleration**

Miscellaneous …

- **Shutdown Message Time Stamps**
  - **Shutdown progress (HCPWRP963I)**
  - **Dump information (HCPDMP9252I)**
  - **Dump progress (HCPDMP9260I)**
  - **Dump complete (HCPDMP9261I)**
  - **Shutdown complete (HCPWRP961W)**
  - **System termination complete (HCPWRP9277I)**

Miscellaneous …

- **SYSEVENT Query Virtual Server**
  - **Diagnose X'2E0'**
  - **Provides compatible equivalent of MVS SYSEVENT QVS service**
    - **Returns capacity in MSUs of CEC, LPAR, and virtual machine**
    - **Useful for sub-capacity pricing**
    - **Requested by ISV**
  - **APAR VM64122 for z/VM 5.1 and 5.2**

Miscellaneous …

- **TRSOURCE for LDEVs**
  - **Option of TRSOURCE TYPE IO**
    - **LDEV must exist when TRSOURCE issued**
  - **Produces TRF file showing traffic between CP and LDEV host**
  - **QUERY TRSOURCE shows “LDEV” and associated address(es)**
  - **TRACERED output reports “LDEV” and associated address**

TRSOURCE For LDEVs

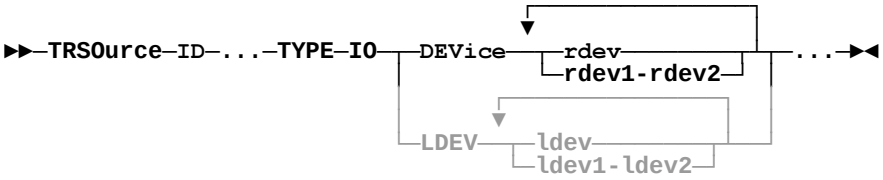

Miscellaneous …

- **QUERY IUCV**
  - **Allows any user to display information about own IUCV connections**
  - **Allows Class B user to display information about IUCV connections**
    - **For a specific user**
    - **For a specific IUCV System Service**

QUERY IUCV

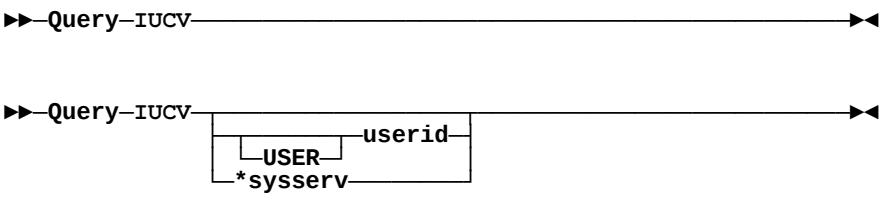

Miscellaneous …

- **DS8000 Dynamic Volume Expansion toleration**
  - **Increase logical volume size while online to a host system**
    - **3390-3 expanded beyond 3,339 cylinders becomes 3390-9**
    - **3390-9 can be expanded up to 65,520 cylinders**
    - **SCSI LUN can be expanded up to 2,147,483,640 blocks (2^{31} - 8)**
  - **QUERY DASD DETAILS reports pending expansion**
  - **z/VM cannot use expanded space until device varied offline/online**
  - **APARs VM64305, VM64354**
    - **z/VM 5.2 PTFs UM32230, UM32240**
    - **z/VM 5.3 PTFs UM32231, UM32241**

Statements of Direction

- **Common Criteria Certification**
- **3480 Distribution Medium**
- **Additional support for managing z/VM systems**
- **z/VM LPAR enhancements**

Statements of Direction …

- **Common Criteria Certification**

  IBM intends to evaluate z/VM V5.3 with the RACF Security Server optional feature for conformance to the Controlled Access Protection Profile (CAPP) and Labeled Security Protection Profile (LSPP) of the Common Criteria standard for IT security, ISO/IEC 15408, at Evaluation Assurance Level 4 (EAL4).

  This new SOD represents a modification to IBM's previously expressed Statement of Direction of July 27, 2005, which stated IBM's intent "to evaluate z/VM V5.2 with the RACF for z/VM optional feature for conformance to the Controlled Access Protection Profile (CAPP) and Labeled Security Protection Profile (LSPP) of the Common Criteria standard for IT security, ISO/IEC 15408, at Evaluation Assurance Level 4 (EAL4).“ Based on additional assessment of requirements, IBM no longer intends to evaluate z/VM V5.2.

Statements of Direction …

- **3480 Distribution Medium**

  IBM intends to withdraw 3480 tape as a distribution medium in a future z/VM release. z/VM is planned to continue distribution on 3590 tape and on DVD, and to be available for electronic delivery from ShopzSeries *(Internet delivery made available January 25, 2008)*

- **Additional support for managing z/VM systems**

  IBM intends to further enhance z/VM in a future release to exploit the new Hardware Management Console (HMC) interface that allows the installation of Linux on System z into a z/VM virtual machine. Additionally, future support is planned for z/VM and the HMC to provide z/VM hypervisor-configuration tasks.

Statements of Direction …

- **z/VM LPAR enhancements:**

  IBM intends to further enhance z/VM in a future release to exploit the System z10 EC support for a new logical partition (LPAR) mode "z/VM," exclusively for running z/VM. This new LPAR mode allows z/VM to utilize a wider variety of specialty processors in a single LPAR. For instance, in a z/VM mode LPAR, z/VM can manage Linux on System z guests running on IFL processors while also managing z/VSE and z/OS on central processors (CPs), and to offload z/OS system software overhead, such as DB2 workloads, on IBM System z9 and IBM System z10 Integrated Information Processors (zIIPs) and IBM System z9 and IBM System z10 Application Assist Processors (zAAPs).