z/VM 6.3 Security "News and How To's"

Brian W. Hugenbruch, CISSP
z/VM Security Design and Development
bwhugen@us.ibm.com
@Bwhugen

Trademarks

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

Not all common law marks used by IBM are listed on this page. Failure of a mark to appear does not mean that IBM does not use the mark nor does it mean that the product is not actively marketed or is not significant within its relevant market.

Those trademarks followed by ® are registered trademarks of IBM in the United States; all others are trademarks or common law marks of IBM in the United States.

For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml:

*, IBM Systems, IBM System z10®, IBM System Storage® , IBM System Storage DS®, IBM BladeCenter®, IBM System z®, IBM System p®, IBM System i®, IBM System x®, IBM IntelliStation®, IBM Power Architecture®, IBM SureOne®, IBM Power Systems™, POWER®, POWER6®, POWER7®, POWER8®, Power ®, IBM z/OS®, IBM AIX®, IBM i, IBM z/VSE®, IBM z/VM ®, IBM i5/OS®, IBM zEnterprise®, Smarter Planet™ ,Storwize®, XIV® , PureSystems™, PureFlex™, PureApplication™ , IBM Flex System™ , Smarter Storage

The following are trademarks or registered trademarks of other companies.

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.
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.
Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. 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.
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.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.
IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

* 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.

Disclaimer

The information contained in this document has not been submitted to any formal IBM test and is distributed on an "AS IS" basis without any warranty either express or implied. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

In this document, any references made to an IBM licensed program are not intended to state or imply that only IBM's licensed program may be used; any functionally equivalent program may be used instead.

Any performance data contained in this document was determined in a controlled environment and, therefore, the results which may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data for their specific environments.

It is possible that this material may contain reference to, or information about, IBM products (machines and programs), programming, or services that are not announced in your country. Such references or information must not be construed to mean that IBM intends to announce such IBM products, programming or services in your country.

Agenda

- **A Public Service Announcement** ***Updated***
- **z/VM Security Certifications** ***Updated***
- **z/VM 6.3 Security News and "How-To's"**
  - TLS 1.2 Support in the z/VM SSL Server ***Updated***
  - CP Changes and Virtual Networking Updates ***Updated***
  - CryptoExpress 4S Support
  - RACFVM Support for z/VM Single System Image Clusters
- **Discussion / Questions**

z/VM Security News:

A Public Service Announcement

"Is z/VM vulnerable to _____?"

"Is z/VM vulnerable to _____?"

- IBM System z Security policy prohibits the general disclosure of vulnerability analyses (negative or positive). In part this is to prevent any inadvertent or malicious exploitation of vulnerabilities in System z environments which have not yet been updated to current levels of service. To stay current, your company can register with the IBM System z Security Portal in order to receive up to date lists regarding APAR/PTF information and CVSS scoring for SEC/INT service as it becomes available. In addition, Security Notices will be published through this website in order to address high-profile security issues, notifications and possible warnings.

- Access to the portal can be obtained through the following website: http://www-03.ibm.com/systems/z/solutions/security_subintegrity.html

Common Vulnerability Scoring System (CVSS v2)

- z/VM provides a CVSS Score and Vector for Security-related z/VM APARs ("**ResourceLink**" information) for subscribed customers

- An open-standard metric for vulnerability measurement
  - http://www.first.org/cvss/cvss-guide.html
  - Not to be confused with a "threat rating system" or vulnerability catalogue

- IBM Internet Security Systems, similarly, includes CVSS base and temporal scores in its X-Force bulletins: http://www.iss.net/threats/ThreatList.php

Common Vulnerability Scoring System (CVSS v2)

- Comprised of three scores:
  - A **base metric** which measures complexity, levels of authentication, access vectors, and impacts to various aspects of security;
  - A **temporal metric** which measures the exploitability of the threat and availability of a fix; and
  - An **environmental metric** which determines a vulnerability's impact to a specific configuration, including the potential for collateral damage and percent of a business that might be under threat.

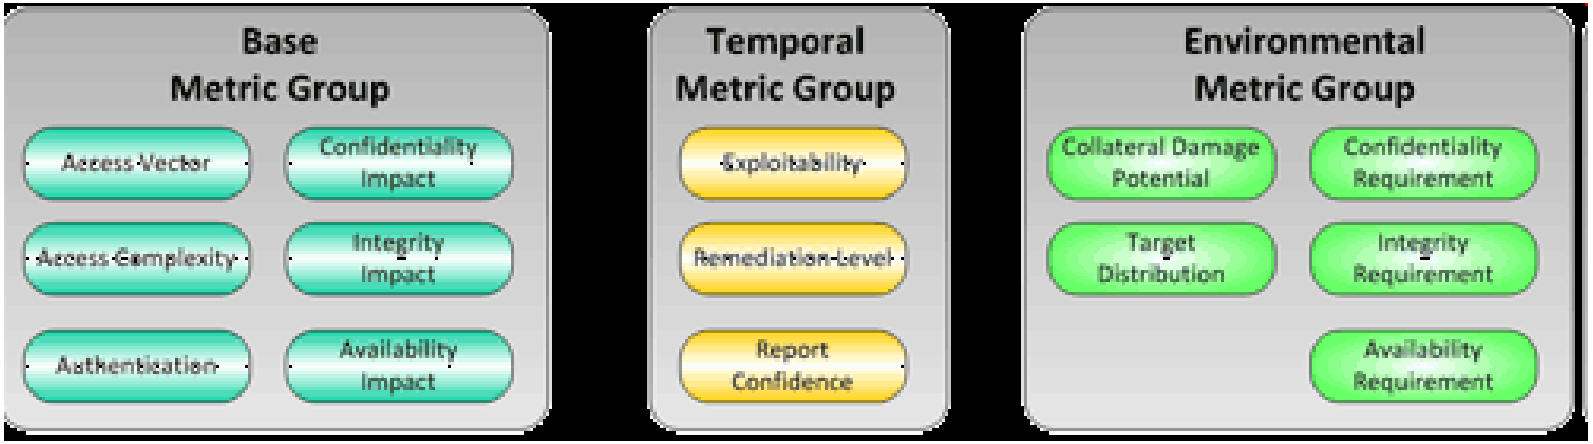

Example: an SSL "Man-in-the-Middle" Exploit

(Sample analysis. Does not represent a formal IBM analysis, or represent actual IBM service.)

Given the following vectors:

(AV:N/AC:M/Au:N/C:P/I:P/A:N/E:ND/RL:OF/RC:C)

Where:

AV:	N	-- access through wide network, not local traffic
AC:	M	-- Access requirements are medium. Complicated, but not esoteric.
Au:	N	-- No system authentication is required.
C:	P	-- There is a partial threat to information confidentiality. (Hacker may steal data.)
I:	P	-- There is a partial threat to data integrity. (Hacker may change, corrupt data.)
A:	N	-- The hacker can't actually bring down the system, though.
E:	ND	-- Exploitability isn't defined.
RL:	OF	-- There is an official fix available
RC:	C	-- Report Confidence is set to Confirmed

This exploit is rated as a 5.0 out of 10.0. (Base Score 5.8; Temporal Score 5.0.)

If the TLS Server is not defined on your system, Overall CVSS Score may be 0.

This score is for z/VM only; makes no statement about guest configuration!

Example: Susceptibility to a DDoS packet storm

(Sample analysis. Does not constitute a formal IBM analysis, or represent actual IBM service.)

Given the following vectors:

(AV:N/AC:L/Au:N/C:N/I:N/A:C/E:ND/RL:OF/RC:C)

Where:

AV:	N	-- access through wide network, not local traffic
AC:	L	-- Access requirements are low. This is a script kiddie running software.
Au:	N	-- No system authentication is required.
C:	N	-- There is no threat to information.
I:	N	-- There is no threat to data or system integrity.
A:	C	-- The hacker may knock systems offline or prevent access to services.
E:	ND	-- Exploitability isn't defined.
RL:	OF	-- There is an official fix available
RC:	C	-- Report Confidence is set to Confirmed

This exploit is rated as a 6.8 out of 10.0. (Base Score 7.8; Temporal Score 6.0.)

If your server requires 24/7 availability, the Overall CVSS Score may be 8.7.

z/VM Security News:

Certifications and Statements of Direction

z/VM Security Certification Discussion

z/VM Level	Common Criteria	FIPS 140-2
z/VM 6.3	In process: ***OSPP with Labeled Security and Virtualization at EAL 4+***	**FIPS 140-2 L1**
z/VM 6.2	Designed to comply to the technical and procedural standards involved in the certification, but not formally certified.	
z/VM 6.1	***OSPP with Labeled Security and Virtualization at EAL 4+*** • *BSI-DSZ-CC-0752*	**FIPS 140-2 L1**
z/VM 5.4	Designed to comply to the technical and procedural standards involved, but not formally certified.	**n/a**
z/VM 5.3	**CAPP/LSPP at EAL 4+**	FIPS "mode" for the Linux-hosted SSL Server, but it was not certified

TM: A Certification Mark of NIST, which does not imply product endorsement by NIST, the U.S. or Canadian Governments.

z/VM Security Certification: Statements of Direction

IBM intends to pursue an evaluation of the Federal Information Processing Standard (FIPS) 140-2 using National Institute of Standards and Technology's (NIST) Cryptographic Module Validation Program (CMVP) for the System SSL implementation utilized by z/VM V6.3.

- **April 30, 2014: FIPS 140-2 evaluation is now completed!**
  - http://csrc.nist.gov/groups/STM/cmvp/documents/140-1/1401val2014.htm#2139
  - Requires the PTF for APAR PI04999
- See http://www.vm.ibm.com/security/ for the latest in z/VM Security information.

z/VM Security Certification: Statements of Direction

IBM intends to evaluate z/VM V6.3 with the RACF Security Server feature, including labeled security, for conformance to the Operating System Protection Profile (OSPP) of the Common Criteria standard for IT security, ISO/IEC 15408, at Evaluation Assurance Level 4 (EAL4+).

- Common Criteria evaluation is listed as **In Certification** by BSI
  - https://www.bsi.bund.de/EN/Topics/Certification/incertification.html
  - BSI-DSZ-CC-0903

Other Security-Related Statements of Direction

Enhanced RACF password encryption algorithm: In the future, an enhanced RACF password encryption algorithm is planned. This support will be designed to provide improved cryptographic strength in RACF password algorithm processing. This will be intended to help protect RACF password data in the event that a copy of a RACF database becomes inadvertently accessible.

- Statement of Direction issued 24 February 2014
  (IBM zEnterprise System, z/OS, z/VM)

- z/OS has already met this Statement of Direction:
  New Function APARs OA43998 (SAF) / OA43999(RACF)

z/VM Security News:

TLS 1.2 Support for the z/VM SSL-TLS Server

Introducing the z/VM SSL-TLS Server

The z/VM SSL Server:

- A CMS-based server associated with the TCP/IP stack
- Secures point-to-point traffic **to the hypervisor**
- Telnet, FTP and SMTP provide “dynamic” SSL traffic
- Port-based “static” SSL

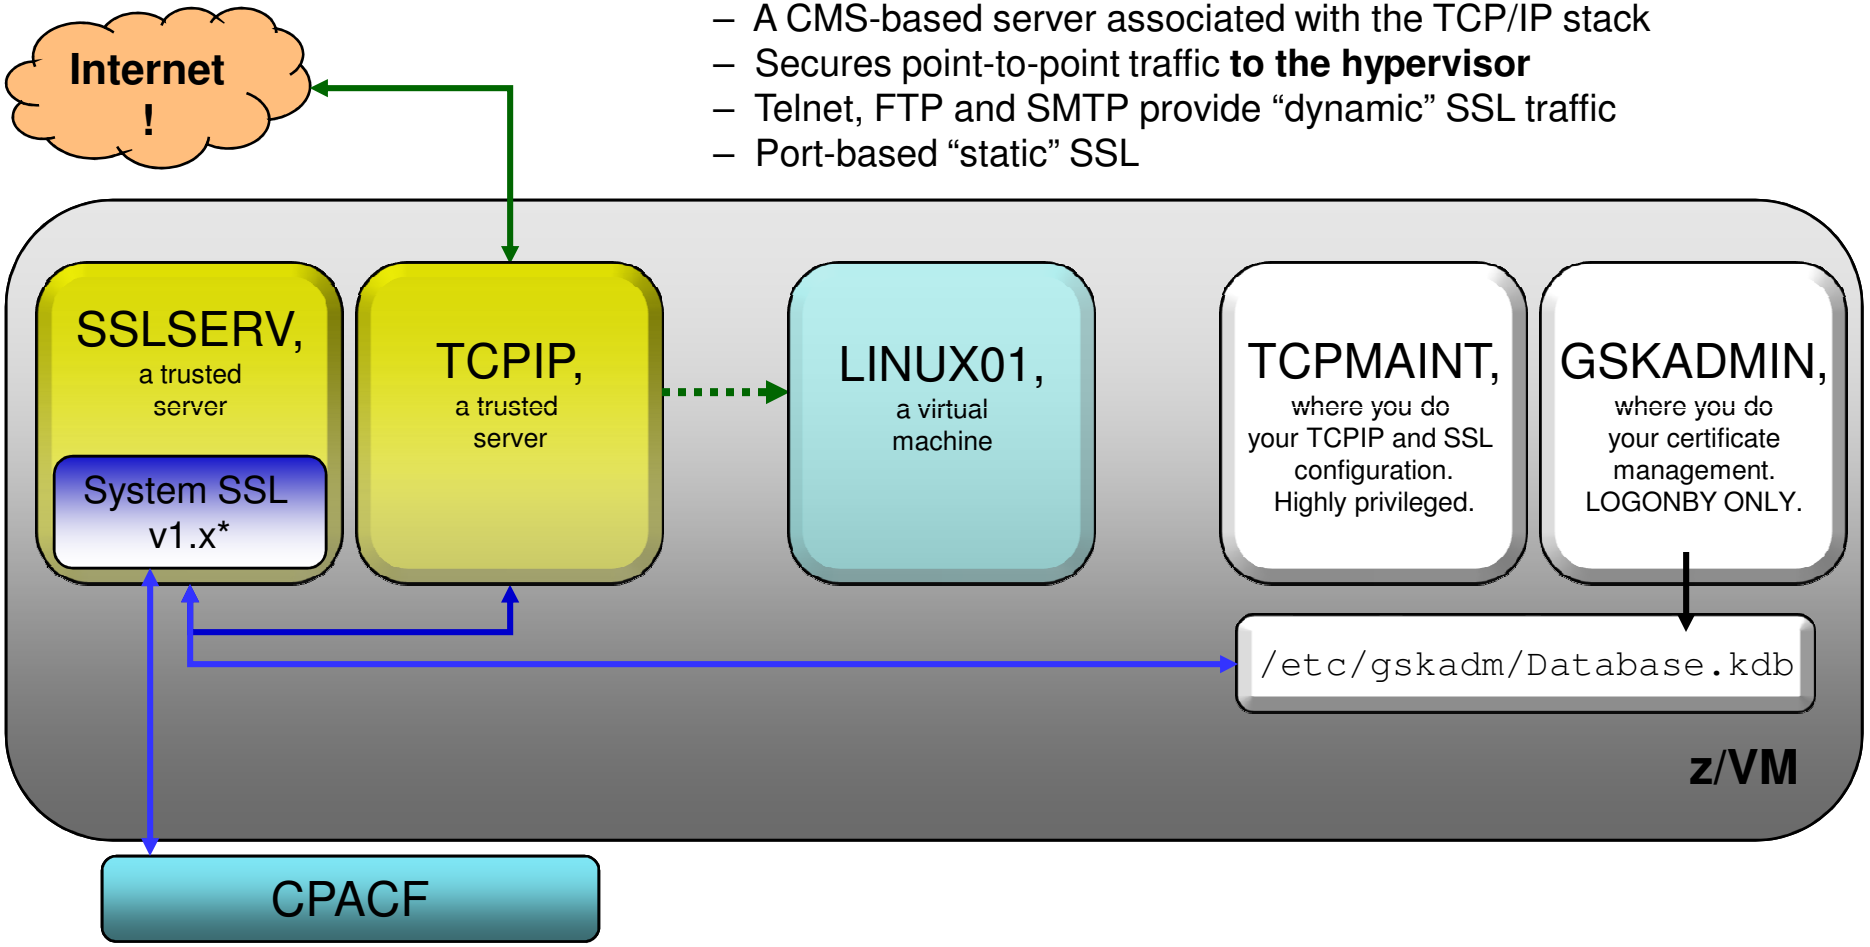

z/VM SSL-TLS Server News: Version 6 Release 3.0 and Service

- *System SSL Update*
  - Port of **z/OS V1.13** equivalency plus z/OS APAR OA39422
  - Enables TLS 1.2 functionality, SHA2 hashing, SHA2 certificates
  - HMAC-SHA256 integrity checking at start-up
  - Pre-Initialization FIPS Compliance (APAR PM95516)
  - FIPS Validated Level (APAR PI04999)

- *SSL Server Upgrades*
  - Client Certificate Validation for Telnet (**APAR PM52716** for z/VM 6.2)
  - IPv6 Support for Secure Telnet, FTP, and SMTP
  - **Support for TLS 1.2 connections**
  - New `'PROTOCOL'` keyword – enable/disable versions of SSL or TLS
  - New `'MODE'` keyword (APAR PM93363) – enable particular crypto compliance modes

z/VM SSL-TLS Server Options

- Specified either on VMSSL (command-line exec) or DTCPARMS
- Persists for the run-time for a server or server pool. Must be consistent for all members of a server pool
- Options:
  - **KEYFILE – BFS location of the certificate database**
  - **CACHELIFE – for secure connections, in hours, minutes, seconds**
  - **CACHECLEANUP – processed every n connections**
  - **MODE – sets a cryptographic compliance mode**
    - **MODE FIPS-140-2**
    - **MODE NIST-800-131A**
    - **FIPS – equivalent to MODE FIPS-140-2**
  - **PROTOCOL – enable or disable SSL/TLS levels.**
    - **TLS 1.0 enabled by default <= change from documented behavior**
    - **Available protocols change based on MODE**
  - **EXEMPT – disable particular cipher suites**
  - **GSKTRACE – enable System SSL tracing**
  - **TRACE/NOTRACE – enable SSL Server tracing**
    - **Can be dynamically manipulated via authorized commands**

z/VM SSL-TLS Server News – TLS 1.2 Support

High	Medium	Low	None
3DES_168_SHA	RC4_128_SHA	RC2_40_MD5	NULL
DH_DSS_3DES	RC4_128_MD5	RC4_40_MD5	NULL_SHA
DH_RSA_3DES	RSA_AES_128	DES_56_SHA	NULL_MD5
DHE_DSS_3DES	RSA_AES_128_SHA256	DH_DSS_DES	NULL_SHA256
DHE_RSA_3DES	DH_DSS_AES_128	DH_RSA_DES	
RSA_AES_256	DH_DSS_AES_128_SHA256	DHE_DSS_DES	
RSA_AES_256_SHA256	DH_RSA_AES_128	DHE_RSA_DES	
DH_DSS_AES_256	DH_RSA_AES_128_SHA256		
DH_DSS_AES_256_SHA256	DHE_DSS_AES_128		
DH_RSA_AES_256	DHE_DSS_AES_128_SHA256		
DH_RSA_AES_256_SHA256	DHE_RSA_AES_128		
DHE_DSS_AES_256	DHE_RSA_AES_128_SHA256		
DHE_DSS_AES_256_SHA256			
DHE_RSA_AES_256			
DHE_RSA_AES_256_SHA256			

Legend:
TLS 1.2 only
Not in TLS 1.2
Not in TLS 1.1 or 1.2

Note 1: Cipher suites can be exempted from processing based on either cipher name or by strength set, per the above (but not both).

Note 2: Exempting by strength automatically exempts a lower strength!

Note 3: Ciphers are negotiated on a per-handshake basis and are protocol-dependent.

z/VM SSL-TLS Server Updates – Mode Selection

- **MODE FIPS-140-2**
  - Replaces 'FIPS' keyword
  - Minimum Protocol of TLS 1.0
  - Export ciphers restricted
  - Minimum key exchange value of 1024
  - FIPS-compliant database required
    - Integrity checking (HMAC-SHA256): Digitally signs the crypto modules and database against tampering
    - Known Answer Tests – verify integrity after initialization
- z/VM has been FIPS-compliant since V6R1

- ***NEW*** **MODE NIST-800-131A**
  - Minimum Protocol of TLS 1.2
  - Minimum key exchange value of 2048
    - DSA certificate usage prohibited!
  - Minimum hash of SHA2
  - No certificate database requirements
    - Integrity checking only (HMAC-SHA256)
  - Supersedes FIPS-140-2 where applicable
- Requires **APAR PM93363** (z/VM 6.3 only)

When running in either mode, the cipher suites available adjust according to security settings …

“How To”: Select Protocols and Modes for the SSL-TLS Server

If we specify ...

[Default Settings]:
PROTOCOL +TLSV1_0

[New Protocols]:
PROTOCOL +TLSV1_1
PROTOCOL +TLSV1_2

MODE FIPS-140-2

MODE NIST-800-131A

EXEMPT MEDIUM

```
SSL00001  Enabled   TLSV1_2
SSL00001  Disabled  TLSV1_1 TLSV1_0 SSLV3 SSLV2

RSA_AES_256_SHA256 DH_RSA_AES_256_SHA256
DHE_RSA_AES_256_SHA256 RSA_AES_256 DH_RSA_AES_256
DHE_RSA_AES_256 DHE_RSA_3DES DH_RSA_3DES
```

- MODE FIPS-140-2 and MODE NIST-800-131A have additional restrictions:
  - Certificate key minimum of 1024 for FIPS, and FIPS-mode database required
  - Certificate key minimum of 2048 for NIST, and SHA-2 only
- MODE overrides specified PROTOCOL statements
  - FIPS requires a minimum protocol level of TLS 1.0
  - NIST requires a minimum protocol level of TLS 1.2
- Plan ahead if MODE support is a requirement for your configuration!

z/VM Security News:

CP Changes and Virtual Networking Updates

Security Policy Updates for z/VM 6.3

- IBM Supplied User Directory:
  Default passwords have been modified for new installs
  - Common string for easy search/replace
  - Remember to change your default passwords!

- User-Class Restructure (UCR) capability has been removed
  - per Statement of Direction, UCR and the OVERRIDE utility have been discontinued
  - `CP MODIFY COMMAND` and `CP MODIFY DIAGNOSE` available for decades
  - `CVTOVRID.XEDIT` macro available to translate UCR spool files to System Configuration statements

Virtual Networking Improvements

- Live Guest Relocation support for port-based virtual switches built on existing support:
  - Allow relocation of port-based interface
  - Prevent relocation of an interface that will be unable to establish proper network connectivity
  - Adjust the destination virtual switch configuration, when possible, by inheriting virtual switch authorization from the origin

- Virtual Switch recovery and stall prevention
  - New SET VSWITCH UPLINK SWITCHOVER command
  - Change from current device to one of the configured backup devices

- Virtual Ethernet Port Aggregator (VEPA) mode
  - Moves switching of traffic out of the virtual switch and into physical hardware
  - A more stringent mechanism for separating data flow in a virtual environment
  - …

VEB mode - VSwitch internal switching of Guest port traffic (default)

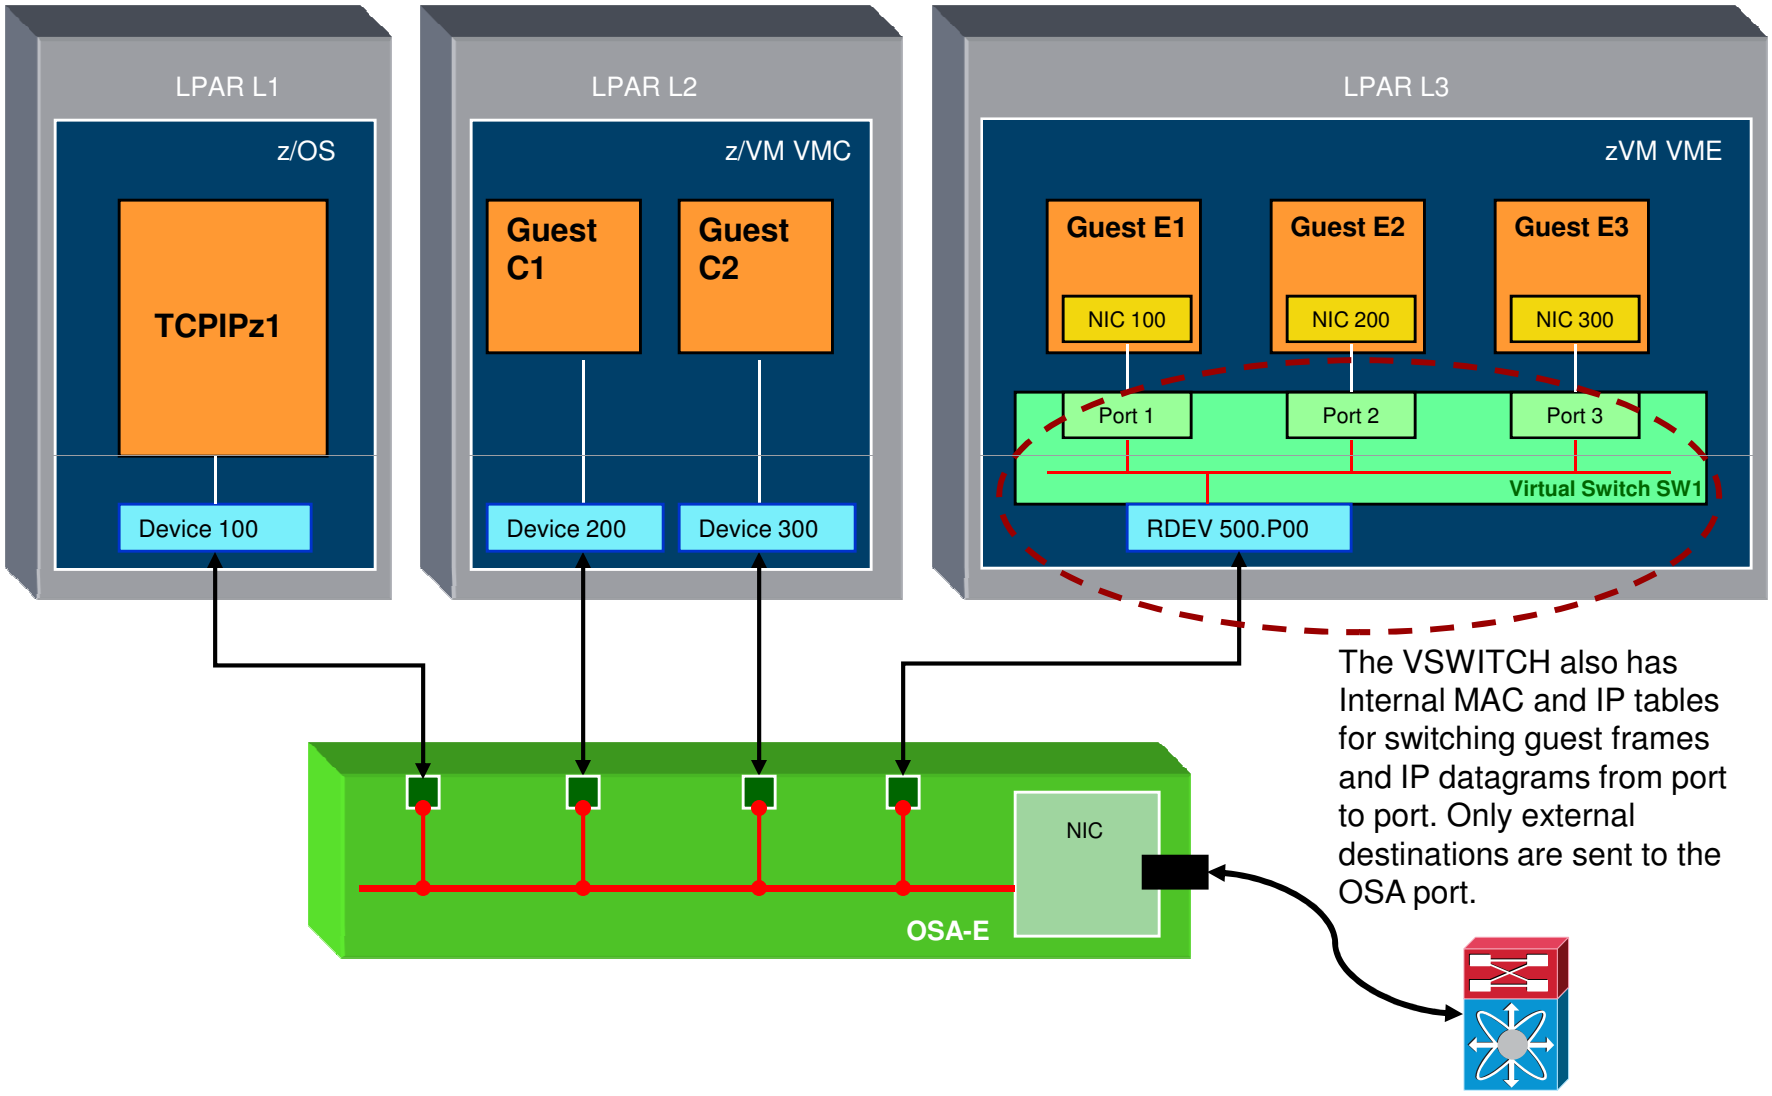

VSWITCH Port Isolation Mode

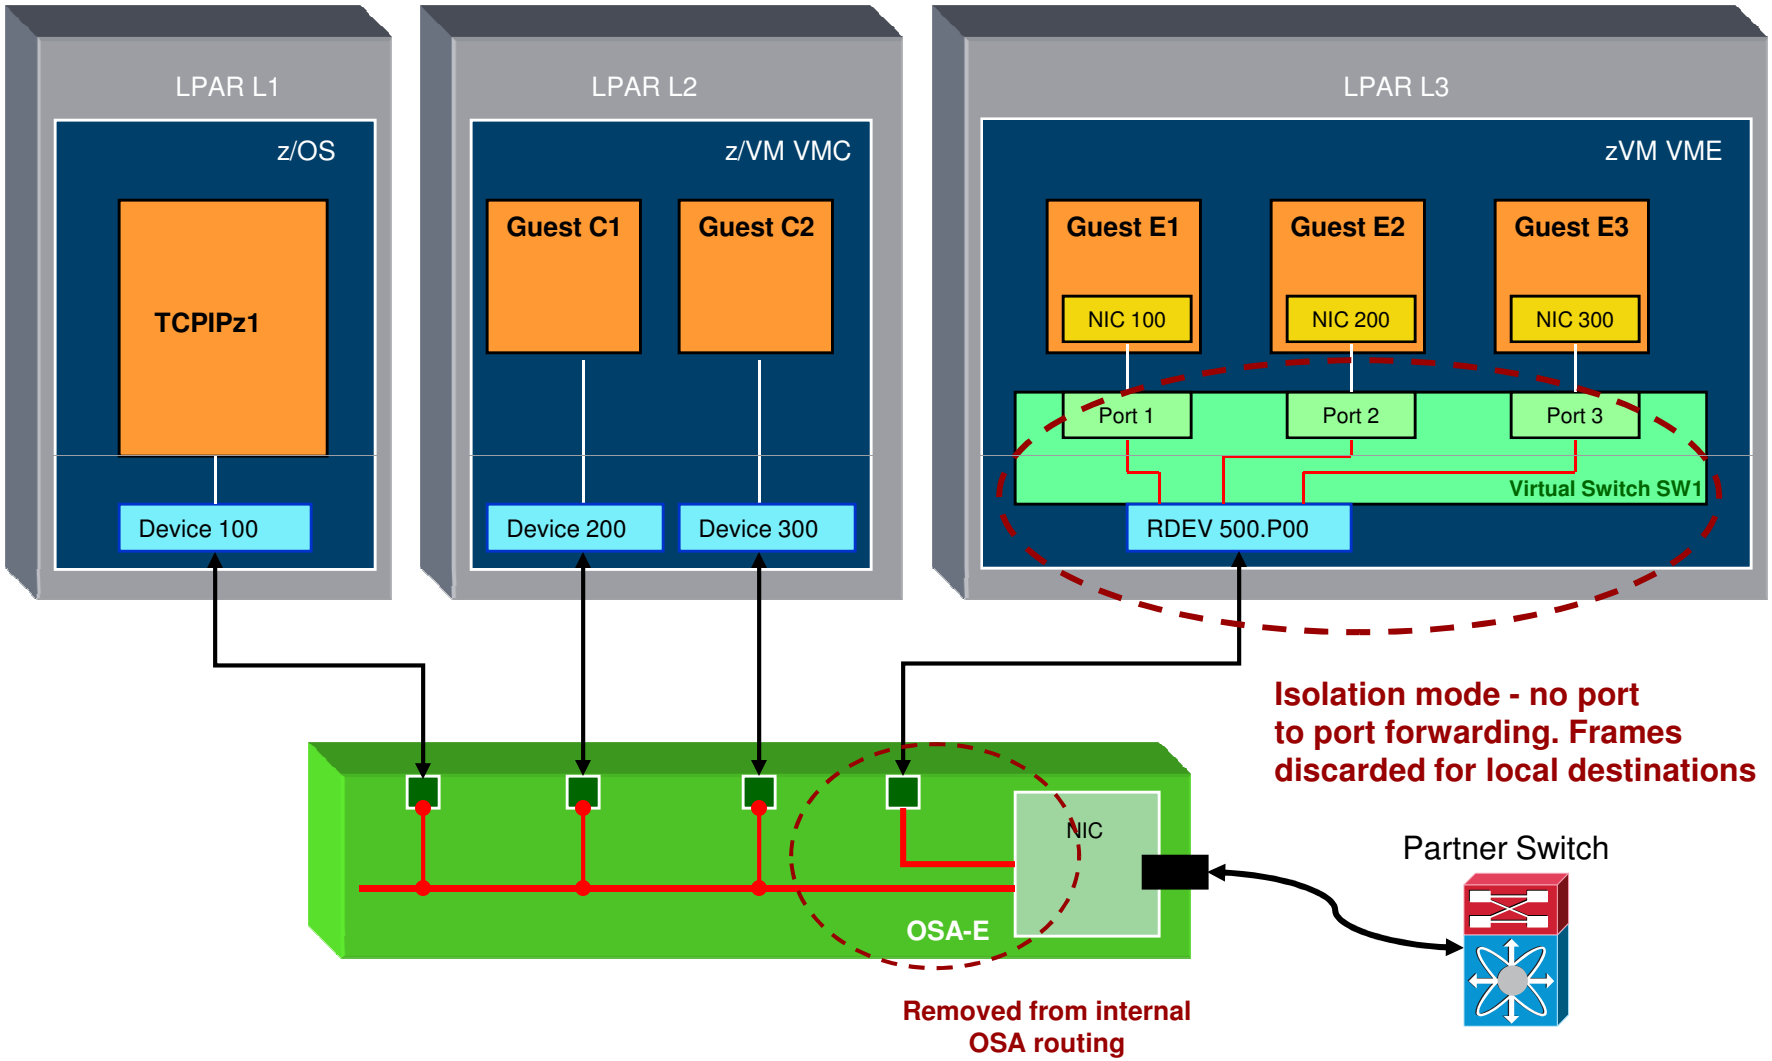

z/VM VSWITCH VEPA Mode

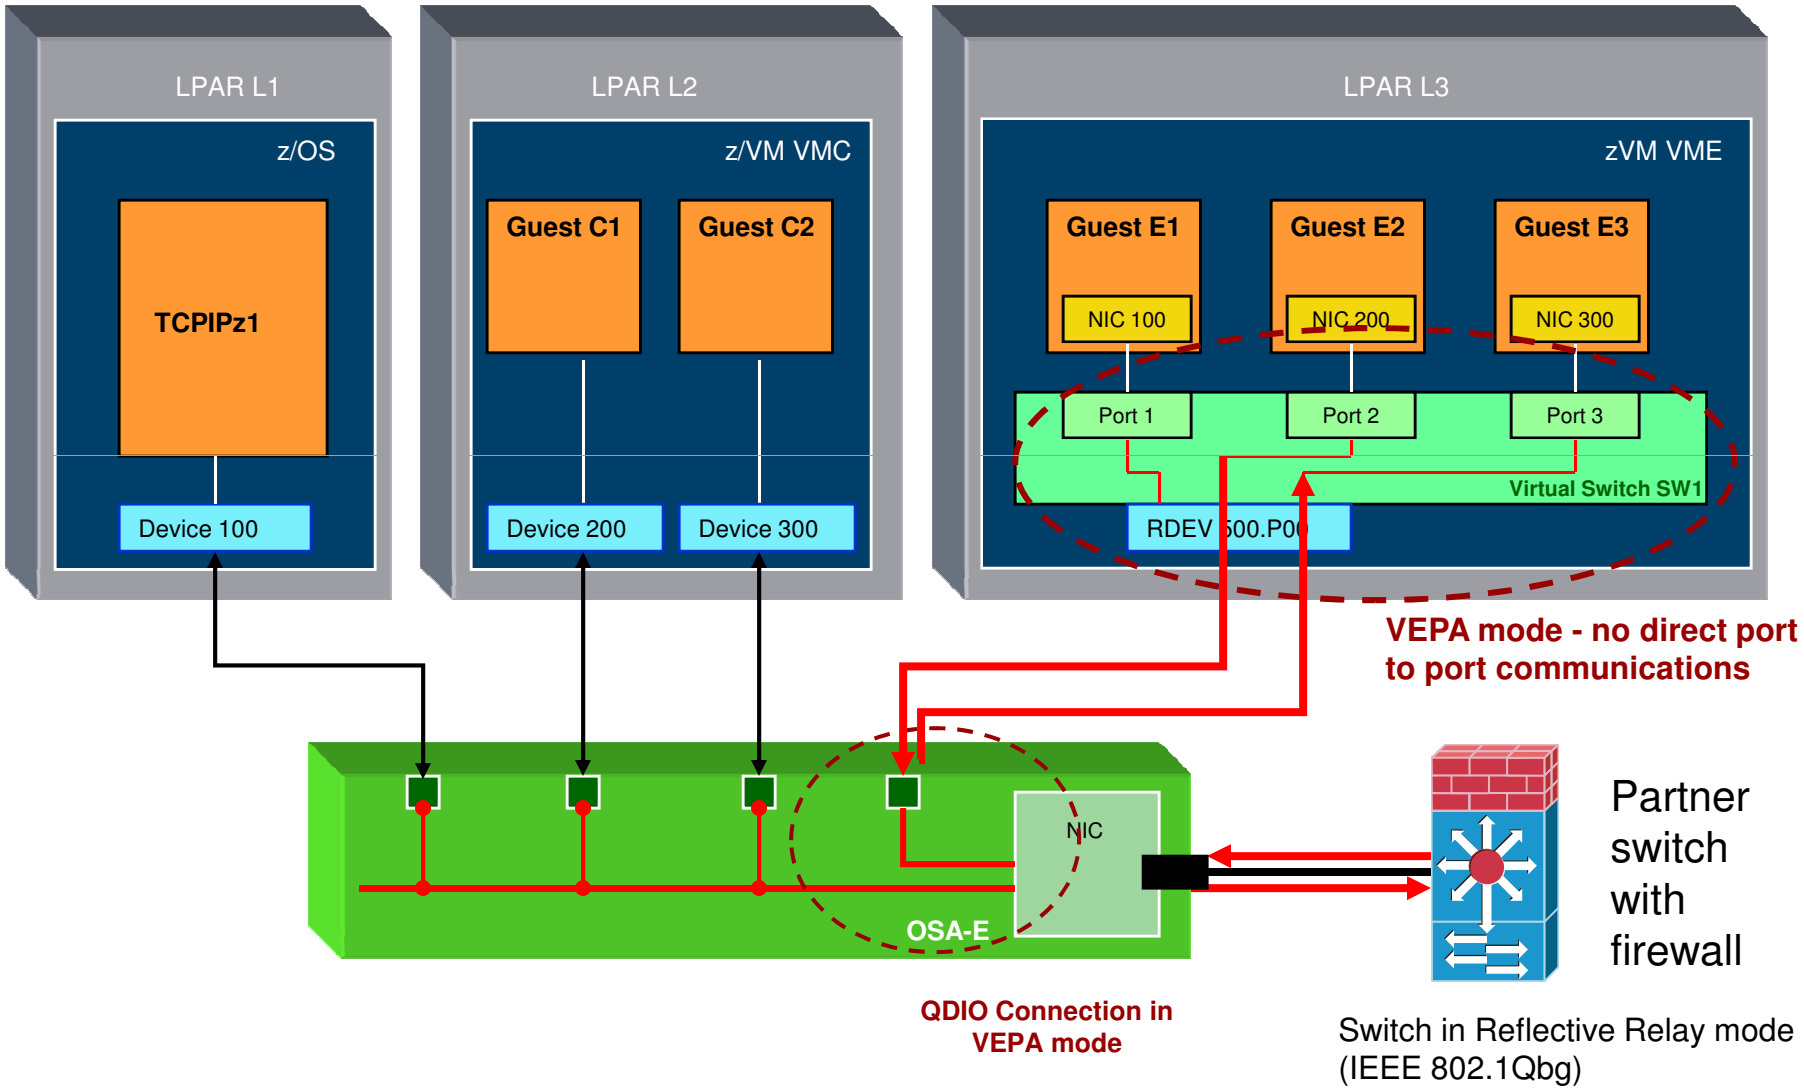

SET VSWITCH

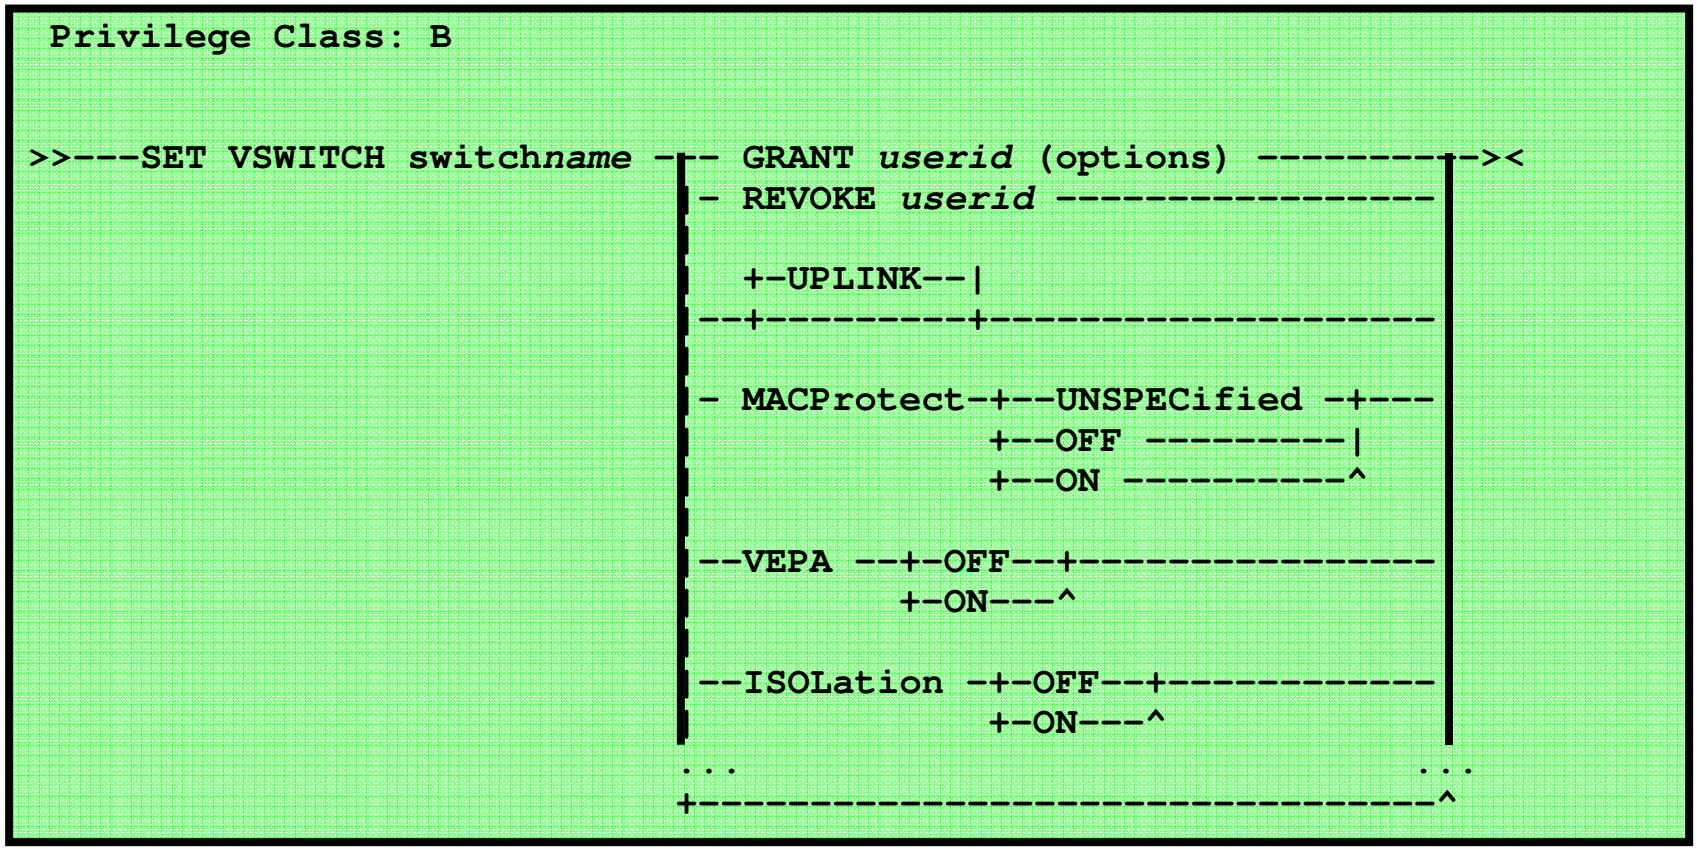

- By default, no virtual machine can access a VSWITCH
- GRANT provides capability to COUPLE.
- Note that VEB is the default setting; Port Isolation and/or VEPA would need to be enabled!

Some useful diagnostic commands for the VSWITCH

- **CP QUERY VMLAN**
  - to get global VM LAN information (e.g. limits)
  - to find out what service has been applied
- **CP QUERY LAN ACTIVE**
  - to find out which users are coupled
  - to find out which IP addresses are active
- **CP QUERY NIC DETAILS**
  - to find out if your adapter is coupled
  - to find out if your adapter is initialized
  - to find out if your IP addresses have been registered
  - to find out how many bytes/packets sent/received
- **CP QUERY PORT GROUP**
  - To determine the members of a particular groupname
  - To determine which groups are active or inactive

z/VM Security News:

Virtualizing the CryptoExpress4S

z/VM Hardware Crypto Support Updates

z/VM Guest Support for the Crypto Express4S feature

- Guest support for Crypto Express4S (available on zEC12 and zBC12)
- Can be configured in one of three ways:
  - IBM Common Cryptographic Architecture (CCA) Coprocessor
  - IBM CCA Accelerator
    - z/VM supports dedicated and shared modes for CEX4C and CEX4A
  - IBM Enterprise Public Key Cryptographic Standards (PKCS) #11 (EP11) Coprocessor
    - Usable for dedicated cryptographic services for a virtual machine (APDED)
    - No sharing of CEX4P domains

- APAR VM65007 for z/VM 5.4, 6.1, and 6.2 support
- APAR VM65308 for CEX4C sharing (clear key)

“How to:” Virtualize Hardware Crypto Features for z/VM Guests

- **CPACF** – CP-Assisted Cryptographic Facility
  - Feature 3863 (disabled by default, but free to enable on the HMC/SE)
  - On-CPU cryptographic operational assistance
  - Clear key operations only
  - Symmetric algorithms only (DES, 3DES, AES, SHA, SHA-2)

- CPACF is available to any virtual machine if the feature is enabled
  - All modern System z hardware supports this feature
  - The z/VM SSL-TLS Server will use CPACF *automatically*

“How to:” Virtualize Hardware Crypto Features for z/VM Guests

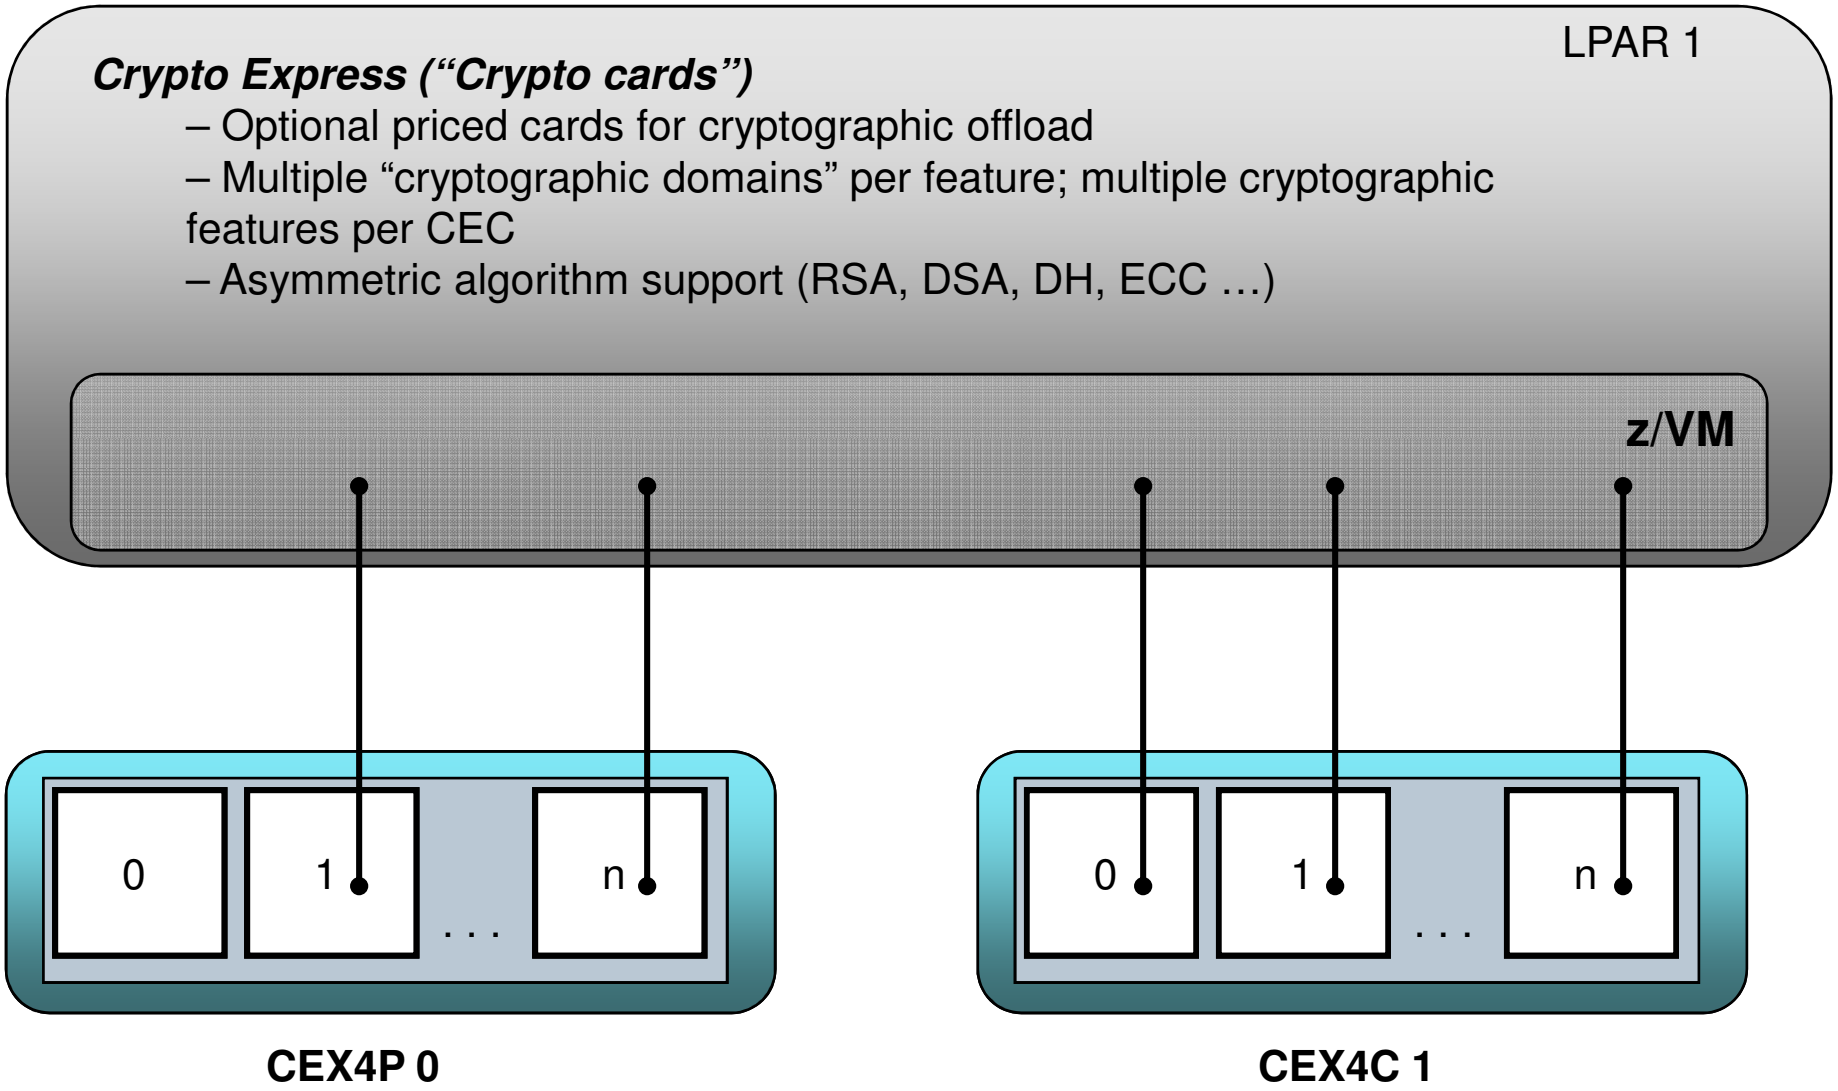

“How to:” Virtualize Hardware Crypto Features for z/VM Guests

The CRYPTO User Directory statement can associate domains/APs from the CryptoExpress features associated with the z/VM instance and assign them to a virtual machine for use:

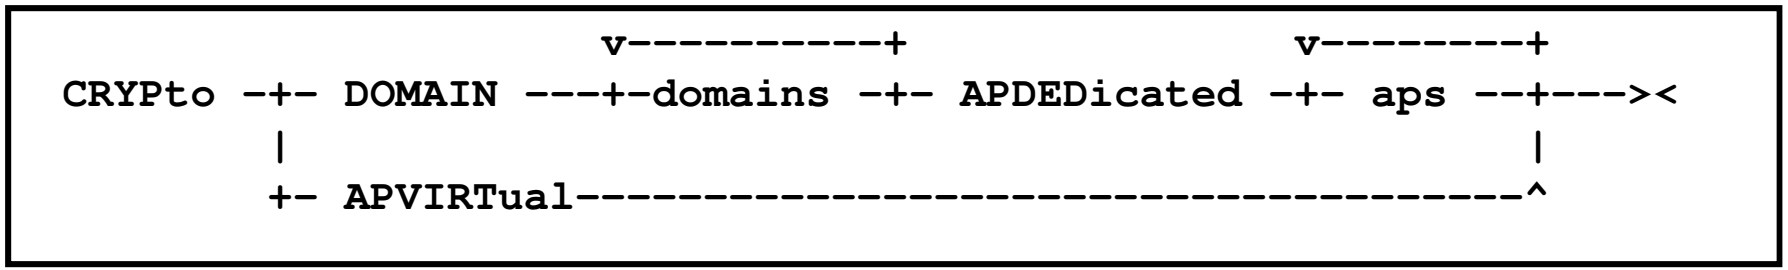

APDED

Domains granted in the directory are “reserved for dedication”; they are not actually in-use until the virtual machine logs on. Then, they are for exclusive use of a single virtual machine.

APVIRT

Access makes use of shared queues controlled by the system. These domains are controlled by the hypervisor, and do not support secure-key operations.

“How to:” Virtualize Hardware Crypto Features for z/VM Guests

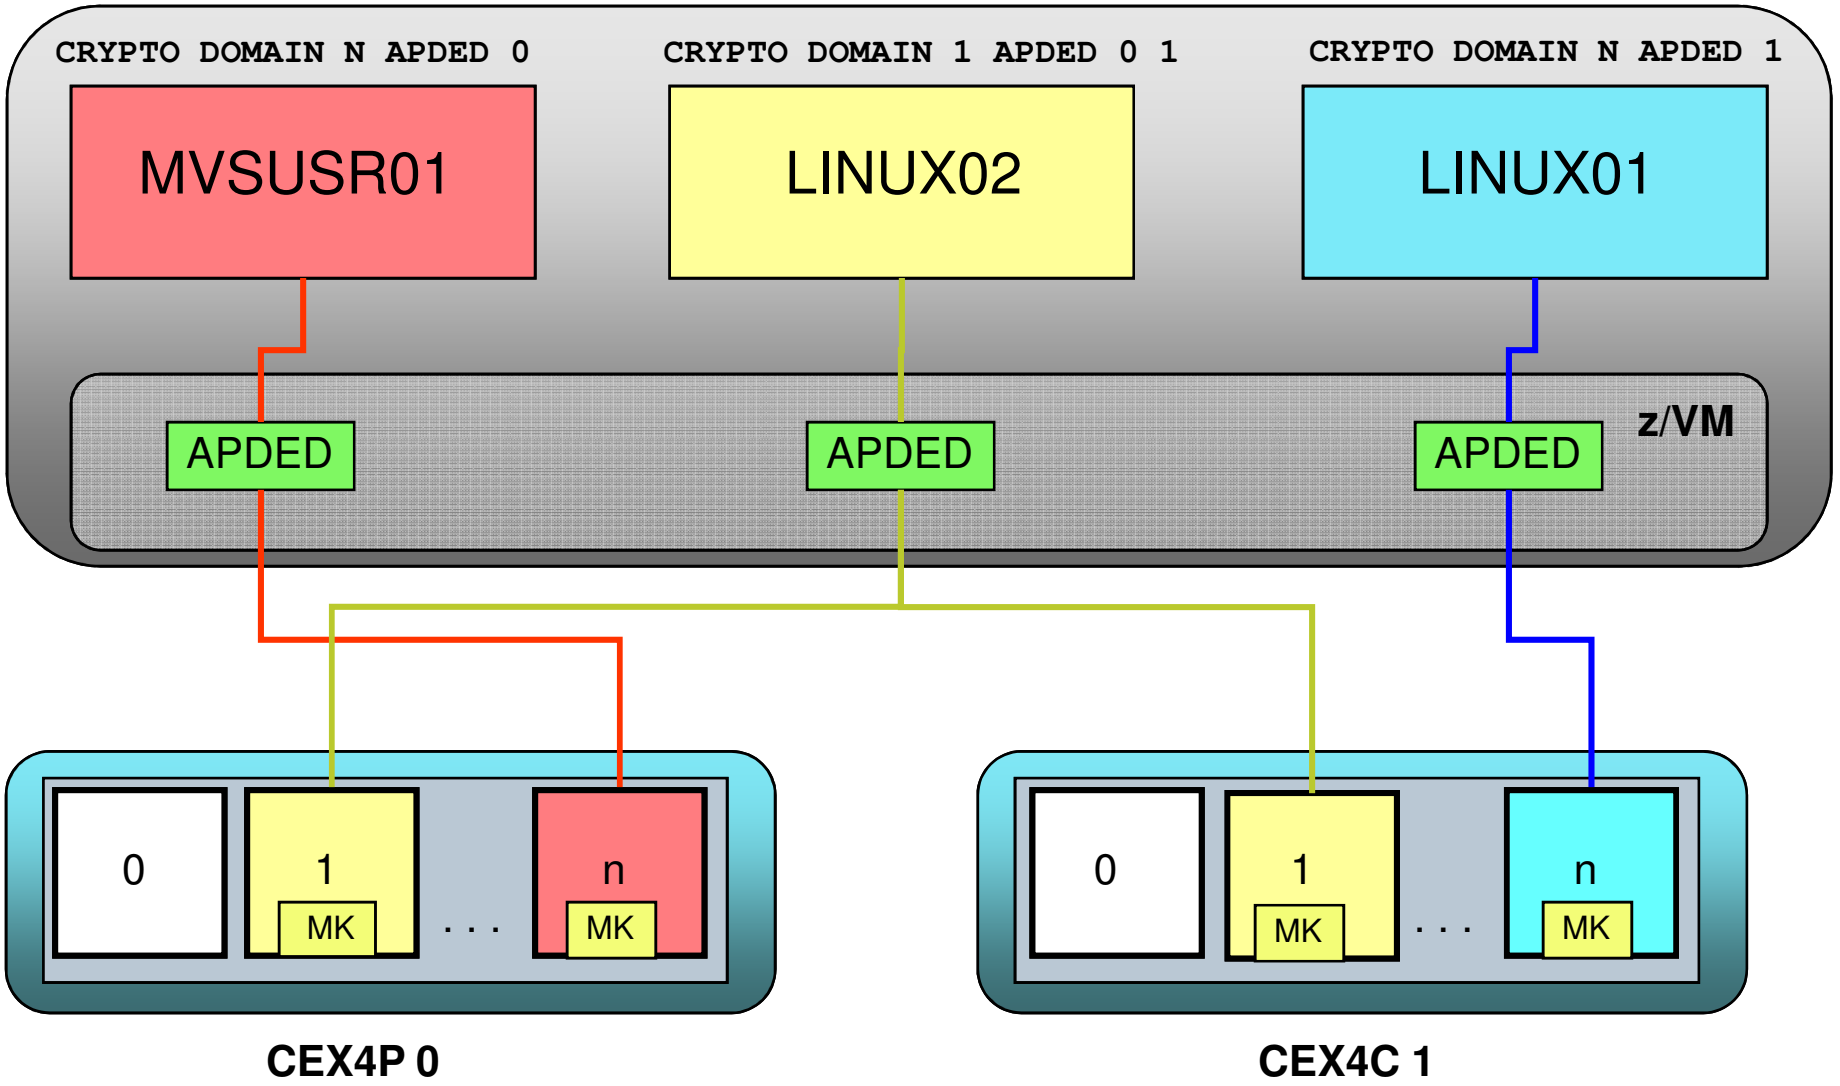

“How to:” Virtualize Hardware Crypto Features for z/VM Guests

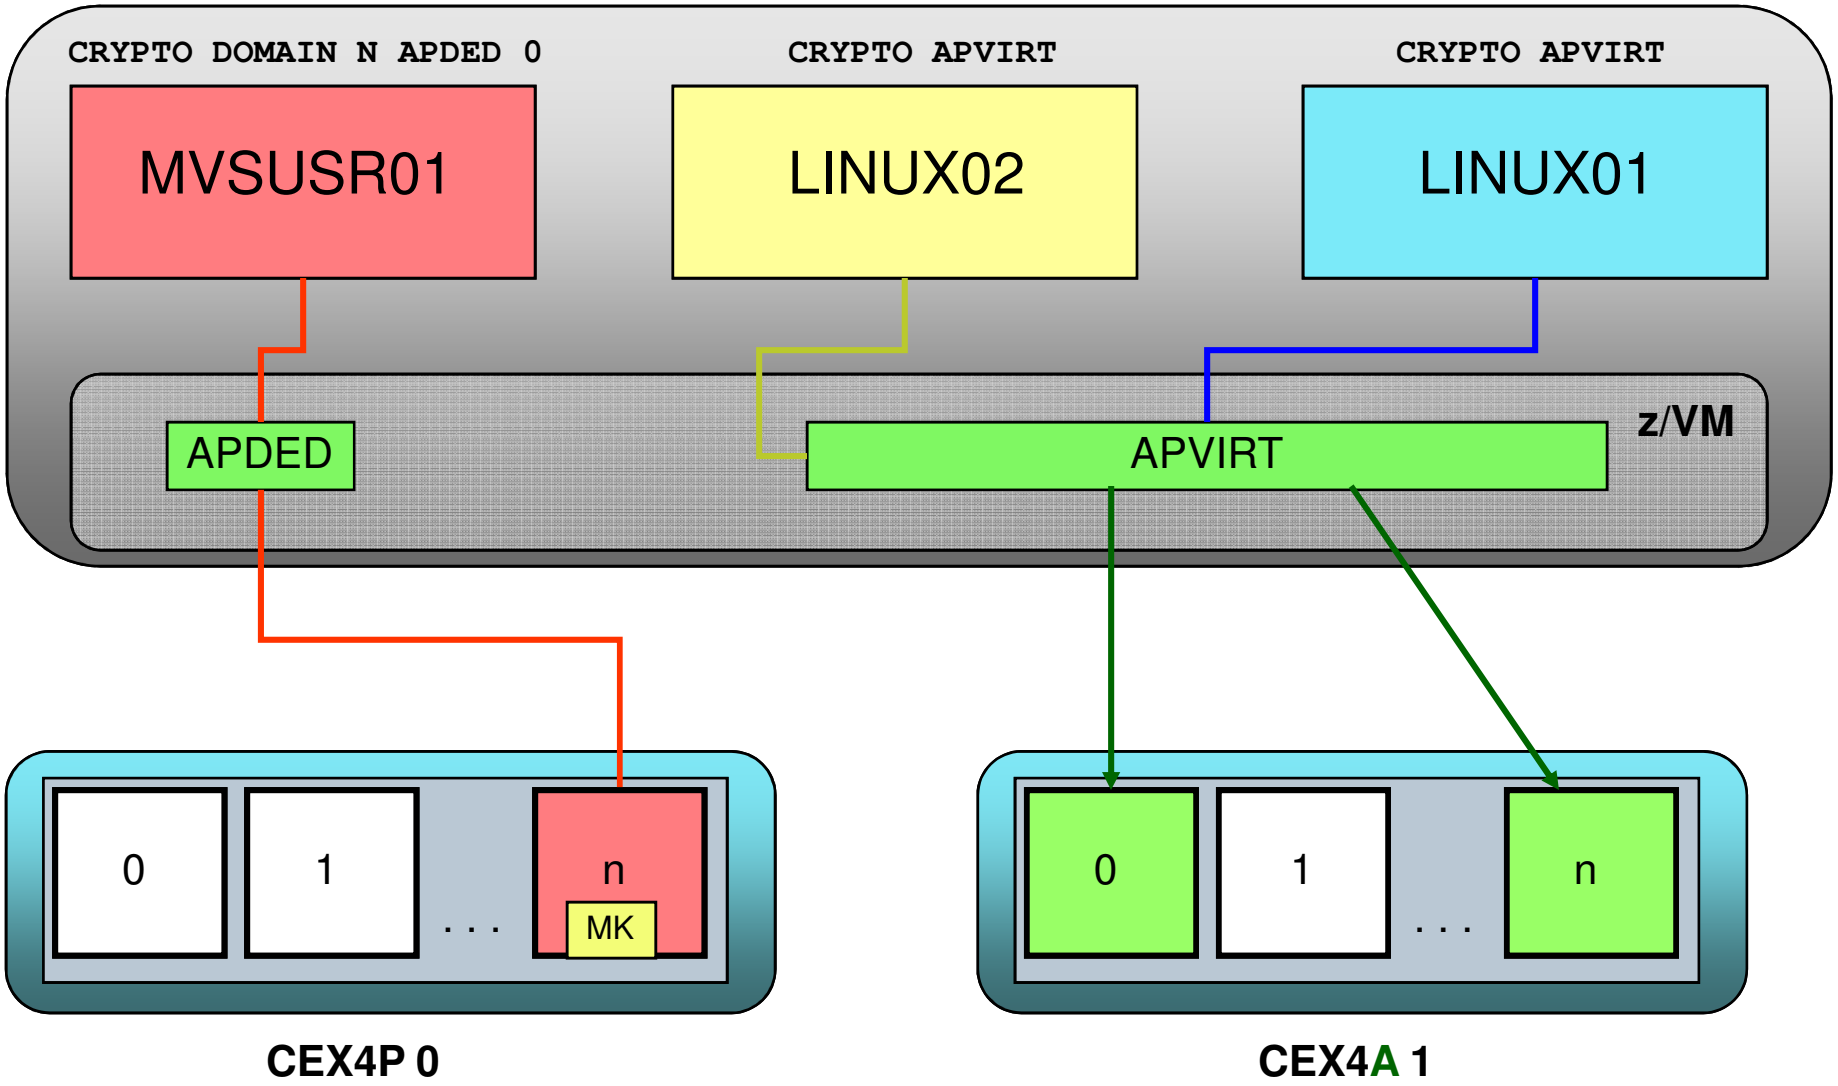

“How to:” Virtualize Hardware Crypto Features for z/VM Guests

QUERY CRYPTO

(Class A, B, C, or E) will display which domains/APs are available. Note that this list will be limited to devices available to a z/VM instance.

```
>>-Query--CRYPto--+--------------------+-----------------><
                  '-DOMains--+-------+-'
                             '-Users-'
```

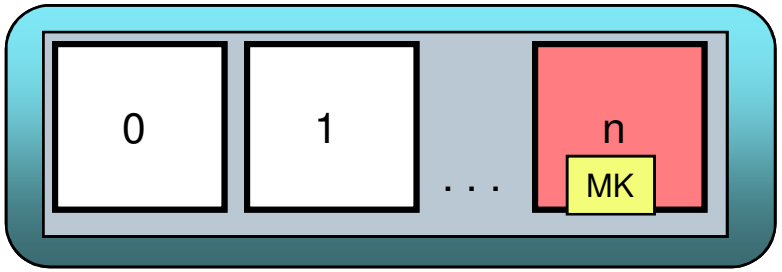

CEX4P 0

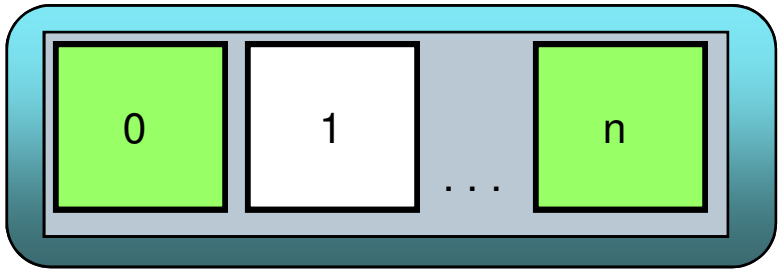

CEX4A 1

“How to:” Virtualize Hardware Crypto Features for z/VM Guests

```
QUERY CRYPTO DOMAINS USERS
```

```
                AP  device   Domain nn  device status   system usage            planned usage
01: AP 02 CEX3C  Domain 08 available    free                    unspecified
01: AP 03 CEX3A  Domain 06 available    dedicated to BWHUGEN    dedication
01: AP 03 CEX3A  Domain 07 available    free                    unspecified
01: AP 03 CEX3A  Domain 08 available    free                    unspecified
01: AP 04 CEX4C  Domain 06 available    free                    dedication
01: AP 04 CEX4C  Domain 07 available    free                    dedication
01: AP 04 CEX4C  Domain 08 available    free                    unspecified
There are no shared-crypto users.
Ready;
```

“How to:” Virtualize Hardware Crypto Features for z/VM Guests

QUERY VIRTUAL CRYPTO

(Class G) will display virtual crypto facilities for your guest.

Keyword “virtual” required for Guests with A, B, C, or E privileges.

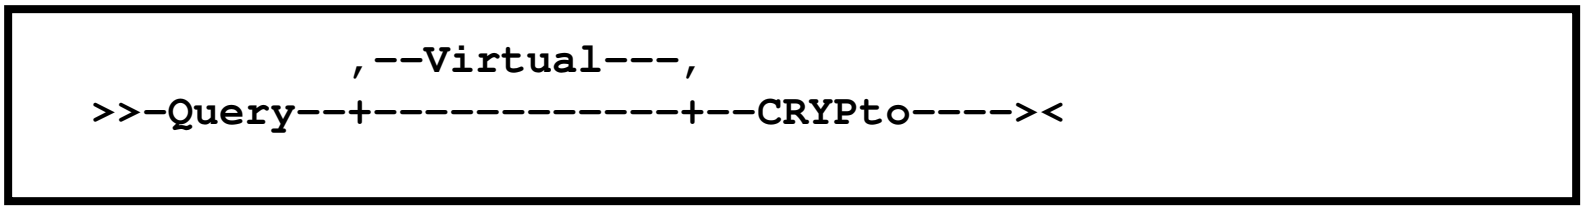


```
QUERY VIRTUAL CRYPTO
```

```
AP 03 CEX3A Domain 06 dedicated
Ready;
```

“How to:” Virtualize Hardware Crypto Features for z/VM Guests

- **The Big Question: Which type of domain do I want to assign to my guest?**

- **It depends:**
  - Do you need **secure** or **protected** key operations? (APDED)
  - Does your security policy require physical isolation? (APDED)
  - **New* Do your guests need to exploit EP11 mode? (APDED only)*
  - Do you need to relocate your guest? (APVIRT*)
  - Can you share your domains without impact to security or performance? (APVIRT)
  - Are you running out of domains attached to the LPAR?
  - Are your guests similar, cloned, or tied to HA solutions?

- Different guests will have different needs, based upon their drivers and configuration requirements.

*Note: some restrictions apply. Consult the *CP Planning and Administration Guide* or *Getting Started With Linux* manuals.

z/VM Security News:

Using RACFVM in a z/VM Single System Image (SSI) Cluster

Security in an z/VM Single System Image Cluster

- A userid has the same password on all systems (Single- or Multi-Configuration)
- A Single Configuration Virtual Machine can only log onto one member of the cluster
  - Error message just like logging onto a userid on the same system
- A Multi-Configuration Virtual Machine is a distinct construct on each system

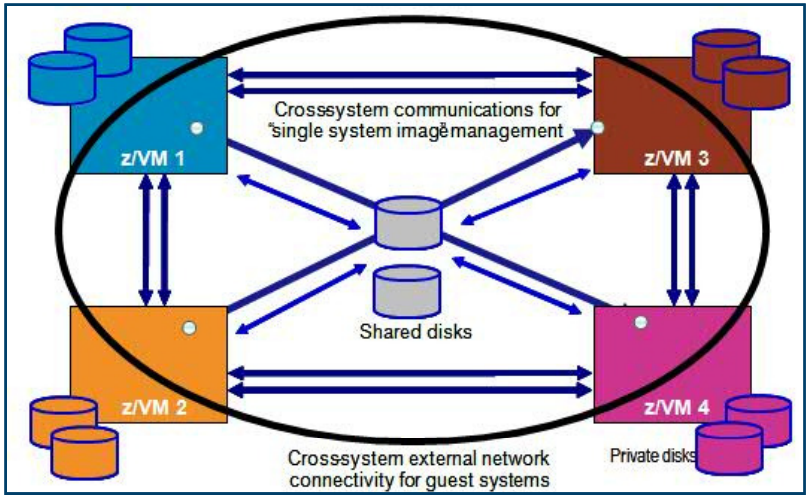

- A userid's privilege classes are the same on every system
  - A common source directory definition
- The cluster maintains a **single security context** for the entire system
  - And an ESM, as with stand-alone systems, extends these capabilities

RACF in a z/VM Single System Image Cluster

- When installed in an SSI, RACF creates *a single security context* for the cluster
  - Shared database and definitions
  - Handshaking of RACFVM instances
  - Cluster-aware auditing

- RACF for SSI is for the entire cluster, it's not something you can enable one step at a time.

- RPIDIRCT has been updated to handle both single-configuration and multi-configuration virtual machines
  - **New* In a mixed-level cluster, use the highest-level RPIDIRCT*

- The virtual machines have been modified to operate both in and out of an SSI …

RACF in a z/VM Single System Image Cluster

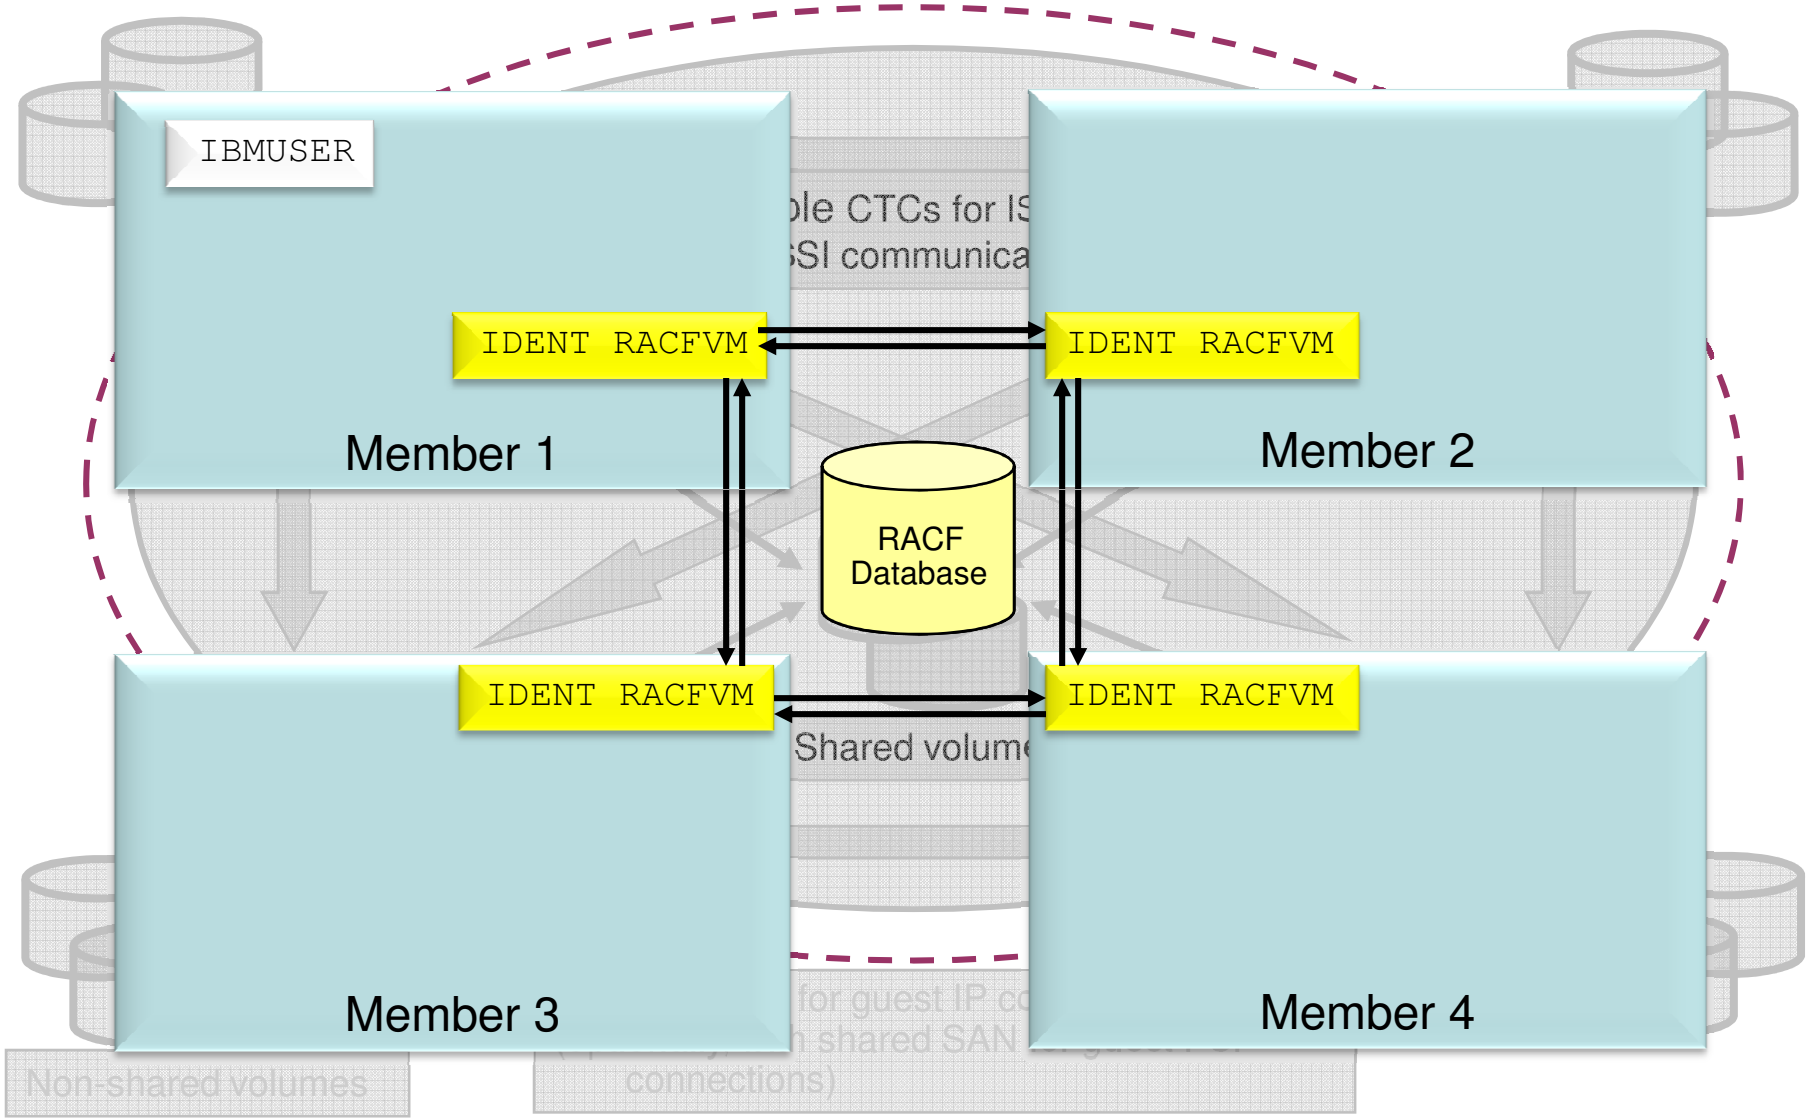

RACF Virtual Machines in an SSI cluster

Handshaking and Command Propagation

- Locking done to ensure RVARY submissions are handled sequentially

- Commands that create broader changes need to be propagated across the cluster
  - SETROPTS
  - RVARY
  - SETEVENT

- RACF will suppress “extra” messages and marshal output when executing “remotely.”

- RAC command, ISPF panels, and R_Admin API (used by LDAP) are interfaces which support command propagation (not the RACF command sessions)

- The propagated commands output from each RACF server on each system is bracketed by the lines:
  - `OUTPUT FROM <racfname> ON SYSTEM <ssinode>`
  - `END OF OUTPUT`

The RACF Database in an SSI

- All RACF servers in SSI must **share** the same RACF database
  - Databases are shareable today
  - Maintain a single security context; no confusion in security policy

- RACF database in SSI must be fullpack minidisk, must support reserve/release and can't be an FBA device
  - Full-pack 3390s for both the primary (200) and backup (300)
  - RDEVICE statements for each in the System Configuration file
  - Minidisk caching is automatically turned off

```
RDEVICE 200 TYPE DASD SHARED YES     /* Default RACFVM db */
RDEVICE 300 TYPE DASD SHARED YES     /* Backup RACFVM db  */
```

- Database synchronization
  - When a member joins, CP+RACF will ensure that the joining server has identical database datasets to those being used and active in the SSI
  - Automatic propagation of RVARY commands

Auditing in RACFVM (An Overview)

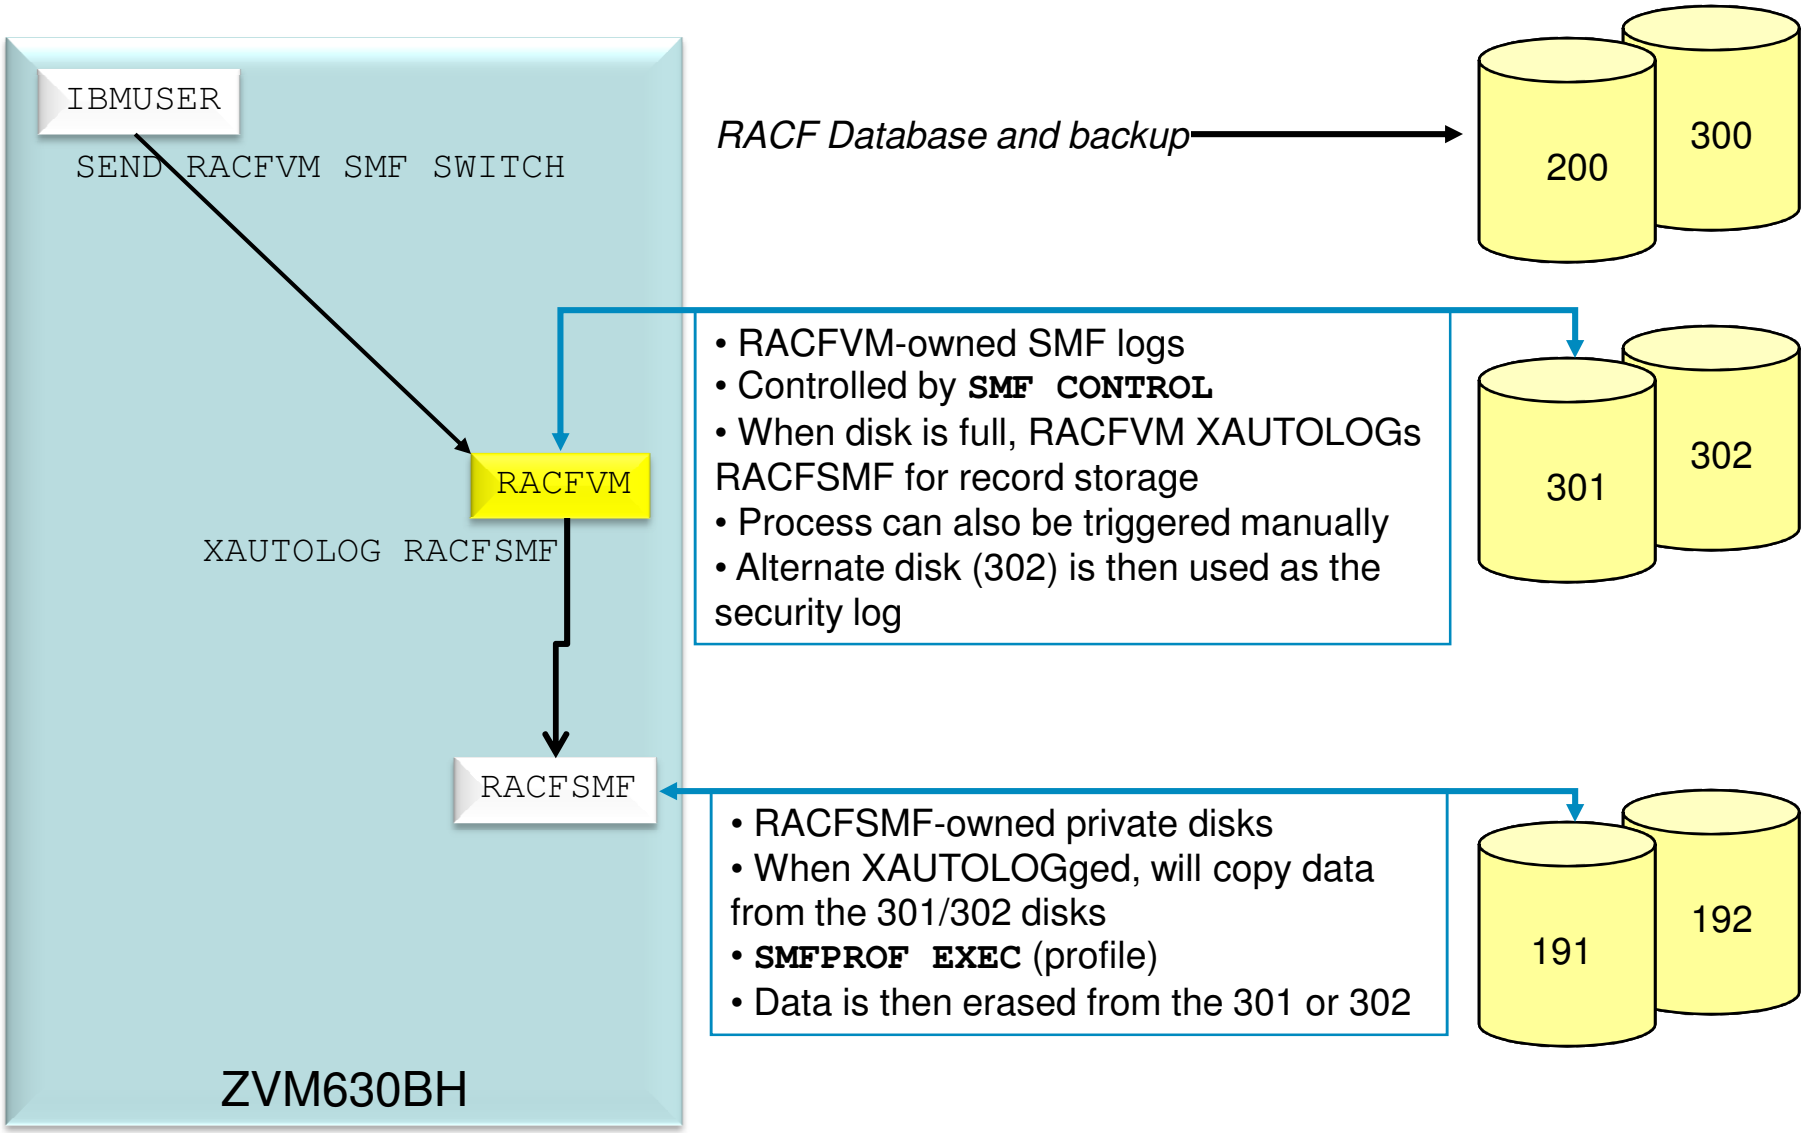

Auditing RACFVM (Cluster View)

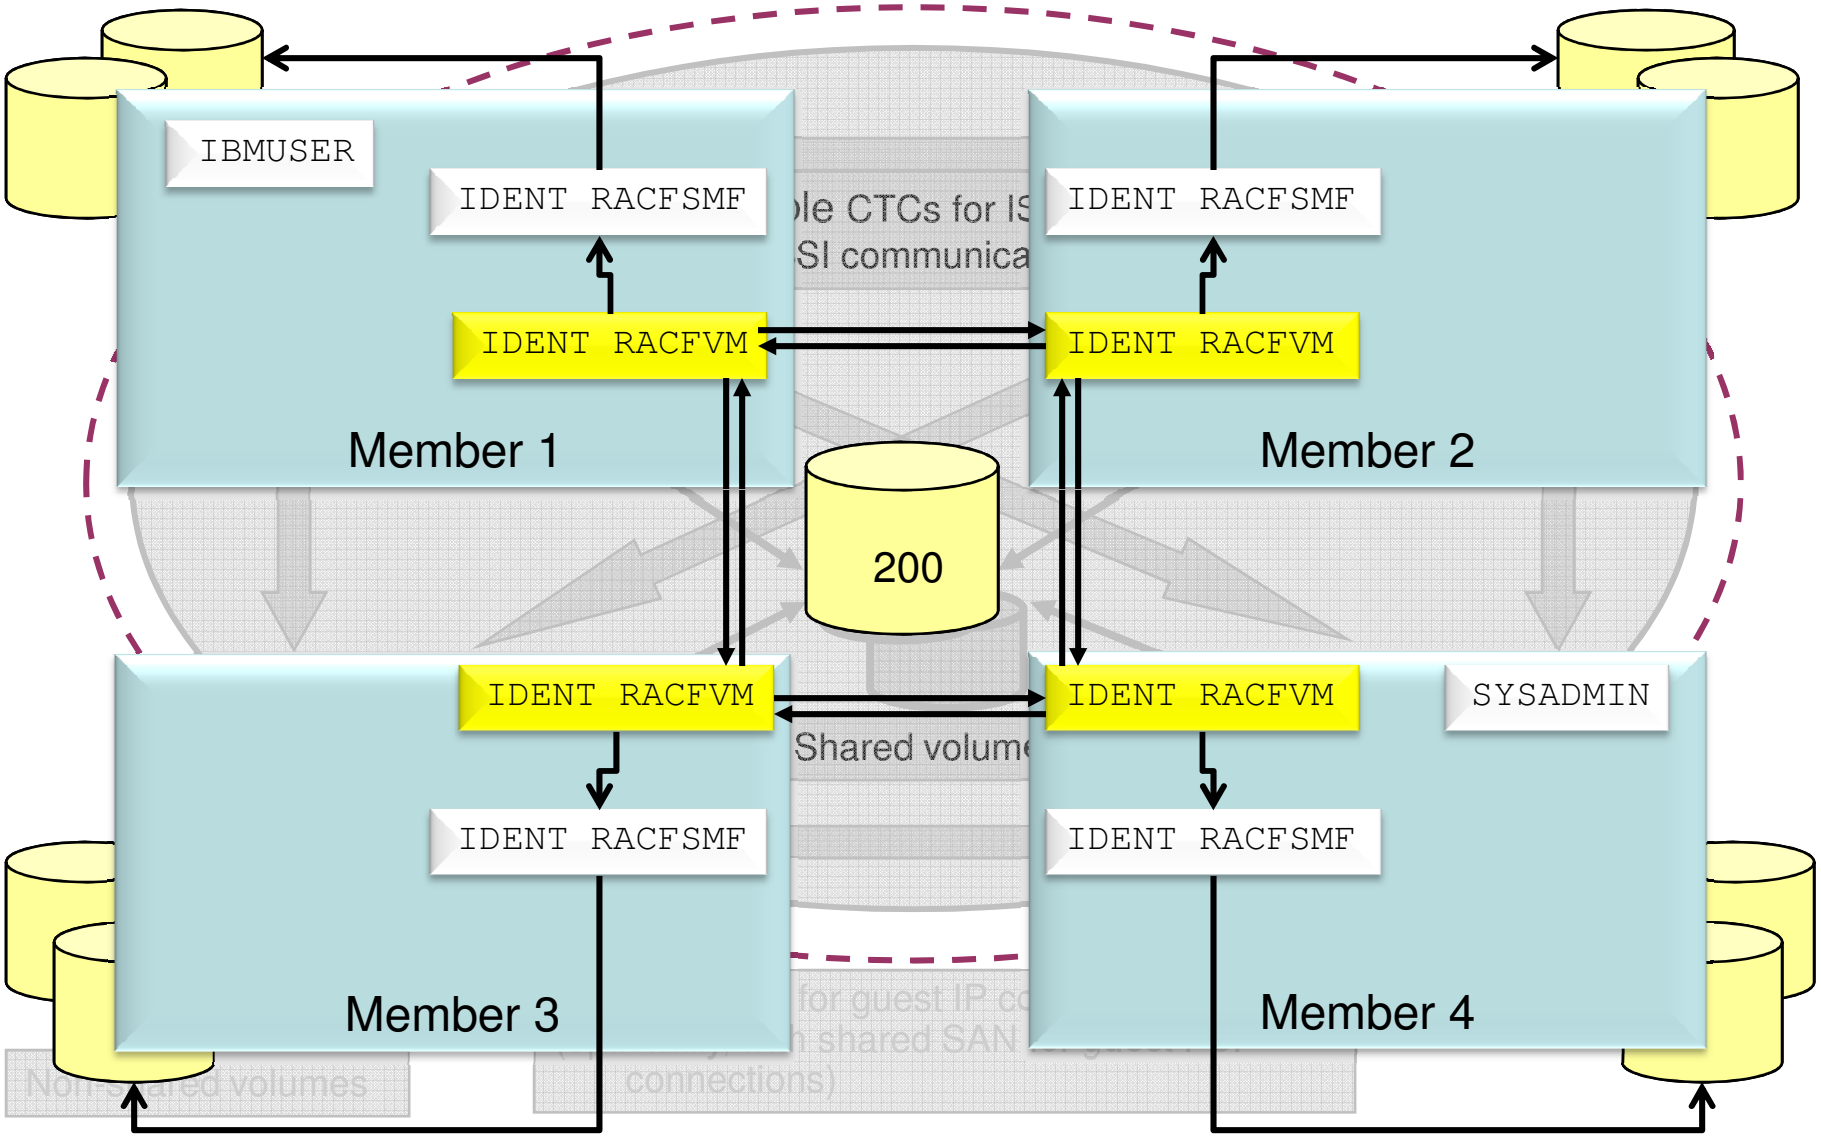

Auditing RACF in a Single System Image cluster

- RACFVM and RACFSMF are multi-configuration virtual machines
  - Shared RACF database
  - ***All other disks are local*** – including 301 and 302 for auditing
  - RACFVM: Separate SMF CONTROL files operating against a single security context
  - RACFSMF: Separate SMFPROF EXEC files, 191 and 192 disks

- In the case of some commands – the AT command in particular – auditing records will appear on the destination system
  - `AT_LOGON, AT_FROM, AT_LOGOFF`

- Auditing automation should account for this disparity to gather all pertinent audit records
  - Make sure all SMF CONTROL files are modified as appropriate
  - Make sure auditing policy and SMF records are managed accordingly

RACF and Live Guest Relocation

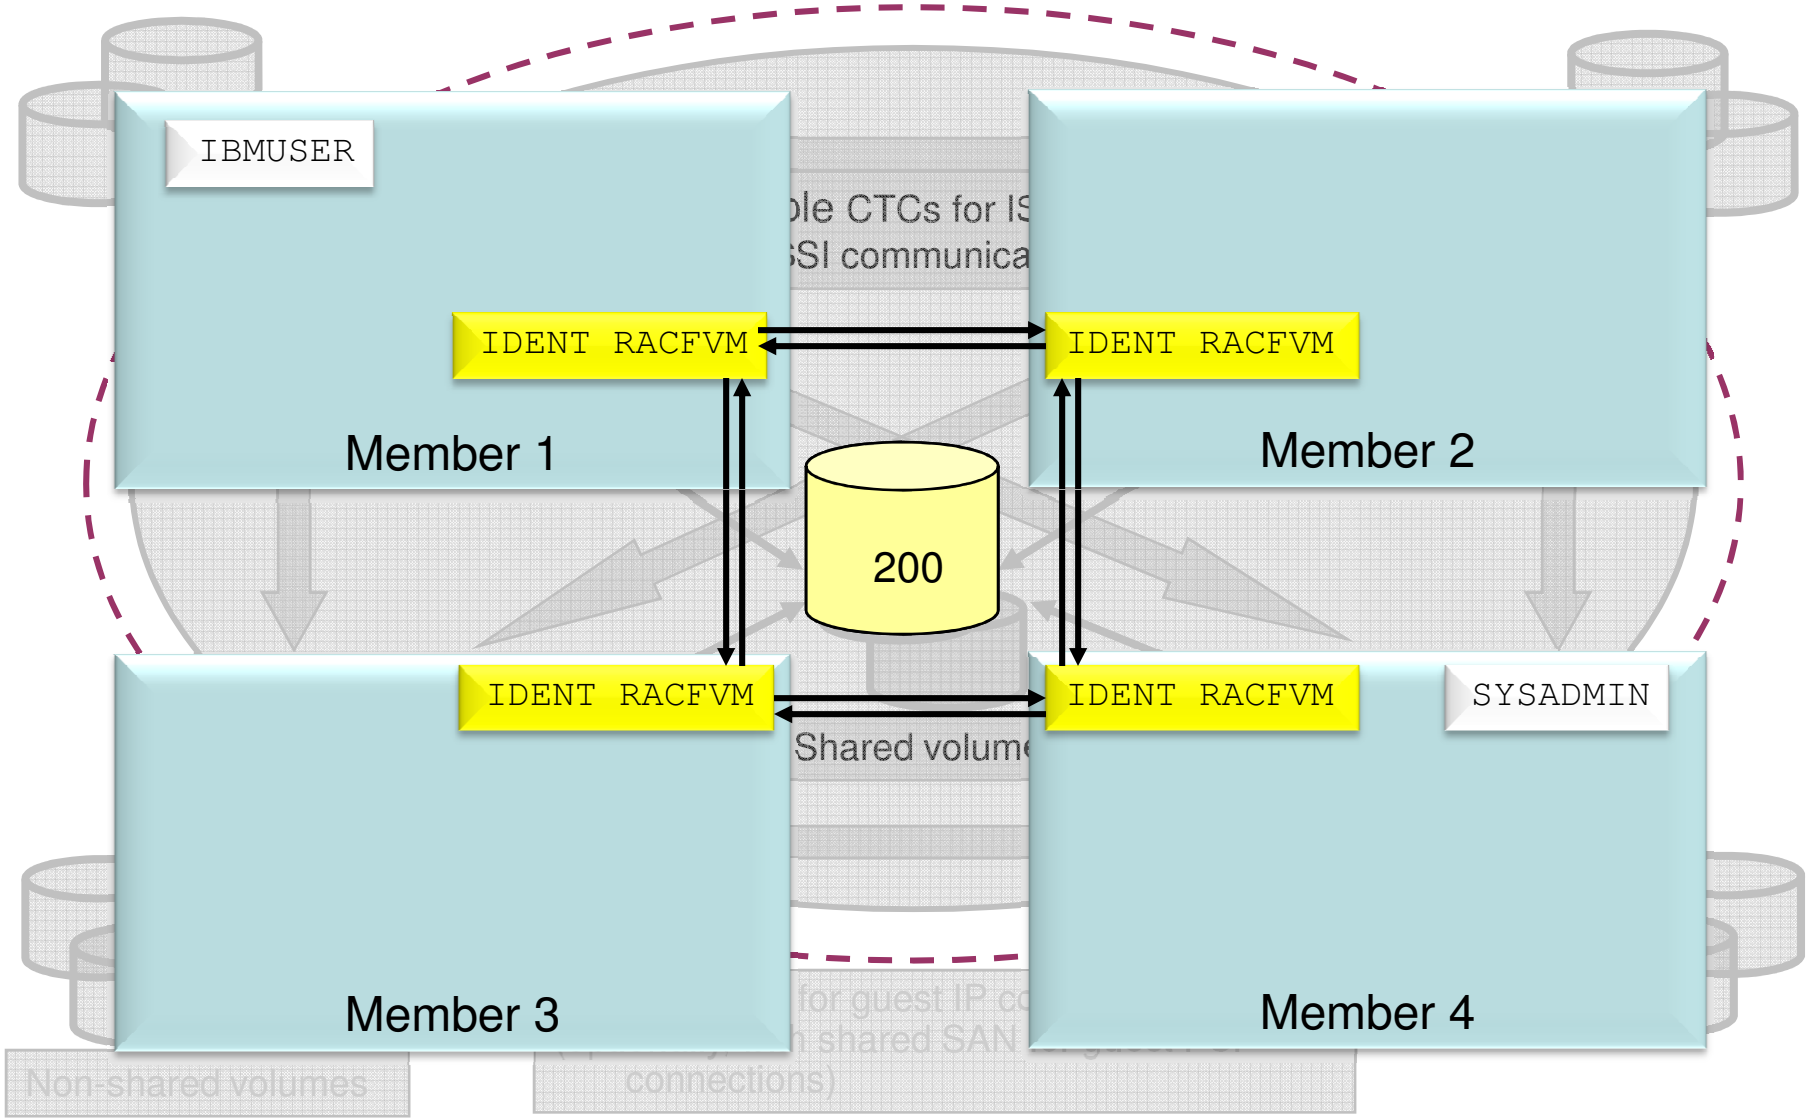

RACF and Live Guest Relocation

Live Guest Relocation

- `VMRELOCATE MOVE USER` *`userid`* `TO` *`sysid`*
  - Class A command

- RACF cleans up a user's presence on the source system, and prepares for the target system for the relocate-logon of the user

- Generate LOGOFF/LOGON auditing events on source/target system, to note the transition

- RACF perspective of relocate events:
  - User data is created for *userid* on *sysid* with all the above
  - User resources are allocated on sysid
  - Associated authorization calls are approved without a RACF check
  - Relocate-logon is requested for *userid* on *sysid* when the inbound relocation is complete

Security Zones in an SSI

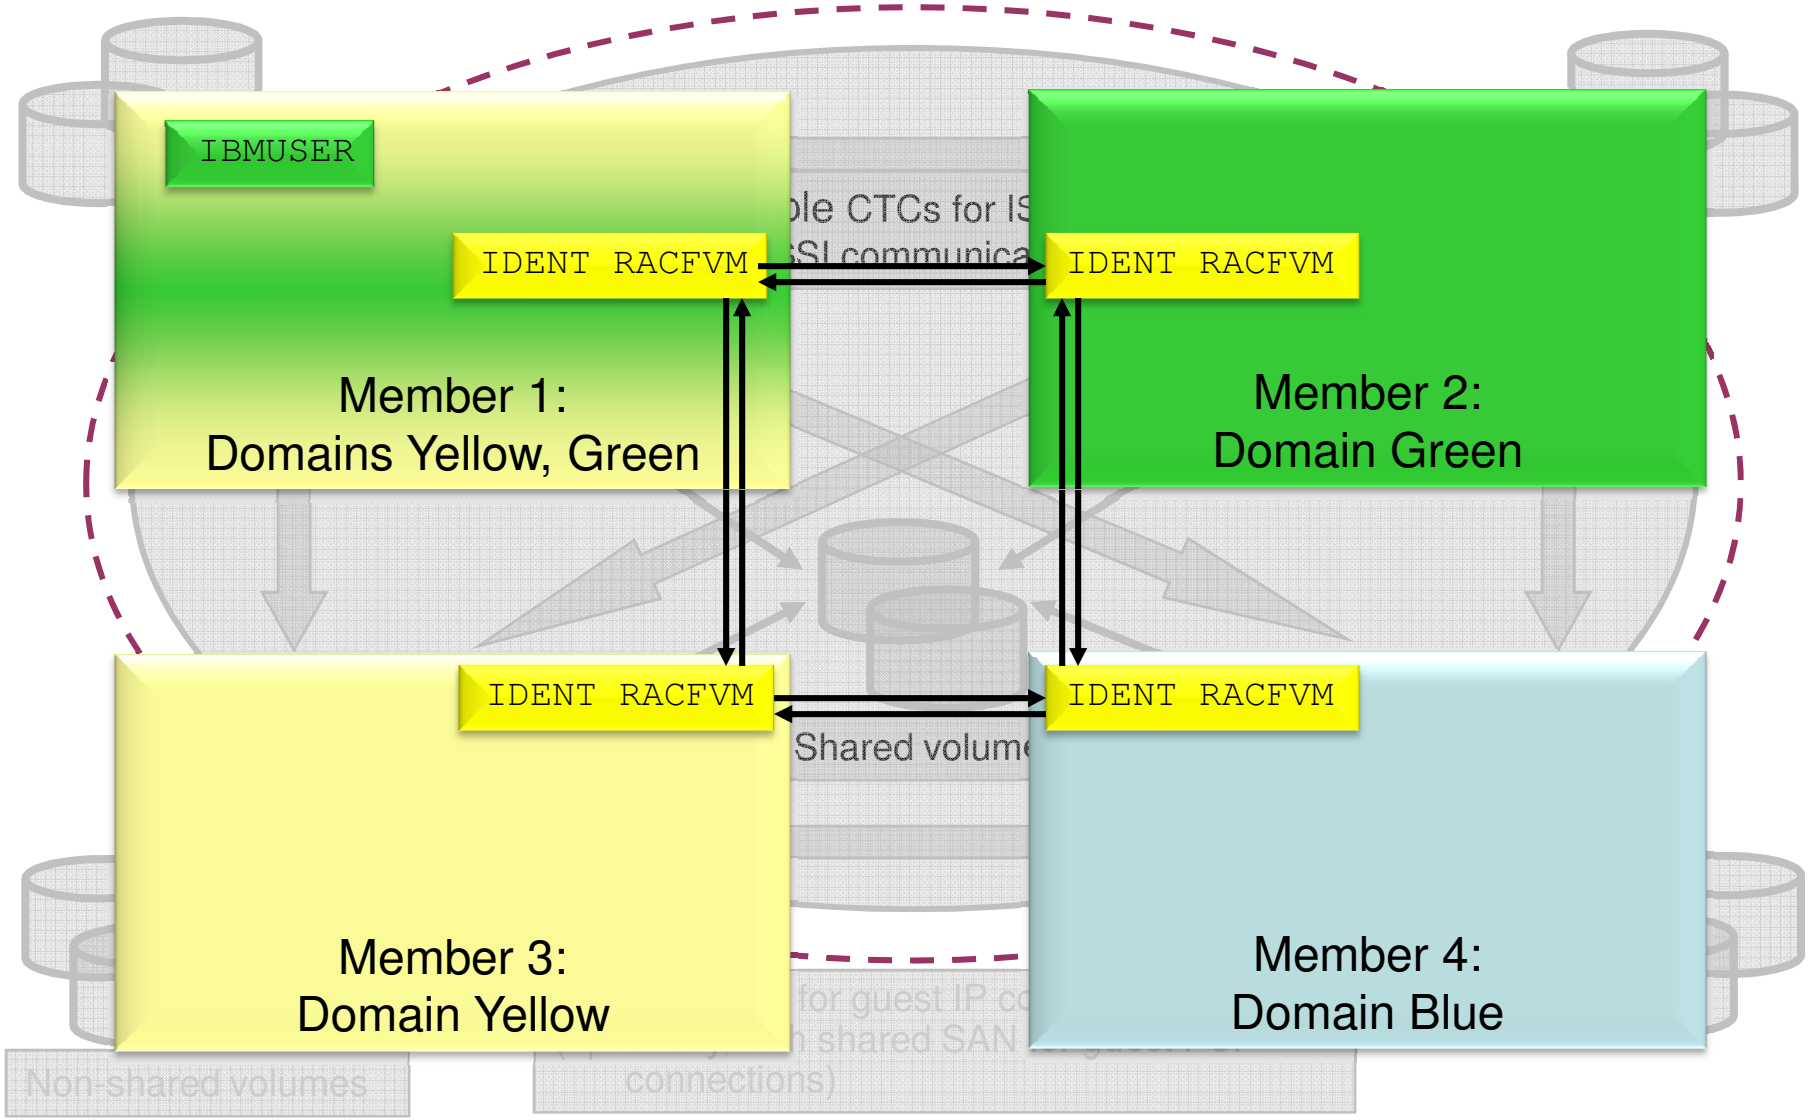

Security Zones in an SSI

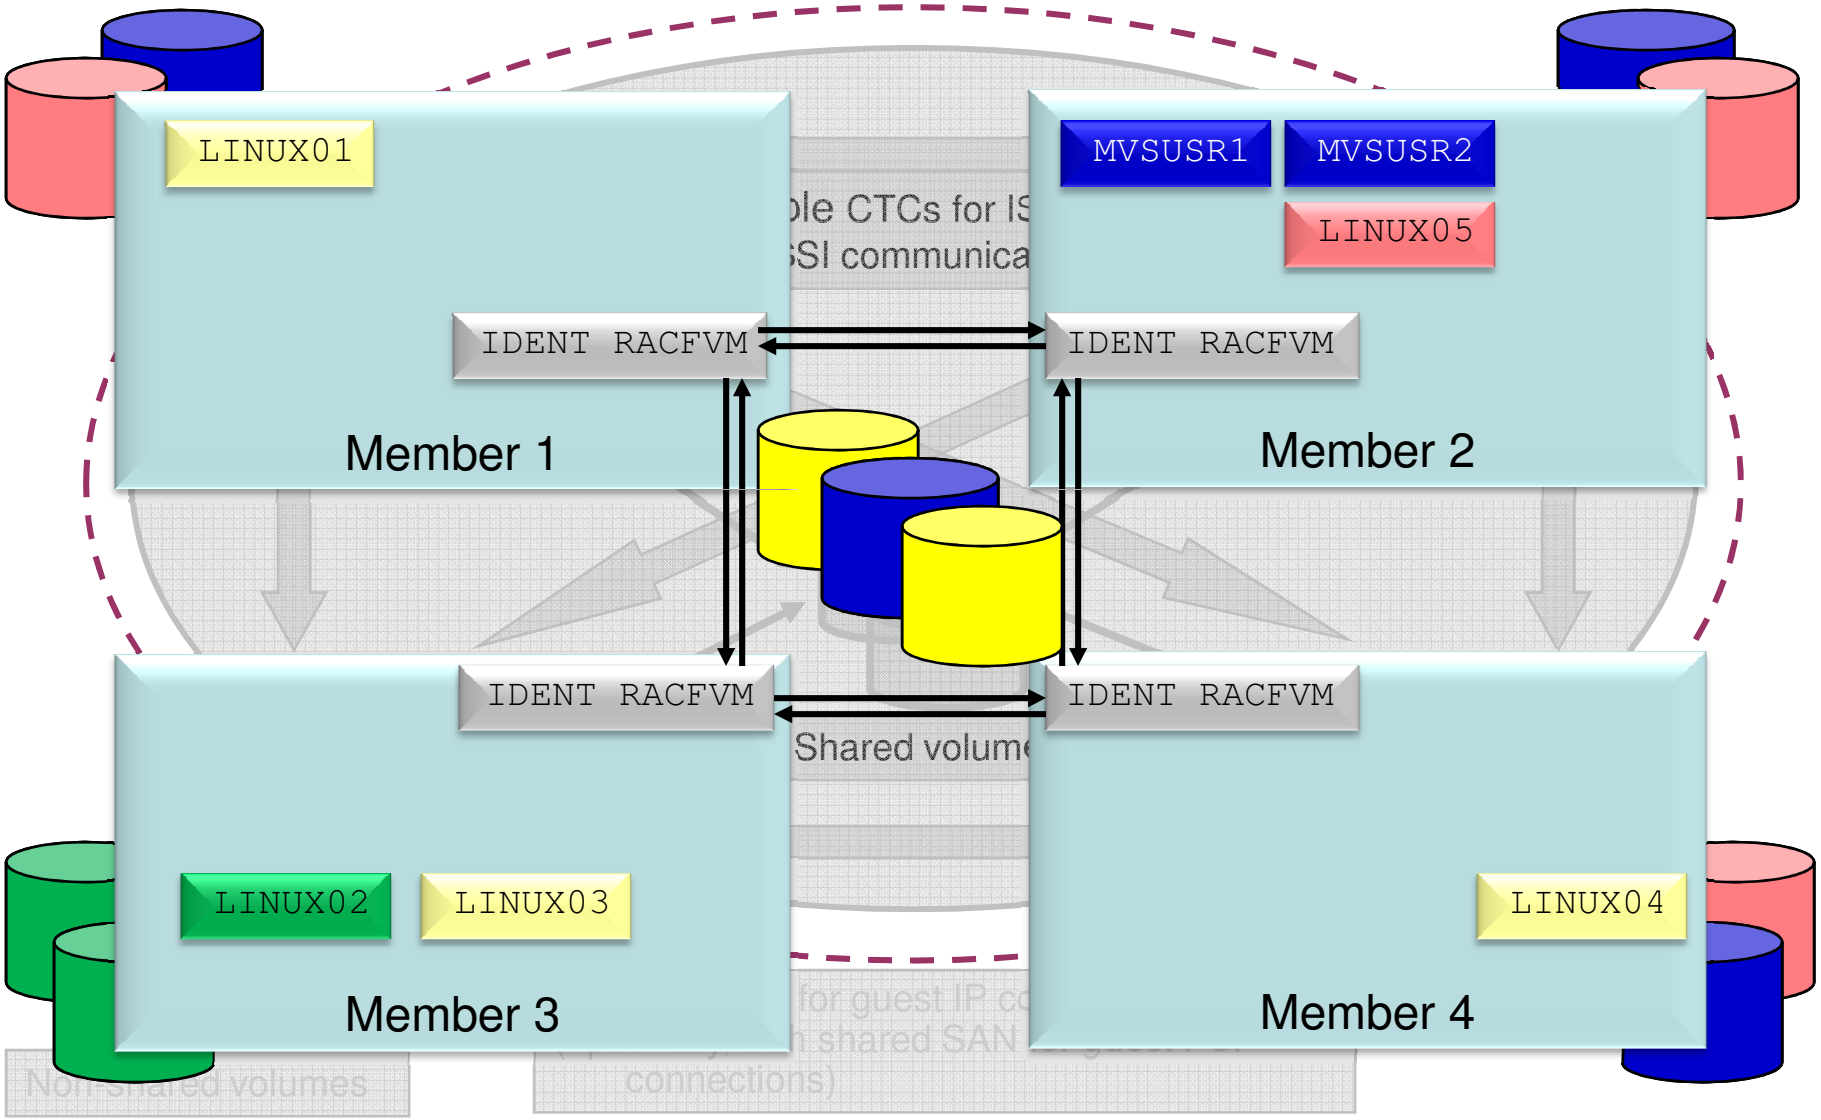

Migrating to RACF in an SSI

- **You can have an ESM and still migrate to SSI!**
  - Step 1: If you don't have an ESM, get one.
  - Line up the shared DASD required for the database; remember that this needs to be a fullpack minidisk!
  - If you're converting one or more ESM-controlled systems into an SSI:

Migrating to RACF in an SSI

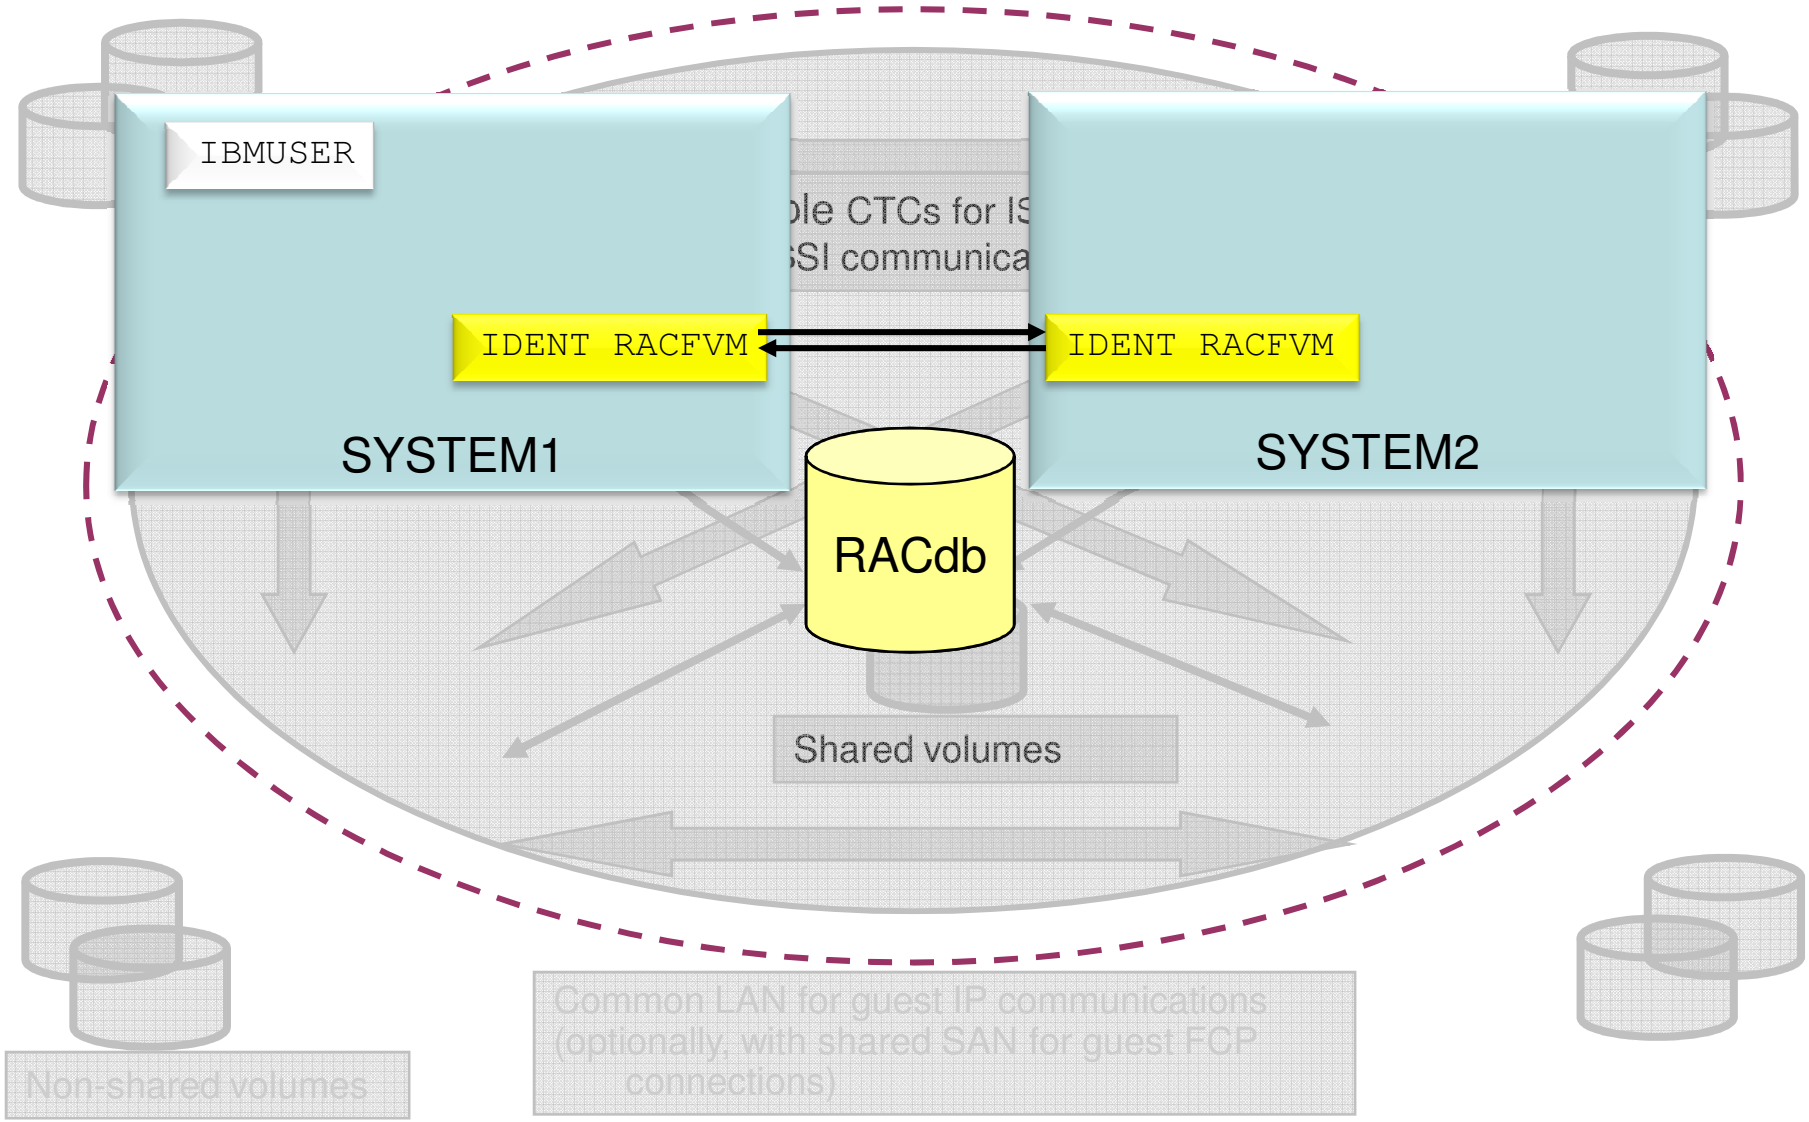

Migrating to RACF in an SSI

- **You can have an ESM and still migrate to SSI!**
  - Step 1: If you don't have an ESM, get one.
  - Line up the shared DASD required for the database; remember that this needs to be a fullpack minidisk!
  - If you're converting one ESM-controlled systems into an SSI:
    - Migrate to 6.2 in a **non-SSI format**
    - Convert associated resource profiles to 6.2 format, using RPIDIRCT as necessary
    - Take the steps to enable SSI; turn on RACFVM as part of the outlined process
  - Note for z/VM 6.3: RPIDIRCT is now located on the PMAINT.551 minidisk
    - For mixed-release z/VM clusters
    - Change to existing RACFVM machine definitions

Notes on Using a Fullpack Minidisk for RACFVM Database

- If you're migrating from a non-SSI RACF-secured system, you may need to **convert your database** to reside on a Fullpack Minidisk
  - DDR the Database (and backup) to the Fullpack minidisk
  - Increase database allocation if pertinent (e.g., if merging multiple systems)
  - Remember to issue RACFCONV to upgrade if installing new support at the same time

- Refer to the *RACF System Programmer's Guide* (SC24-6219-02) for more information
  - Chapter 4: "Operating Considerations Unique to z/VM"
  - Subheading: "RACF Database on Fullpack Minidisk"

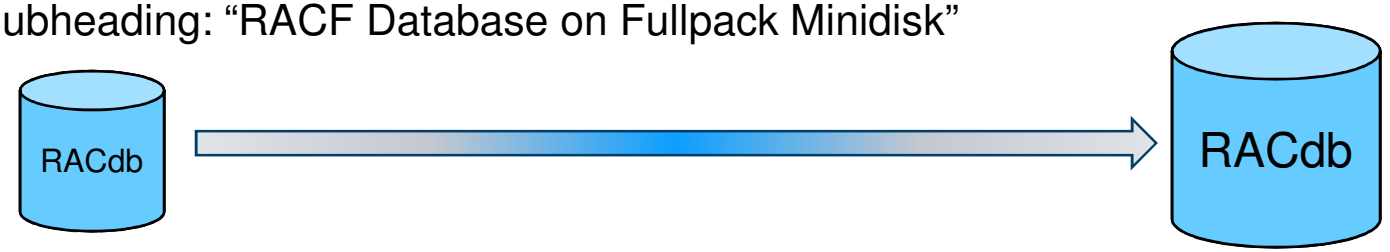

Migrating to RACF in an SSI

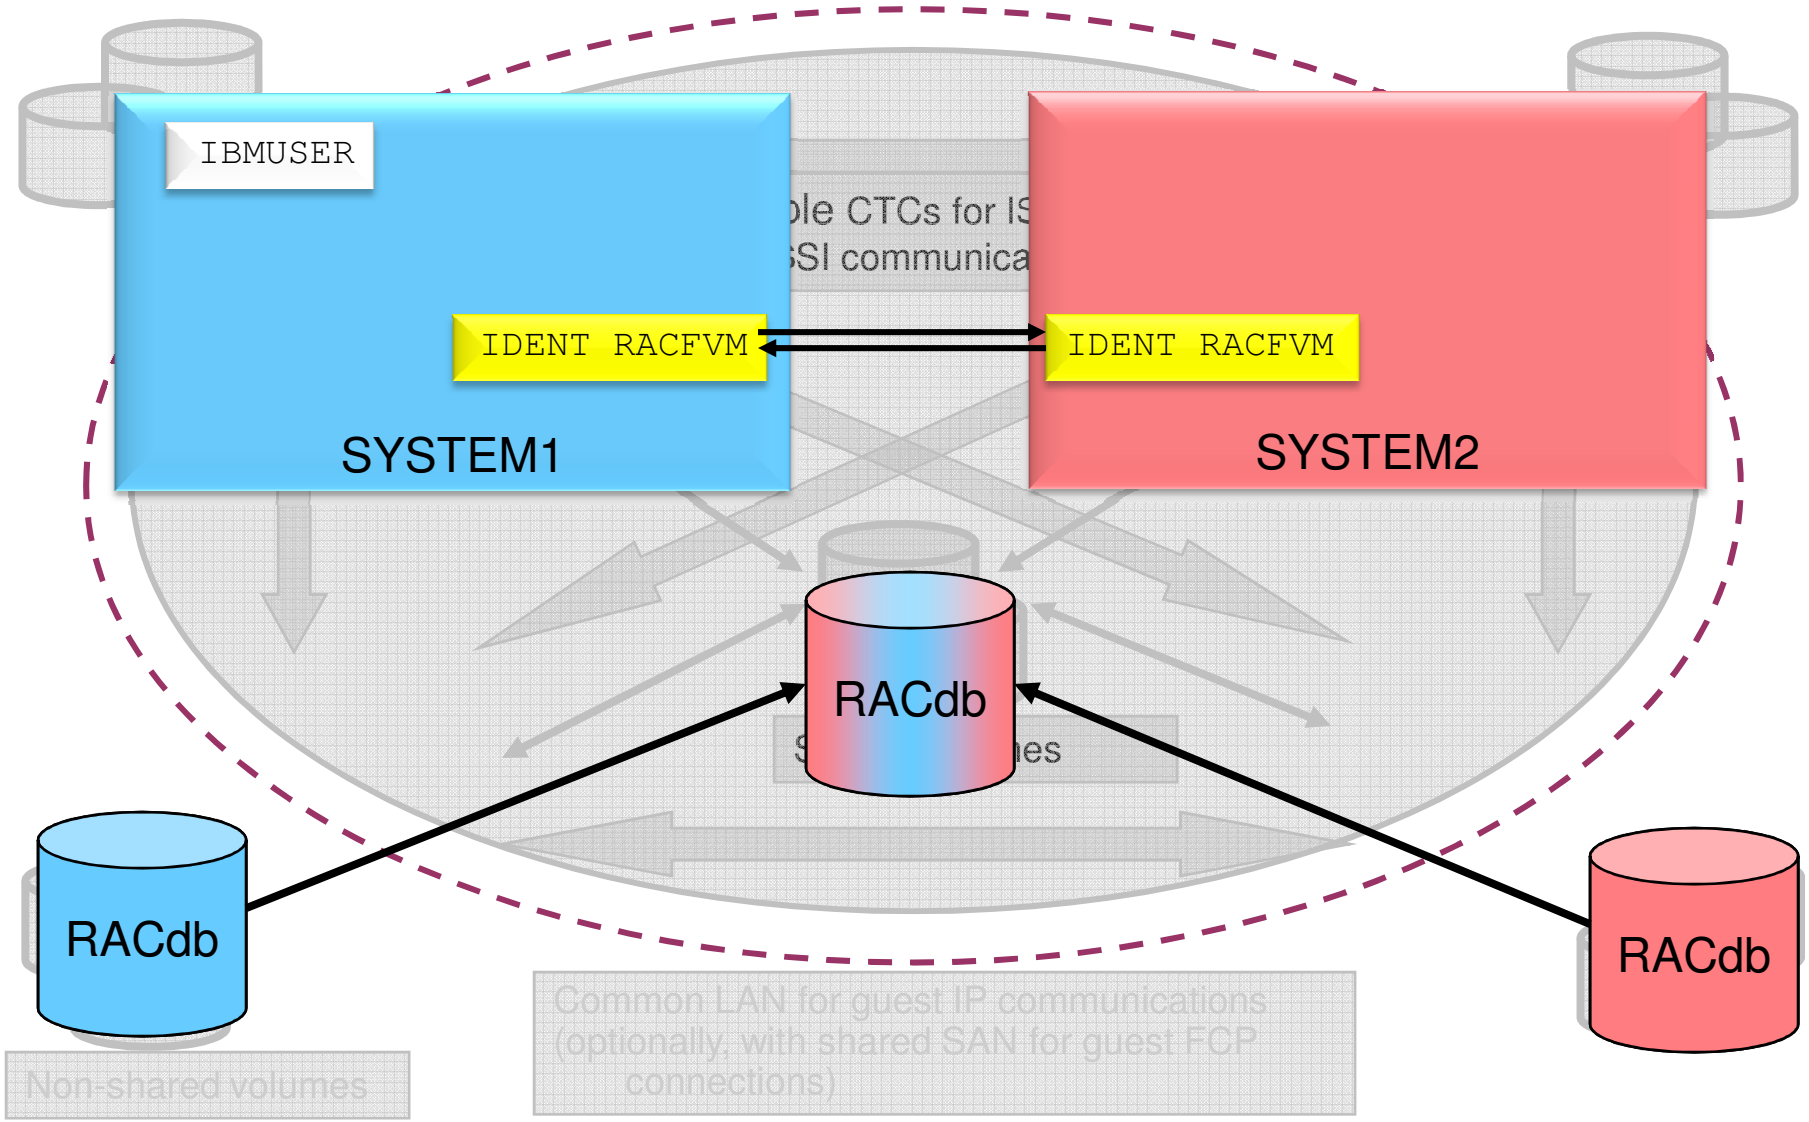

Migrating to RACF in an SSI

- **You can have an ESM and still migrate to SSI!**
  - Step 1: If you don't have an ESM, get one.
  - Line up the shared DASD required for the database; remember that this needs to be a fullpack minidisk!
  - If you're converting one or more ESM-controlled systems into an SSI:
    - Migrate to 6.2 in a non-SSI format
    - Convert associated resource profiles to 6.2 format, using RPIDIRCT as necessary
    - Take the steps to enable SSI; turn on RACFVM as part of the outlined process
  - If you're converting two (or more) distinct ESM-controlled systems to an SSI
    - You will need to merge the databases
    - You may want to consider which of your 2+ systems has the most complex security context before choosing which one is the "master" system
    - After one system is enabled, make directory and RACF database updates for the secondary system

Any

Questions?

Compliance

Certifications validate the high standard of z/VM security. z/VM is continually updated to keep pace with changes in security standards, and will pursue security certifications for z/VM 6.3, including FIPS 140-2 (now complete) and Common Criteria (in process).

Confidentiality

z/VM continues to deliver functional updates to keep pace with modern security requirements, via support for new cryptographic developments and the maintenance of security policy in a Single System Image environment.

.

Confidence

With over forty years of security design, delivery and evaluation experience, z/VM continues to secure the road to Smarter Computing.

For More Information …

On the web:

- z/VM Security: http://www.VM.ibm.com/security
- System z Security: http://www.ibm.com/systems/z/advantages/security/
- ***Security for Linux on System z*** (SG24-7728), IBM RedBooks
- *z/VM Secure Configuration Guide*: http://publibz.boulder.ibm.com/epubs/pdf/hcss0b30.pdf

Contact Information:

Brian W. Hugenbruch, CISSP
z/VM Security Design and Development
bwhugen at us dot ibm dot com

@Bwhugen

Dank u
Dutch

Merci
French

Спасибо
Russian

Gracias
Spanish

شكراً
Arabic

감사합니다
Korean

Tack så mycket
Swedish

धन्यवाद
Hindi

תודה רבה
Hebrew

Obrigado
Brazilian Portuguese

谢谢
Chinese

Dankon
Esperanto

Thank You

ありがとうございます
Japanese

Trugarez
Breton

Tak
Danish

Danke
German

Grazie
Italian

நன்றி
Tamil

děkuji
Czech

ขอบคุณ
Thai

go raibh maith agat
Gaelic