July 2009

Cryptographic Hardware Use Cases for Web Servers on Linux on IBM System z

Introduction

Table of Contents

Introduction

The following use cases show examples of using IBM System z® cryptographic hardware with Linux®. The scenarios described involve one of the following Web servers:

- Apache2
- IBM Tivoli® Access Manager for e-Business WebSEAL 6.1
- IBM HTTP Server 6.1.

The intent of showing these scenarios is to provide examples of some of the common tasks one might need to perform when configuring one of these Web servers to use cryptographic hardware. Hopefully there is enough variety among the use cases so that, if the scenario you are attempting is not an exact step-by-step match with a scenario shown, there is enough information among all the scenarios to be able to infer the required steps.

Of the three Web servers discussed, Apache differs from WebSEAL and IHS, in that Apache accesses cryptographic hardware on Linux on IBM System z using an OpenSSL 'IBMCA engine', which interfaces to the IBM Crypto Adapter library (libica). The IBMCA engine support was included in OpenSSL prior to version 0.9.8. As of version 0.9.8, the IBMCA engine is shipped as a dynamic engine in a separate openssl-ibmca RPM.

WebSEAL and the IBM HTTP Server both use the PKCS#11 API to access cryptographic hardware on Linux on IBM System z. Therefore, several of the use cases include examples of configuring the PKCS#11 environment, using the utility provided by the openCryptoki package, which provides PKCS#11 support. WebSEAL and the IBM HTTP Server both also use IBM's Global Security Toolkit (GSKit), for cryptographic operations. Therefore, several of the scenarios show the use of the key management tool provided by GSKit version 7 – named gsk7ikm, to do common key management tasks that deal with cryptographic hardware.

The scenarios shown are the following:

- Configuring Apache2 to use Cryptographic Hardware Devices on Novell SUSE Linux Enterprise 9 (SLES 9)
- Configuring Tivoli Access Manager for eBusiness WebSEAL 6.1 to use Cryptographic Hardware on Linux for IBM System z
  - Configuring WebSEAL 6.1 to use cryptographic hardware on Novell SLES 9
  - Configuring WebSEAL 6.1 to use cryptographic hardware on Red Hat Enterprise Linux 5 (RHEL 5)
  - Configuring WebSEAL 6.1 to use cryptographic hardware on Novell SUSE Linux Enterprise 10 (SLES 10)
- Creating a certificate request and receiving a signed certificate using a PKCS#11 Cryptographic Token
- Using a Test certificate from Thawte with the PKCS#11 Token and IBM HTTP Server 6.1 on Novell SLES 10
- Importing a PKCS12 file containing a personal certificate and private key, plus intermediate and root CA certificates, into a PKCS#11 Token
- Copying PKCS#11 Token Information from one Linux on IBM System z Image to another
  - Copying PKCS#11 Token Information - Novell SLES 9
  - Copying PKCS#11 Token Information - Novell SLES 10

Notes and References

These use cases involve clear-key cryptography, with cryptographic hardware used for acceleration. The private keys used in the cryptographic operations are stored encrypted, within the Linux file system, and are decrypted before being used by the hardware to accelerate the particular cryptographic operation being performed.

There are two types of IBM System z cryptographic hardware used in the scenarios:

Crypto Express2 Feature: This is a pluggable feature containing two co-processors. The co-processors (or 'cards') can provide secure key storage and acceleration when configured in CEX2C mode, or, acceleration only when configured in CEX2A mode. In the following scenarios, the Crypto Express2 Feature is used only as an accelerator for the asymmetric cryptographic operations that take place during the SSL handshake, to establish an SSL session. The Web server's private key is stored encrypted, in the Linux file system, and is decrypted and transferred to the Crypto Expresss2 feature, where the premaster secret decrypt part of the SSL handshake is performed.

CP Assist for Cryptographic Function (CPACF): This is a set of problem-state machine instructions that can accelerate symmetric cryptographic operations:

- DES,TDES on IBM eServer™ zSeries® 990, 890 (z990, z890)
- AES-128 on IBM System z9®
- AES-256 on IBM System z10™.

Also MAC hashing algorithms:

- SHA1 on z990/z890
- SHA-256 on z9®
- SHA-384, SHA-512 on System z10.

Also, Pseudo Random Number Generator (PRNG) beginning with the IBM System z9.

Included in the use cases are examples of configuring Apache2 and WebSEAL to use CPACF to accelerate symmetric encryption and decryption of messages when using certain high-strength ciphers.

There are several good references for information on IBM System z cryptographic hardware:

An IBM Redbooks® publication with a good explanation of configuring cryptographic features, cryptographic domains, and LPARS: *IBM eServer zSeries 990 (z990) Cryptography Implementation, SG24-7070*

An IBM Redbooks publication with Crypto updates for z9, and discussions of CPACF and the Crypto Express2 Feature: *z9-109 Crypto and TKI V5 Update, SG24-7123*

The IBM HTTP Server 'Must Gather' page has good information on debugging problems when using cryptographic hardware: http://publib.boulder.ibm.com/httpserv/ihsdiag/gather_certificate_doc.html

An IBM Redbooks paper that discusses configuring cryptographic hardware specifically for Linux on IBM System z: http://www.redbooks.ibm.com/redpapers/pdfs/redp4131.pdf

Making Cryptographic Hardware available to Linux on IBM System z

In the scenarios that follow, the Linux systems are running as guest operating systems under z/VM®. To access the cryptographic hardware on the machine (CEX2 features and CPACF), we have done the configuration steps that follow.

Step 1: Verify that CPACF is enabled and if not, enable it

The CPACF feature code (3863) is enabled on our machines. This can be verified from the Hardware Management Console. From the HMC, login to the support element for the CPC, and double-click the CPC icon. The Instance Information Tab has the status of the CPACF feature. The example below shows that the feature is enabled:

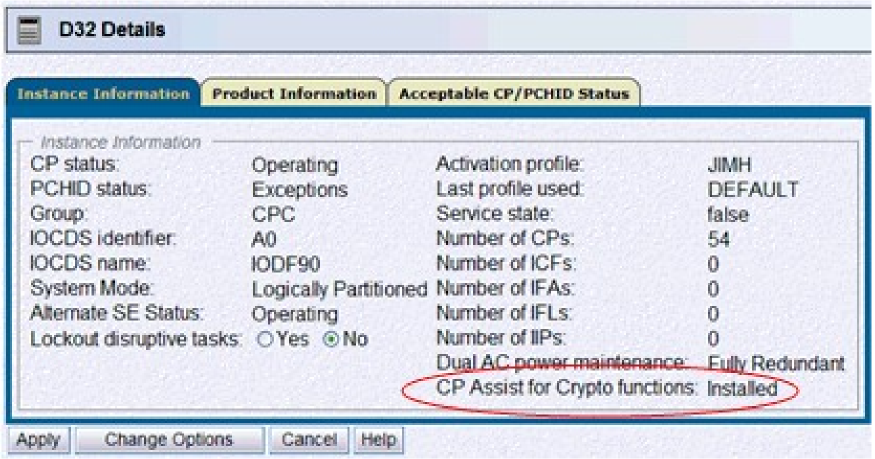

Instance Information

CP status:	Operating	Activation profile:	JIMH
PCHID status:	Exceptions	Last profile used:	DEFAULT
Group:	CPC	Service state:	false
IOCDS identifier:	A0	Number of CPs:	54
IOCDS name:	IODF90	Number of ICFs:	0
System Mode:	Logically Partitioned	Number of IFAs:	0
Alternate SE Status:	Operating	Number of IFLs:	0
Lockout disruptive tasks:	○ Yes ● No	Number of IIPs:	0
		Dual AC power maintenance:	Fully Redundant
		CP Assist for Crypto functions:	Installed

No other configuration is required to make the CPACF instructions available to a Linux guest running under z/VM.

Step 2: Make Crypto Express2 devices available to Linux guests

For the cryptographic devices supplied by a Crypto Express2 Feature, several steps are required to make the devices available to a Linux guest.

The first step is to configure the mode of each of the two cards contained in a Crytpo Express2 feature. Each card can be configured as a co-processor, or an accelerator.

In co-processor (CEX2C) mode, the card can be used for secure-key operations: by storing keys in the card itself so that the keys used for cryptographic operations never appear unencrypted outside the boundary of the card.

In accelerator (CEX2A) mode, the card can only be used to accelerate clear-key RSA operations.

Note that, a co-processor in CEX2C mode can also accelerate clear-key RSA operations, but, processing power is divided between acceleration and key-management. In CEX2A mode, all the processing power of the card is devoted to acceleration, providing better throughput for clear-key RSA acceleration. Since our scenarios will use the Crypto Express2 for clear-key RSA acceleration, we will make CEX2A configured devices available to our Linux systems.

From the Hardware Management Console, when logged into the CPC support element, the Cryptographic Management panel shows the status of the currently installed Crypto Express2 features:

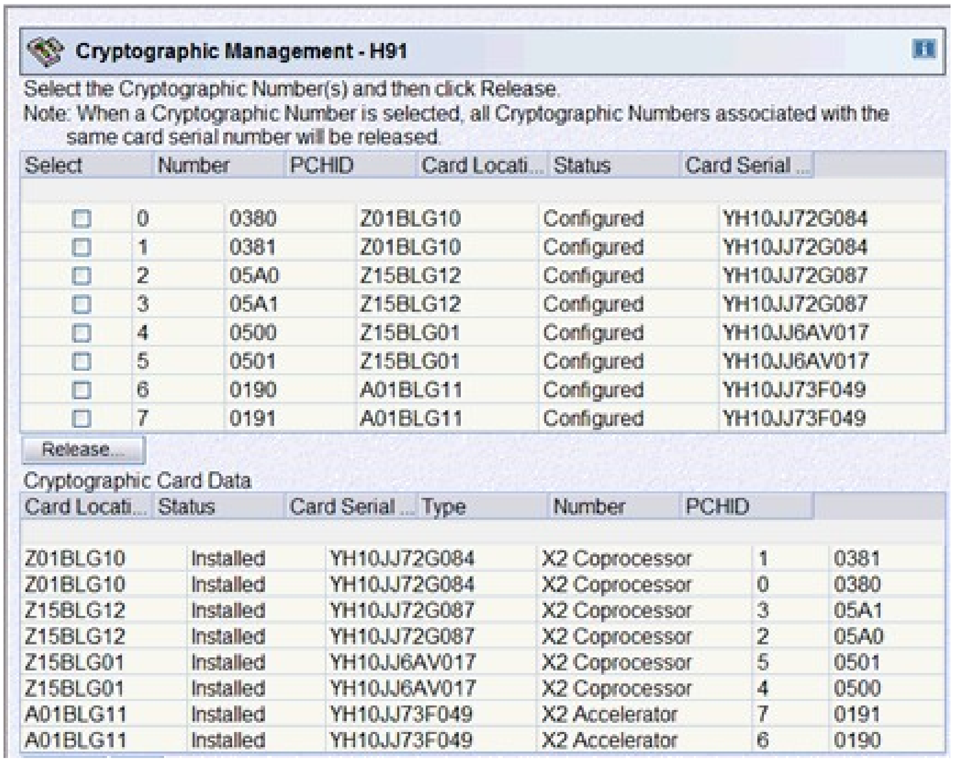

Cryptographic Management - H91

Select the Cryptographic Number(s) and then click Release.
Note: When a Cryptographic Number is selected, all Cryptographic Numbers associated with the same card serial number will be released.

Select	Number	PCHID	Card Locati...	Status	Card Serial ...
□	0	0380	Z01BLG10	Configured	YH10JJ72G084
□	1	0381	Z01BLG10	Configured	YH10JJ72G084
□	2	05A0	Z15BLG12	Configured	YH10JJ72G087
□	3	05A1	Z15BLG12	Configured	YH10JJ72G087
□	4	0500	Z15BLG01	Configured	YH10JJ6AV017
□	5	0501	Z15BLG01	Configured	YH10JJ6AV017
□	6	0190	A01BLG11	Configured	YH10JJ73F049
□	7	0191	A01BLG11	Configured	YH10JJ73F049

Cryptographic Card Data

Card Locati...	Status	Card Serial ...	Type	Number	PCHID
Z01BLG10	Installed	YH10JJ72G084	X2 Coprocessor	1	0381
Z01BLG10	Installed	YH10JJ72G084	X2 Coprocessor	0	0380
Z15BLG12	Installed	YH10JJ72G087	X2 Coprocessor	3	05A1
Z15BLG12	Installed	YH10JJ72G087	X2 Coprocessor	2	05A0
Z15BLG01	Installed	YH10JJ6AV017	X2 Coprocessor	5	0501
Z15BLG01	Installed	YH10JJ6AV017	X2 Coprocessor	4	0500
A01BLG11	Installed	YH10JJ73F049	X2 Accelerator	7	0191
A01BLG11	Installed	YH10JJ73F049	X2 Accelerator	6	0190

In this example, the Crypto Express2 feature with Serial YH10JJ73F049 has both of its co-processors configured in accelerator mode. Note also that each installed cryptographic device is assigned a number, referred to as the Adjunct Processor (AP) number, and, a PCHID number.

To make an Adjunct Processor available to an LPAR, the crypto tab in the activation profile for the LPAR image is used. Each AP can be assigned to the Candidate List or Online List for the LPAR image:

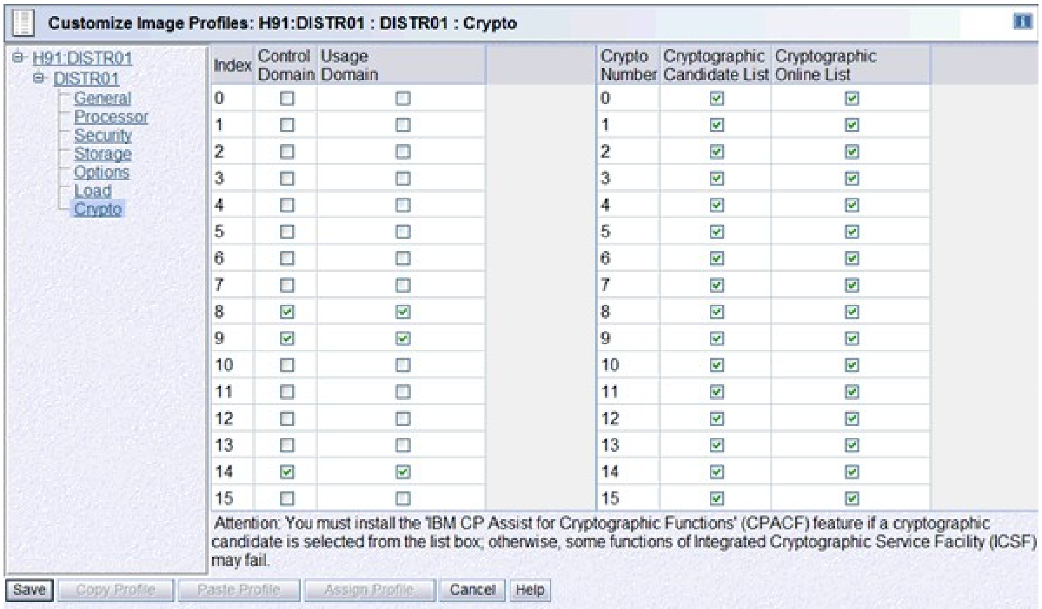

Customize Image Profiles: H91:DISTR01 : DISTR01 : Crypto

Index	Control Domain	Usage Domain	Crypto Number	Cryptographic Candidate List	Cryptographic Online List
0	□	□	0	☑	☑
1	□	□	1	☑	☑
2	□	□	2	☑	☑
3	□	□	3	☑	☑
4	□	□	4	☑	☑
5	□	□	5	☑	☑
6	□	□	6	☑	☑
7	□	□	7	☑	☑
8	☑	☑	8	☑	☑
9	☑	☑	9	☑	☑
10	□	□	10	☑	☑
11	□	□	11	☑	☑
12	□	□	12	☑	☑
13	□	□	13	☑	☑
14	☑	☑	14	☑	☑
15	□	□	15	☑	☑

Attention: You must install the 'IBM CP Assist for Cryptographic Functions' (CPACF) feature if a cryptographic candidate is selected from the list box; otherwise, some functions of Integrated Cryptographic Service Facility (ICSF) may fail.

In addition, the set of Usage Domains and Control Domains for those AP's assigned to the LPAR, are also assigned. An IBM Redbooks publication with a good discussion of Cryptographic Domains, AP's and LPARs is: *IBM eServer zSeries 990 (z990) Cryptography Implementation, SG24-7070*

In summary:

- Each AP can have up to 16 'usage domains' assigned
- Each usage domain has a separate set of master key registers (for secure-key cryptography)
- Each domain is associated with a separate 'AP queue'
- The AP queues reside in HSA – providing access to an AP from a CP
- AP Numbers are assigned to a 'candidate list' or 'online list' in an LPAR activation profile
- Each LPAR is assigned 'usage domains', which apply to all of the AP's configured to that LPAR
- An AP can be shared among 16 LPARs
- A usage domain / AP combination must be unique among active LPARs
- LPAR with Control Domain assignment allows key management of domains from the LPAR (using a TKE workstation)

The following picture illustrates the relationship of AP's, usage domains and AP queues when assigned to LPARs (LPAR1 to LPARx):

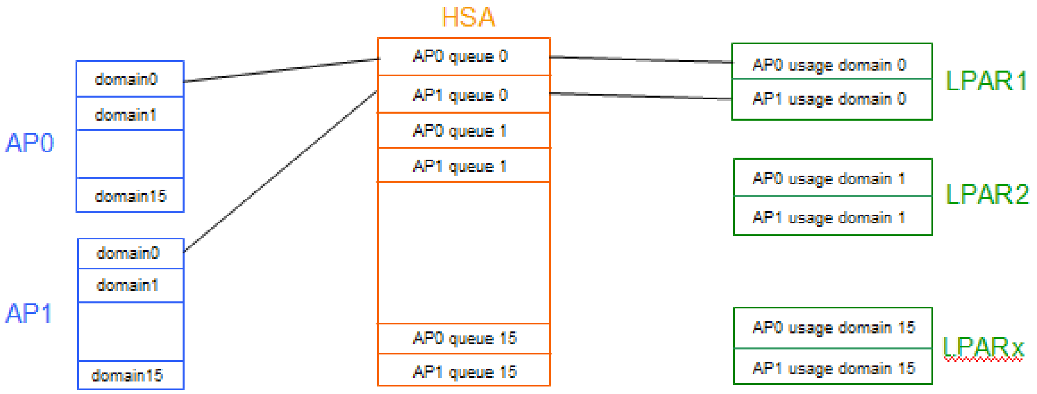

After activation of an LPAR with an appropriate profile, the AP's in the 'Online list' should be Online & Operating. From the HMC, when logged onto the CPC Support Element, from the

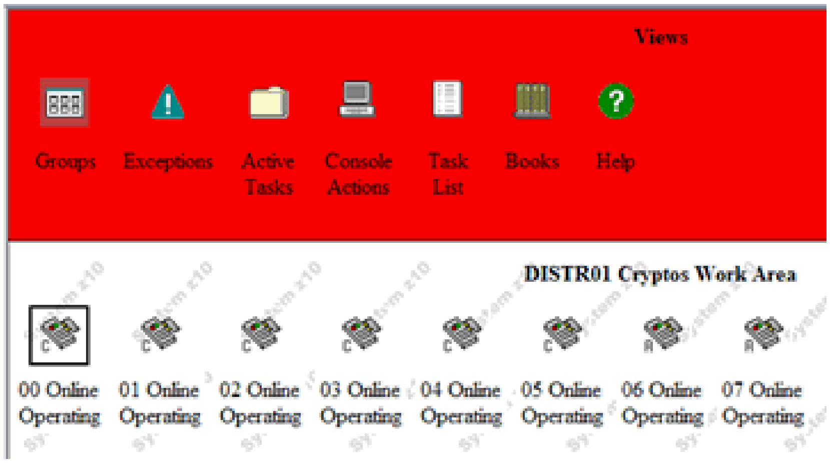

Images Work Area, select the LPAR's image icon, right-click, and select 'Cryptos' to view the status of the AP's assigned to the LPAR:
We see that AP's 00 through 05 are online, and running in co-processor mode, while AP's 06 and 07 are online and running in accelerator mode.

Crypto AP's online to z/VM can be accessed by a Linux guest if the directory statement includes:

```
CRYPTO APVIRT
```
or
```
CRYPTO DOMAIN <domain(s)> APDED <AP #>
```

For example:
```
CRYPTO DOMAIN 5 APDED 1
```

Unless there is a reason to dedicate domains (for secure-key cryptography, for example), it is preferable to use CRYPTO APVIRT and not CRYPTO APDED. With CRYPTO APVIRT:

- The Linux guest gets one random virtual queue on one random virtual AP.
- If there are multiple AP types available to z/VM, z/VM chooses the 'best' AP type for acceleration to show to the Linux guest. The prioritized list is: CEX2A, PCICA, CEX2C, PCIXCC.
- When a type is selected, z/VM will route requests from the guest to how ever many queues/cards of that type are available

From the z/VM OPERATOR, the 'QUERY CRYPTO APQS' command can be used to display the status of the AP's. For example, with AP's 00 through 05 online in co-processor mode and AP's 06 through 07 online in accelerator mode to the 'DISTRO1' LPAR as shown in the Cryptos Work Area panel example above, the 'Q CRYPTO APQS' command on the z/VM system running in the DISTRO1 LPAR shows:

```
q crypto apqs
AP 00 CEX2C Queue 08 is superseded by CEX2A
AP 00 CEX2C Queue 09 is superseded by CEX2A
AP 00 CEX2C Queue 14 is reserved for dedicated use
AP 01 CEX2C Queue 08 is superseded by CEX2A
AP 01 CEX2C Queue 09 is superseded by CEX2A
AP 01 CEX2C Queue 14 is reserved for dedicated use
AP 02 CEX2C Queue 08 is superseded by CEX2A
AP 02 CEX2C Queue 09 is superseded by CEX2A
AP 02 CEX2C Queue 14 is reserved for dedicated use
AP 03 CEX2C Queue 08 is superseded by CEX2A
AP 03 CEX2C Queue 09 is superseded by CEX2A
AP 03 CEX2C Queue 14 is reserved for dedicated use
AP 04 CEX2C Queue 08 is superseded by CEX2A
AP 04 CEX2C Queue 09 is superseded by CEX2A
AP 04 CEX2C Queue 14 is reserved for dedicated use
AP 05 CEX2C Queue 08 is superseded by CEX2A
AP 05 CEX2C Queue 09 is superseded by CEX2A
AP 05 CEX2C Queue 14 is reserved for dedicated use
AP 06 CEX2A Queue 08 is installed
AP 06 CEX2A Queue 09 is installed
AP 06 CEX2A Queue 14 is installed
AP 07 CEX2A Queue 08 is installed
AP 07 CEX2A Queue 09 is installed
AP 07 CEX2A Queue 14 is installed
```

We see that z/VM will use the two AP's configured as CEX2A's for Linux guests with 'CRYPTO APVIRT' statements in their directory entry. Each Linux guest will see one virtual CEX2A device and one virtual queue. In each of the following Use case scenarios, the directory entry for the z/VM userid of the Linux system contains 'CRYPTO APVIRT'.

The following picture illustrates how the Web Servers used in our scenarios access the IBM System z cryptographic hardware. In the case of the CPACF instructions, the instructions are issued by the IBM Crypto Adapter library (libica). In the case of the Crypto Express , the z90crypt driver sends requests to 'virtual' AP's, which are routed by z/VM to one of the real,

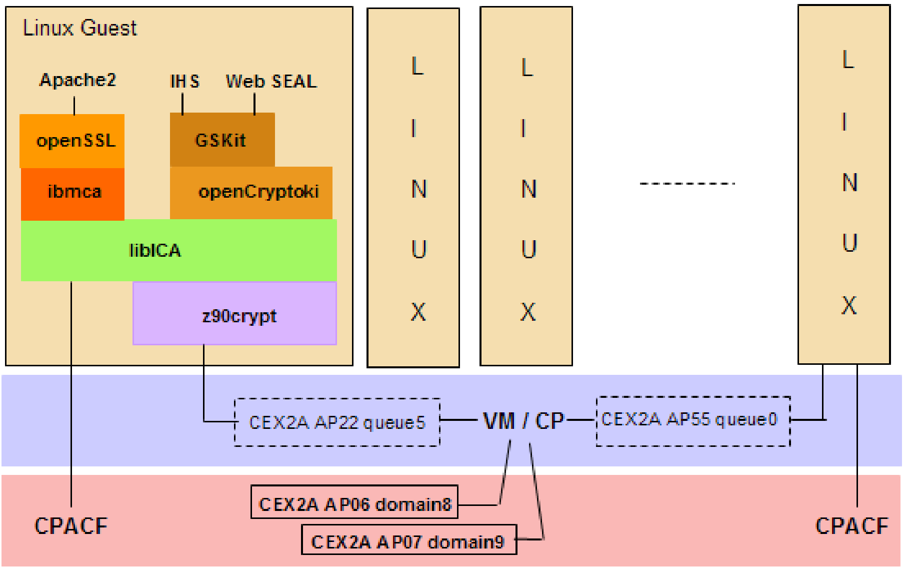

online, AP's:

IBM Systems and Technology Group
Route 100
Somers, New York 10589
U.S.A.
Produced in the United States of America,
07/2009

All statements regarding IBM’s future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

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.