



IBM Americas, ATS, Washington Systems Center

# Database Encryption

## Share 11488

### Anaheim, CA August 2012

Greg Boyd ([boydg@us.ibm.com](mailto:boydg@us.ibm.com))

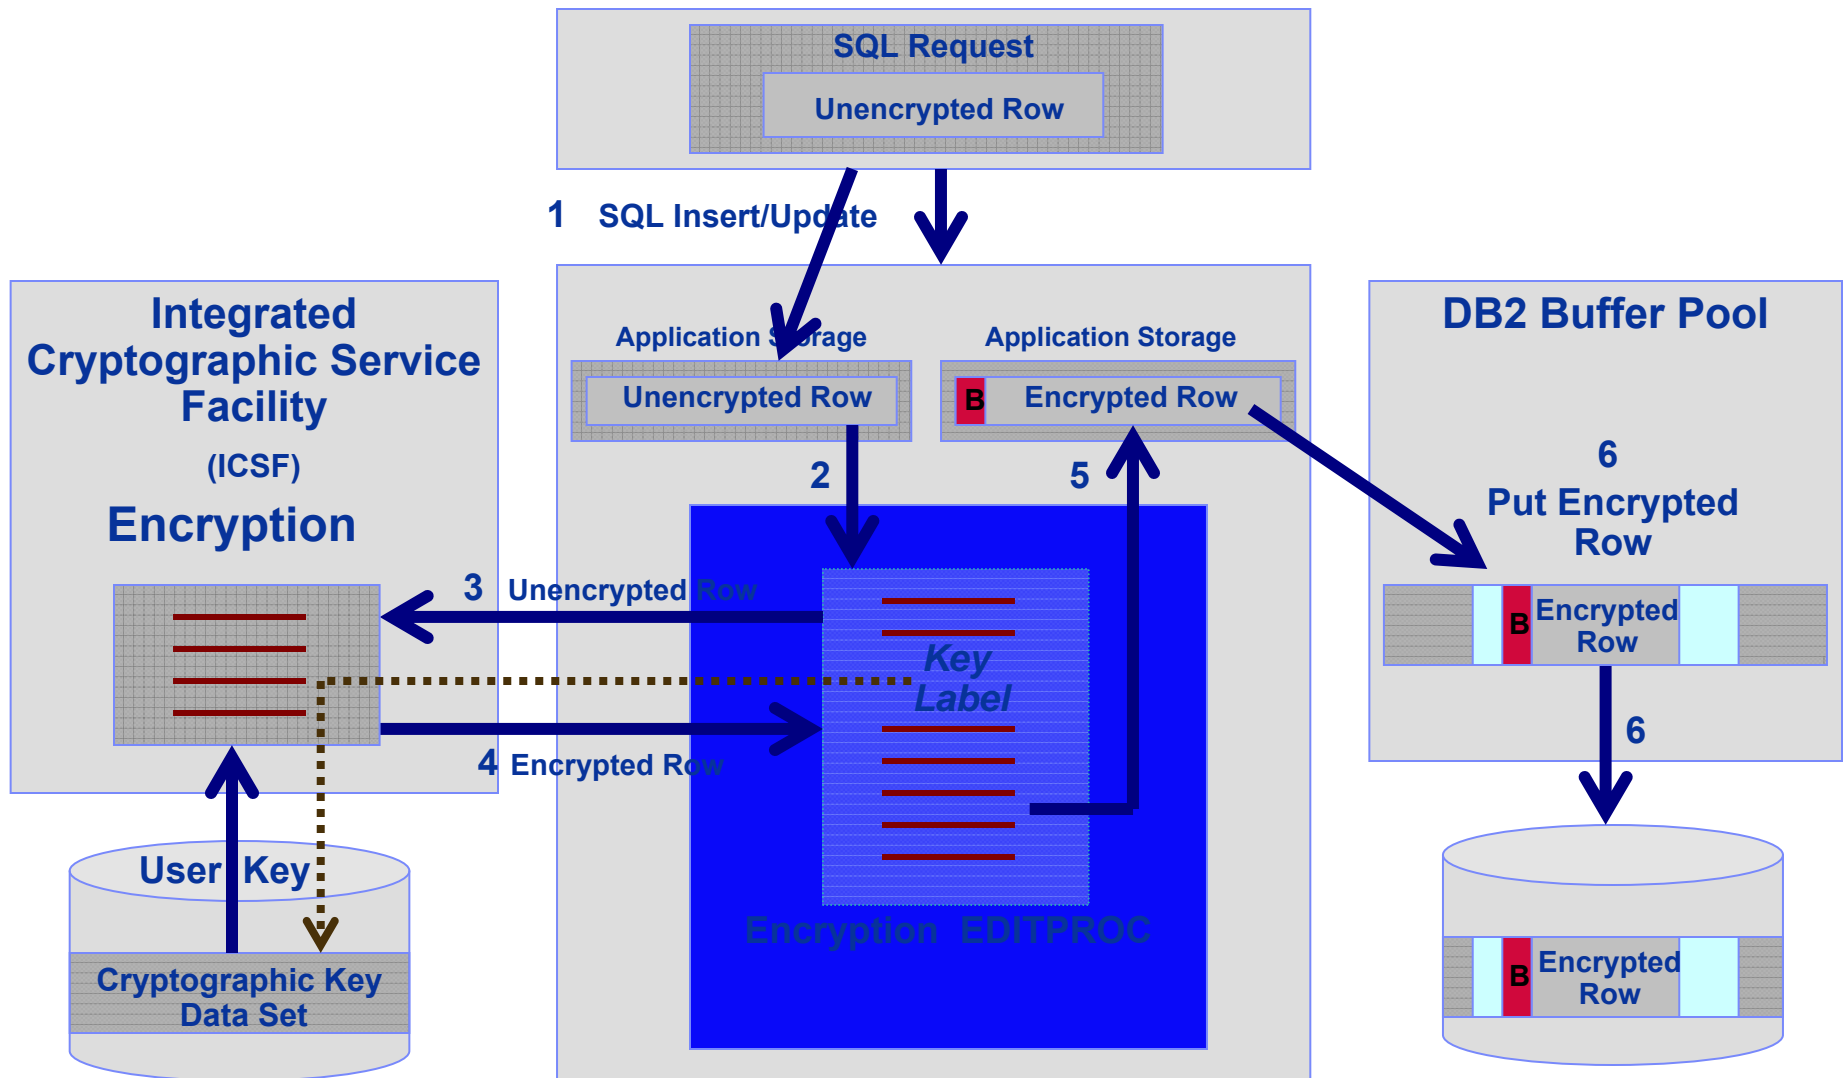
# Database Encryption

- **DB2 Built-In Functions**
- **IMS Data Encryption Tool for IMS & DB2 Databases (5799-P03)**
- **Other database encryption support**
  - Encrypting Tape
  - Encrypting DASD
  - Network encryption

## How does the Data Encryption Tool do encryption?

- **Via an EDITPROC, for every row processed by any SQL Utility for DB2 or IMS**
  - Encrypted row same length as clear row
  - No application changes required
  - One key per table or segment specified in the EDITPROC
  - Can use Clear Key, Secure Key or Protected Key
    - Protected key requires HCR7770 or later and CEX3

## DB2 Data Encryption Flow – Insert / Update



## How do the DB2 Built-In Functions do encryption?

- **Within the application, for every field that contains encrypted data**
  - ex. `encrypt(data,'password for encryption',hint)`**
  - 'Password for Encryption' is hashed to generate a unique key
  - Hint can be used as a prompt for remembering the key
  - Encrypted field must be defined as VARCHAR (since it will contain binary data once its encrypted) and
  - The encrypted field will be longer (next multiple of 8 bytes + 24 bytes of MetaData + 32 bytes for optional hint field)

## Cryptographic Keys

- **Data Encryption Tool**

- Clear Key or Secure Key or Protected Key
- Key must be stored in the CKDS
- When the table with an EDITPROC is in use, the key is available in the DB2 address space

- **DB2 BIF**

- Clear key only (it's calculated from the password for encryption in software) – so it's available in the DB2 address space
- Keys are not stored in a dataset, but the password for encryption is stored in the table

## Changing Cryptographic Keys

- **Data Encryption Tool**

- Unload, change EDITPROC to reference new key, reload
- Unload, change current key, DB2 restart, reload

- **DB2 BIF**

- Under application control

# Encryption and Indexes

## ■ Data Encryption Tool

- EDITPROC encrypts the entire row, so the data is encrypted, but the index is not
  - Bad for security, good for performance

INDEX	SSN NAME ADDRESS
223491398	F{(œ(•´ú □ — GÿP# □ ¥†%ojliÑÆ

## ■ DB2 BIF

- Application encrypts the field, if that field is an index, then the index is encrypted
  - Good for security, bad for performance

INDEX	SSN NAME ADDRESS
F{(œ(•´ú	F{(œ(•´ú □ — GÿP# □ ¥†%ojliÑÆ



## Data Encryption Tool – Hardware Requirements

### ■ Clear Key

–z196/z114/z10/z9 CPACF (& PCIXCC z9/z890/z990 or CEXnC for CKDS\*)

### ■ Secure Key

–z890/z990 Requires a PCIXCC or CEX2

–z9 Requires a CEX2C

–z10 Requires a CEX2C or CEX3C

–z196/z114 Requires a CEX3C

### ■ Protected Key

–z10/z196/z114 Requires a CEX3C\*\*

\*Prior to HCR7750, a CEXnC is required to create and use a CKDS, beginning with HCR7751 ICSF supports a clear key only CKDS

\*\*Protected Key support requires HCR7770 or higher

## DB2 BIFs - Hardware Requirements

- **z196/z114/z10/z9/z990/z890 (CPACF)**
  - Uses MSA instructions, not the ICSF APIs, but ICSF must be started to provide hashing support
  - TDES only

## Side-by-side Comparison

	Column (DB2 Built-In Functions)	Row/Table (IBM Encryption Tool for IMS and DB2)
<b>DB2 Support</b>	<ul style="list-style-type: none"> <li>▪ V8, V9, V10</li> <li>▪ Data in indexes is encrypted</li> <li>▪ Does not work w/DB2 Load Utility</li> <li>▪ Data type of encrypted columns must be FOR BIT DATA</li> </ul>	<ul style="list-style-type: none"> <li>▪ V7.x, V8.x, V9.x, v10.x</li> <li>▪ DB2 index data is not encrypted.</li> <li>▪ Works with all DB2 utilities</li> </ul>
<b>Application Change Required</b>	<ul style="list-style-type: none"> <li>▪ Application must change to invoke the BIFs for the columns that will be encrypted</li> </ul>	<ul style="list-style-type: none"> <li>▪ No application change, but each table will need to be recreated with an EDITPROC</li> </ul>
<b>Transaction Processing Overhead</b>	<ul style="list-style-type: none"> <li>▪ The cost overhead depends on hardware, DB2 and application access</li> </ul>	<ul style="list-style-type: none"> <li>▪ High overhead due to the amount of data encryptions</li> </ul>
<b>Key Management</b>	<ul style="list-style-type: none"> <li>▪ Application has responsibility for the encryption key</li> </ul>	<ul style="list-style-type: none"> <li>▪ Keys are managed by and accessed through ICSF</li> </ul>
<b>Pre-Reqs</b>	<ul style="list-style-type: none"> <li>▪ ICSF must be active</li> <li>▪ CPACF hardware</li> </ul>	<ul style="list-style-type: none"> <li>▪ ICSF must be active</li> <li>▪ Secure PCI card, unless running HCR7751 or later and clear key only CKDS</li> </ul>

## Enabling Protected Key


- **Install HCR7770 or later**
  - CSFINIT replaces CSFMMAIN
- **Install Crypto Express3 on z10 with Driver 79 or on z196/z114**
  - With master keys loaded
- **Install RACF (OA29193) and SAF (OA29194) APARs**
- **Create secure keys which will be used as protected keys**
- **Create/update RACF profiles for the keys, with SYMCPACFWRAP(YES)**

# z196 Crypto Performance


## From the Crypto Performance Whitepapers

<http://www.ibm.com/systems/z/advantages/security/z10cryptography.html>

### ■ AES Encryption



### ■ TDES Encryption



# Secure Key SQL Performance Results

```
JOBNAME:                Current Thread Detail                DATE: 06/26/08
DB2 V8 :                 TIME: 11:42:17
COMMAND:                 CYCLE: MMSS
```

```
CONN ID :                PLAN      :                CURRENT STATE: INAPP
CORR ID :                AUTH ID  :                THREAD START  : 11:32:49.4763
LOCATION:                 SQLID   :                CONN TYPE    : CALL ATTACH
RQST LOC:               LUWID   :                }
PKG LOC :                ACCT TKN:
PKG NAME: FDB2V600.SQLPCRTN.18386E190EC573D6
```

```
+----- Timings -----+
ELAPSED: 5:42.97 DB2 ELA: 0:02.96
TOT CPU: 0:00.14 DB2 CPU: 0:00.13
I/O WT : 0:00.00- LOCK WT: 0:00.00
SORT   : 0:00.00- TOT WT : 0:00.30
NESTED : 0:00.00
```

```
+----- Event Counts -----+
WAIT      :      27   PACKAGES:      2
IFI       :      0   PARA GRP:      0
RMT CALL:      0   PARA CPU:      0
SORT     :      1   PARA MBR:      0
SQL LOGR:      0   DS OPENS:      1
RID LIST:      0
```

```
+----- SQL Counts -----+
TOTAL    :      2054  PREPARES :      1
SELECT  :           1  OPEN CSR  :      8
FETCH   :      2031  INCR BIND:      0
COMMIT  :           2  SECURITY  :      0
DML     :           0  DDL      :      0
```

```
+----- Buffer Pool/Locking -----+
GETPAGE:      182  MX PG LK:      1
SYNC RD:      9   LOCKESCL:      0
PREFTCH:      4   SUSPENDS:      0
ASYN RD:      4   TIMEOUTS:      0
PGS/IO  :     14.0  DEADLOCK:      0
```

# Clear Key SQL Performance Results

DB2 V8 : DTVB  
 COMMAND:

TIME: 11:44:40  
 CYCLE: MMSS

CONN ID : PLAN : CURRENT STATE: INAPP  
 CORR ID : AUTH ID : THREAD START : 11:43:37.9172  
 LOCATION: SQLID :  
 RQST LOC: LUWID :  
 PKG LOC : ACCT TKN:  
 PKG NAME: FDB2V600.SQLPCRTN.18386E190EC573D6

+----- Timings -----+		+----- Event Counts -----+	
ELAPSED: 1:02.64	DB2 ELA: 0:00.36	WAIT : 13	PACKAGES: 2
TOT CPU: 0:00.03	DB2 CPU: 0:00.03	IFI : 0	PARA GRP: 0
I/O WT: 0:00.01	LOCK WT: 0:00.00	RMT CALL: 0	PARA CPU: 0
SORT : 0:00.00	TOT WT : 0:00.33	SORT : 1	PARA MBR: 0
NESTED : 0:00.00		SQL LOGR: 0	DS OPENS: 1
		RID LIST: 0	

+----- SQL Counts -----+		+----- Buffer Pool/Locking -----+	
TOTAL : 2054	PREPARES : 1	GETPAGE: 182	MX PG LK: 1
SELECT : 1	OPEN CSR : 8	SYNC RD: 3	LOCKESCL: 0
FETCH : 2031	INCR BIND: 0	PREFTCH: 4	SUSPENDS: 0
COMMIT: 2	SECURITY : 0	ASYN RD: 4	TIMEOUTS: 0
DML : 0	DDL : 0	PGS/IO : 26.0	DEADLOCK: 0

## Secure vs. Clear Key: Database Load Results

Database utility loads of 200,000 rows yielded the following results:

(In seconds)	Clear Key	Secure Key
CPU Time	2	8
Elapsed Time	18	259

As you can see from the LOAD and SQL examples, secure key is considerably more CPU intensive.



# Implementation—Example


Table xxx

Encrypted Tables xxDBA					Non Encrypted Tables xxNON				
Utility	Elapsed Time	CPU Time	Init Date	Init Time	Utility	Elapsed Time	CPU Time	Init Date	Init Time
Unload	00:01:37.86	00:01:46.02	Sept. 28	9:15 A	Unload	00:01:42.64	00:01:08.31	Sept. 28	9:15 A
Load	00:04:07.73	00:03:45.13	Sept. 28	11:30 A	Load	00:03:40.55	00:03:12.89	Sept. 28	11:30 A
REORG	00:19:56.46	00:03:33.44	Sept. 28	2:30 P	REORG	00:05:49.17	00:02:12.37	Sept. 28	2:30 P
Index Rebuild	00:03:50.03	00:01:32.04	Sept. 29	9:00 A	Index Rebuild	00:01:20.30	00:00:48.94	Sept. 29	9:00 A
Image Copy	00:07:05.19	00:00:08.10	Sept. 29	1:00 P	Image Copy	00:03:51.43	00:00:07.56	Sept. 29	1:00 P
Recover	00:07:05.19	00:00:08.10	Sept. 29	2:15 P	Recover	00:03:51.43	00:00:07.56	Sept. 29	2:15 P
DSNTIAUL	00:05:42.22	00:04:31.99	Sept. 30	9:30 A	DSNTIAUL	00:05:23.32	00:03:52.42	Sept. 30	9:30 A

Your mileage may vary.

## Disaster Recovery Considerations

- **The major requirement is that the appropriate crypto hardware be available at the DR site**
  - Clear Key / Secure Key / Protected Key
  - Key lengths
- **For the data encryption tool, master keys must be available at the DR site**



## Decisions, Decisions ...

- **Ownership (i.e. politics)**

- Data Administrator - Data Encryption Tool

- Sets up the EDITPROC and specifies the key to be used for the entire table

- Key must be defined to/managed by ICSF (stored in the CKDS)

- Application - DB2

- Application logic determines which key to use for each field/column

- Password is managed by the application

- **Security requirements**

- **Performance requirements**

- **Application/production support**

- **Space considerations**

- **Crypto hardware available**



## Data Encryption Tool – New Functions

- **Customizable UDF PM45364/UK72991**
  - No application changes
  - Minimally disruptive, columns encrypted in place
  - Indexes can be encrypted
  - All data types supported by UDFs can be encrypted
  
- **Customizable FIELDPROC PM55879/UK76423**
  - No application changes
  - Non-disruptive
  - Index can be encrypted
  - Same restrictions as EDITPROC

## IBM Tape Based Encryption


- **LTO4 and LTO5 - Open Systems**
- **TS1120, TS1130, TS1140 - Open Systems and Mainframe**
- **AES-256 bit encryption**
- **All files on the tape are protected using a single key**
  - Which is in turn encrypted using RSA (public/private key algorithms)
- **TKLM, Tivoli Key Lifecycle Manager or in a z/OS environment, ISKLM IBM Security Key Lifecycle Manager is required for DS8000 and recommended for Tapes**

# IBM DS8000 Disk Encryption - Characteristics

- **Customer data at rest is encrypted**
  - Data at rest = data on any disk or in any persistent memory
- **Customer data in flight is not encrypted**
  - Data in flight = on I/O interfaces or in dynamic memories (Cache, NVS)
    - If you can read/write to disk, you get access to clear-text data.
- **Uses Encrypting Disk**
  - Encryption hardware in disk (AES 128)
  - Runs at full data rate (146/300/450 GBs 15K RPM )- **No measurable performance impact**
- **Integrated with Tivoli Key Lifecycle Manager (TKLM) or IBM Security Key Lifecycle Manager (ISKLM)**
  - DS8000 automatically communicates with TKLM when configuring encryption group or at power on to obtain necessary encryption keys to access customer data
  - Each disk has an encryption key
    - Data is always encrypted on write and decrypted on read
    - Encryption key is wrapped with access credential and maintained within the disk
    - Access credential maintained by TKLM/ISKLM
    - Establishing a new encryption key causes cryptographic erasure
- **Key attack vectors prevented:**
  - Disk removed (repair, or stolen)
  - Box removed (retire, or stolen)

## zIIP Assisted IPsec (VPN) on z/OS

- Benefits of having secure channel end-point on z/OS
  - Security regulations compliance - No clear-text data on any network segments
  - End-to-end authentication of secure channel end-points
    - Both end-point authentication and message authentication
  - Key management and storage done on System z by z/OS
  - Compliance with end-to-end security regulations
- System z CPU cost is a concern
  - Encryption/decryption CPU cost can be a significant percentage of overall CPU cost for a given application
  - Especially the case for streaming workloads (file transfer type of workload)
- zIIP processors
  - Specialty processor on System z9 or later hardware
  - zIIPs priced lower than general purpose processors
  - No IBM software charges on zIIPs
- zIIP Assisted IPsec
  - Use zIIP processors for most IPsec encryption/decryption
  - Lower the cost of doing IPsec processing on z/OS



## Closing Thoughts

- **Encryption has a cost**
  - Crypto hardware more efficient with large blocks of data
- **Secure Key on a PCI Card – longer pathlength**
- **Clear Key exists in the DB2 Address Space, Protected Key and Secure Key are too, but they are stored encrypted under the Wrapping Key or Master Key**



## Data Encryption for DB2 - Reference Materials

- **SC18-9549 IBM Data Encryption Tool for IMS and DB2 Databases User Guide**
  - Includes an appendix on activating crypto on your hardware
- **ICSF Manuals**
  - SA22-7520 ICSF System Programmer's Guide
  - SA22-7521 ICSF Administrator's Guide
- **Redbooks**
  - DB2 UDB for z/OS Version 8 Performance Topics – SG24-6465
- **Articles**
  - IMS Newsletter article: “Encrypt your IMS and DB2 data on z/OS” - <ftp://ftp.software.ibm.com/software/data/ims/shelf/quarterly/fall2005.pdf>

## Session #11488 Feedback





# Trademarks

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

AlphaBlox*	GDPS*	RACF*	Tivoli*
APPN*	HiperSockets	Redbooks*	Tivoli Storage Manager
CICS*	HyperSwap	Resource Link	TotalStorage*
CICS/VSE*	IBM*	RETAIN*	VSE/ESA
Cool Blue	IBM eServer	REXX	VTAM*
DB2*	IBM logo*	RMF	WebSphere*
DFSMS	IMS	S/390*	zEnterprise
DFSMSHsm	Language Environment*	Scalable Architecture for Financial Reporting	xSeries*
DFSMSrmm	Lotus*	Sysplex Timer*	z9*
DirMaint	Large System Performance Reference™ (LSPRT™)	Systems Director Active Energy Manager	z10
DRDA*	Multiprise*	System/370	z10 BC
DS6000	MVS	System p*	z10 EC
DS8000	OMEGAMON*	System Storage	z/Architecture*
ECKD	Parallel Sysplex*	System x*	z/OS*
ESCON*	Performance Toolkit for VM	System z	z/VM*
FICON*	PowerPC*	System z9*	z/VSE
FlashCopy*	PR/SM	System z10	zSeries*
* Registered trademarks of IBM Corporation	Processor Resource/Systems Manager		

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.

# System z Social Media

- System z official Twitter handle:
  - [@ibm\\_system\\_z](#)
- Top Facebook pages related to System z:
  - [Systemz Mainframe](#)
  - [IBM System z on Campus](#)
  - [IBM Mainframe Professionals](#)
  - [Millennial Mainframer](#)
- Top LinkedIn Groups related to System z:
  - [Mainframe Experts Network](#)
  - [Mainframe](#)
  - [IBM Mainframe](#)
  - [System z Advocates](#)
  - [Cloud Mainframe Computing](#)
- YouTube
  - [IBM System z](#)



- **Leading Blogs related to System z:**
  - [Evangelizing Mainframe \(Destination z blog\)](#)
  - [Mainframe Performance Topics](#)
  - [Common Sense](#)
  - [Enterprise Class Innovation: System z perspectives](#)
  - [Mainframe](#)
  - [MainframeZone](#)
  - [Smarter Computing Blog](#)
  - [Millennial Mainframer](#)