

LUW-05

Recipes for Db2 Backups with
Encryption, Compression,
Hardware Acceleration and Other
Ingredients

Beck Tang, IBM

becktang@ca.ibm.com



PLEASE NOTE


- IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion.
- Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision.
- The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract.
- The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.
- Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including 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 results similar to those stated here.

Agenda

- ❖ Using the right bowls to mix the pastry
 - The backup process
- ❖ Choosing best ingredients
 - Compression, encryption and other options
- ❖ Monitoring the baking process
 - db2pd and list utilities
- ❖ Tasting the meal
 - Check/validate the backup image
- ❖ Get feedback from the guests
 - Reading the history file



Backup Theory - Process and Involved EDUs



db2 "backup db ABC online to /backup_target_1, /backup_target_2 with 6 buffers parallelism 3"

Note: The backup buffers are coming from the database configuration, Utility memory heap (UTIL_HEAP). For reference, SAP recommendation for DB2 11.5 for util_heap_sz is AUTOMATIC with an initial value of 50000 X 4KB pages.

Db2 Backup – Using Compression

- Advantages:
 - The size of the backup images will be smaller than without compression
 - This has impact on the time to take backups, the time to transfer the image and the size it allocates in the backup infrastructure
- Disadvantages:
 - Takes more resources to compress the data
- Available compression methods:
 - Common Db2 Compression
 - Software based
 - NX842 on Power
 - Using special accelerator units on CPUs of Power 7+ and later
 - ZLib Compression
 - Either software or using additional accelerators in x86, Power 9 and z CPUs



Backup Theory – HW Acceleration – NX842

```
# amepat

Command Invoked      : amepat
Date/Time of invocation : Tue Aug 22 16:36:21 CEST 2023
Total Monitored time  : NA
Total Samples Collected : NA
```

System Configuration:

```
-----
Partition Name      : db6p021004
Processor Implementation Mode : POWER7 Mode
Number Of Logical CPUs : 32
Processor Entitled Capacity : 8.00
Processor Max. Capacity : 8.00
True Memory         : 61.00 GB
SMT Threads         : 4
Shared Processor Mode : Disabled
Active Memory Sharing : Disabled
Active Memory Expansion : Disabled
```

System Resource Statistics:

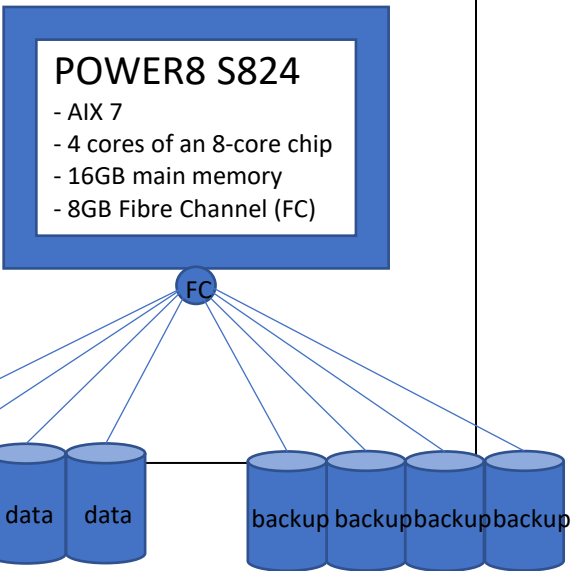
	Current
CPU Util (Phys. Processors)	0.09 [1%]
Virtual Memory Size (MB)	41752 [67%]
True Memory In-Use (MB)	59719 [96%]
Pinned Memory (MB)	9181 [15%]
File Cache Size (MB)	17905 [29%]
Available Memory (MB)	18812 [30%]

```
1> db2pd -osinfo
```

Operating System Information:

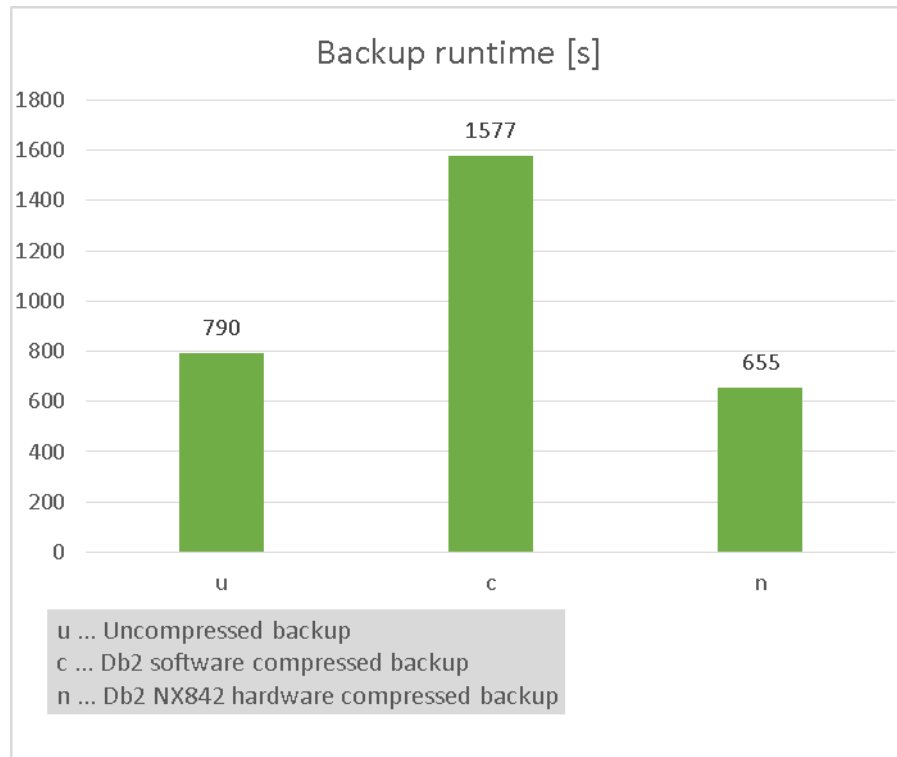
```
OSName:      AIX
NodeName:    db6p021004
Version:     7
Release:     2
Machine:     00CB85FF4C00
AME       : Disabled
AMS       : Disabled
NX842       : 109:Function not implemented
NXZLIB       : n/a
```

Active Memory Expansion Capable	True
Hardware-Accelerated Active Memory Expansion Capable	True




Backup Theory – Compression-Type-Efficiency

- Runtime



- CPU Usage



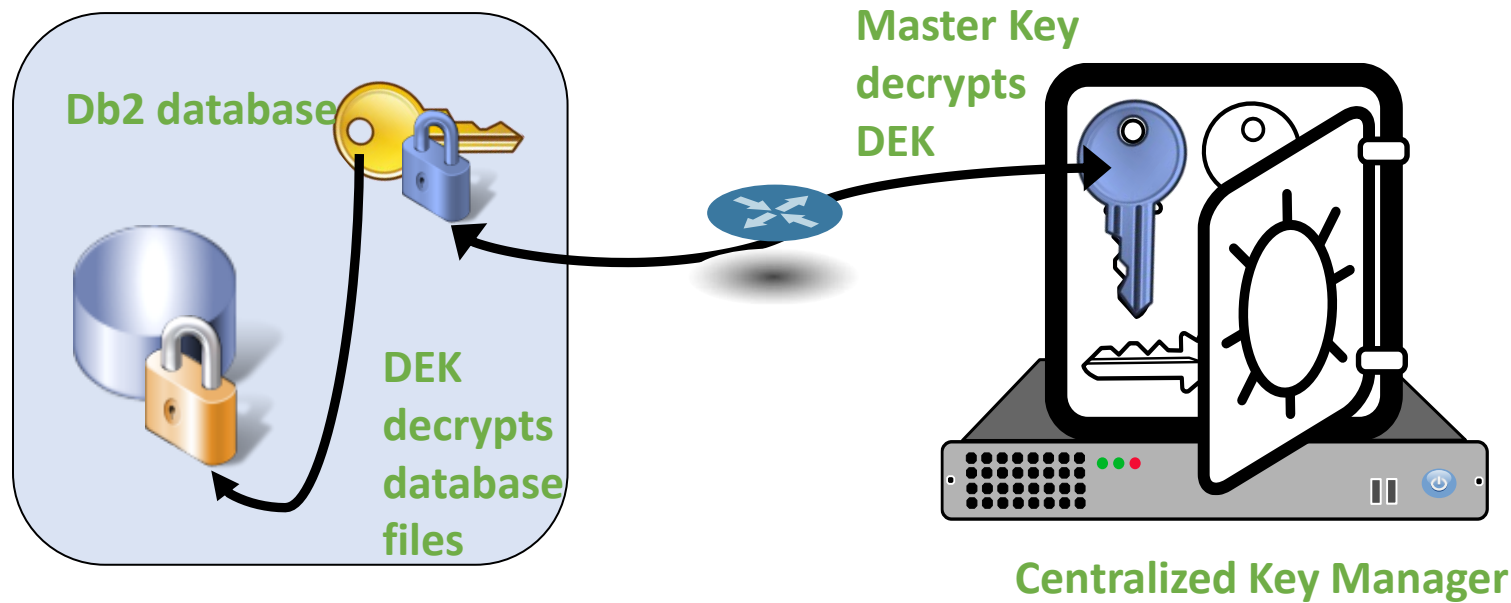
HW-Compression: fast and efficient – use if available

Db2 Backup – Using Encryption



- Advantages:
 - **Data Security:** Encryption helps protect sensitive data in backups from unauthorized access. Even if backup files are stolen or compromised, the data remains encrypted and inaccessible without the encryption key.
 - **Compliance:** Encryption is often required to meet regulatory compliance standards, such as **HIPAA** (Health Insurance Portability and Accountability Act), or **PCI DSS** (Payment Card Industry Data Security Standard), **GIS** (Global Information Security) which mandate the protection of sensitive data, especially for the finance and health sectors.
 - **Safe Data Transfer:** If you need to transfer backup files to a different location or store them in a cloud service, encryption adds an extra layer of security during transit.
- Disadvantages:
 - **Performance Overhead:** Obviously costs additional resources either on general purpose CPU or on hardware accelerators – up to 10 %.
 - **Complexity:** Requires setting up encryption algorithms, managing keys, and ensuring compatibility with your backup and recovery processes.

Backup Theory - Encryption



The master key is accessed whenever a DEK requires decryption

- db2start open keystore
- db start (activate, first connect)
- backup /restore
- ...

DEK is generated by Db2 when
Database is created
Restore into new database
Backup is created

Master Key is created
Manually within the key manager
Automatically by Db2

Backup Theory – Encryption/Compression Options

- COMPRLIB/ENCRLIB options

Operating System	Compression	Encryption	Both
Windows	db2compr.dll	db2encr.dll	db2compr_encr.dll
Linux	libdb2compr.so	libdb2encr.so	libdb2compr_encr.so
AIX	libdb2compr.a	libdb2encr.a	libdb2compr_encr.a

ZLIB without HW acceleration: libdb2zcompr_encr.dll and libdb2zcompr_encr.so

- COMPROPTS/ENCROPTS options

Option	Purpose	Values
Cipher	Type of encryption algorithm to use	AES, 3DES
Key Length	Length of the encryption key	AES: 128, 192, 256 3DES: 168
Master Key Label	Optional name of the Master Key Label used to encrypt the database key	String
Mode	Optional – Cipher Block Chaining	CBC

POWER 7+

AIX NX842 accelerator

- libdb2nx842.a

- libdb2nx842_encr.a

POWER 9

AIX NXZLIB accelerator

- libdb2zcompr.a

- libdb2zcompr_encr.a

z15+


Linux NXU accelerator

- libdb2zcompr.so

- libdb2zcompr_encr.so

Since Db2 11.5.7+

Backup in Action – Test Environment



- IBM XIV storage system
- Db2: 11.5.6, database size: 270GB
- Database overall compression ratio: 52%
- Number of tablespaces: 48
- Automatic backup settings:
 - Parallelization: 30, buffer: 60 (unencrypted “default” backup)
 - Parallelization: 15, buffer: 30 (encrypted backups)

Backup in Action – Encryption

[Malloc heap contention may cause performance degradation when using DB2 on AIX with specific features \(ibm.com\)](#) -

Optimized AIX malloc setting for instance owner

- Change the AIX malloc settings

```
export MALLOCOPTIONS=buckets,multiheap:4
export
MALLOCBUCKETS=number_of_buckets:128,bucket
_sizing_factor:64,blocks_per_bucket:1024
```
- Unlmt the size of the data area of the instance owner in /etc/security/limits:

```
db2vfb:
data = -1
```
- Update the Db2 registry and restart the instance:

```
db2set DB2ENVLIST="MALLOCOPTIONS
MALLOCBUCKETS"
db2stop
db2start
```

Reduce parallelization and buffer

- Backup Command

```
db2 "backup db <dbname> to <backup paths>
with 8 buffers parallelism 4 encrypt
encrlib libdb2encr.a
```
- Manually adapted backup options

Parallelism: Defines the number of tablespaces backed up in parallel

Buffer: At least two times the parallelism value

Backup in Action – Create Encrypted DB

Install Keystore using IBM GSKit

- Create keystore directory
 - `su - instance owner`
 - `mkdir <path>/keystore`
- Create keystore
 - `gsk8capicmd - keydb -create -db <path>/keystore/<keyfile>.p12 -pw <strong password> -string -type pkcs12 -stash`



Encrypt Database

- Configure instance
 - `db2 update db cfg using keystore type pkcs12 Keystore_location <path>/keystore/<keyfile>.p12`
 - `db2stop`
 - `db2start`
 - `db2 backup db <dbname> to <backup paths>`
 - `db2 drop db <dbname>`
 - `db2 restore db <dbname> from <backup paths> taken at <backup timestamp> ENCRYPT`


You can even use HADR to encrypt the database as described here:

<https://blogs.sap.com/2022/05/08/encrypting-an-sap-system-on-a-db2-for-luw-database-reduce-downtime-by-exploiting-the-hadr-feature/>

Backup in Action – Encryption




Backup in Action – Software Compression & Encryption




Malloc =yes

Backup in Action – AIX NX842 Compression & Encryption




Malloc =yes

Backup in Action – Compression & Encryption



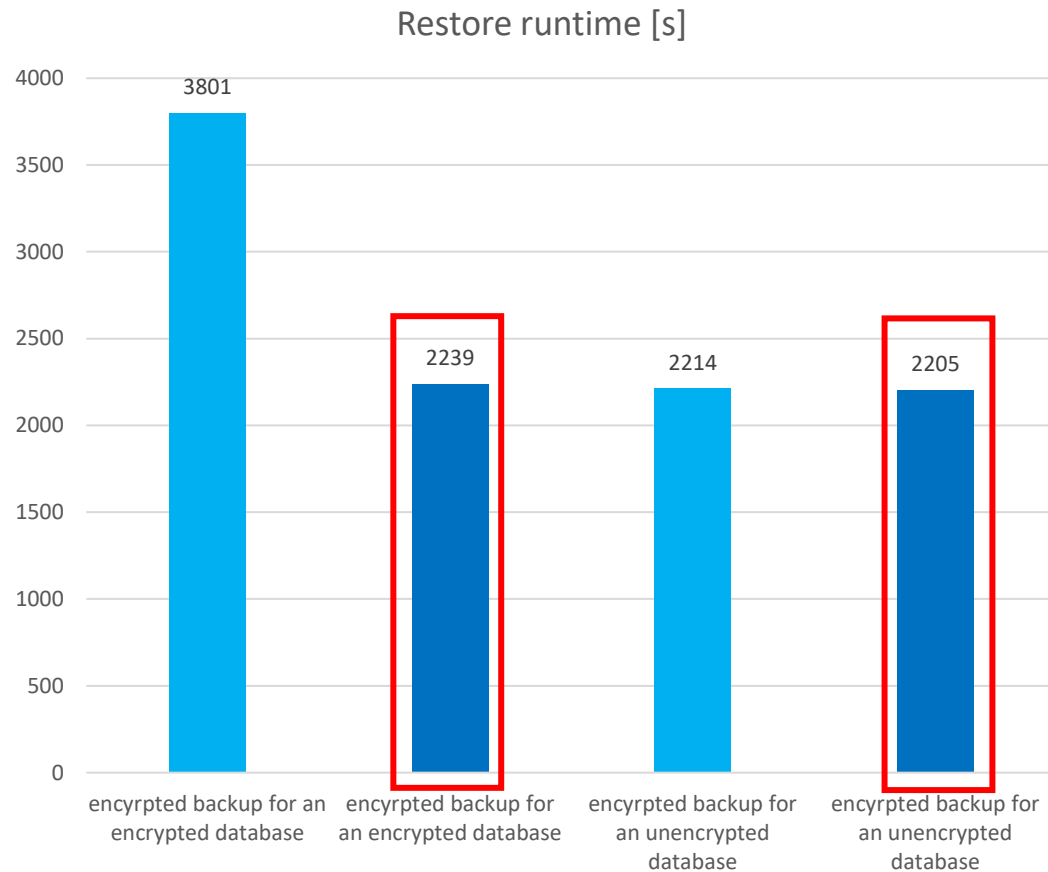
Malloc =yes

Backup in Action – Unencrypted Databases



Malloc =yes

Restore of Encrypted Backups



restoring into encrypted database:

- From encrypted or unencrypted backup
- With MALLOC workaround or not

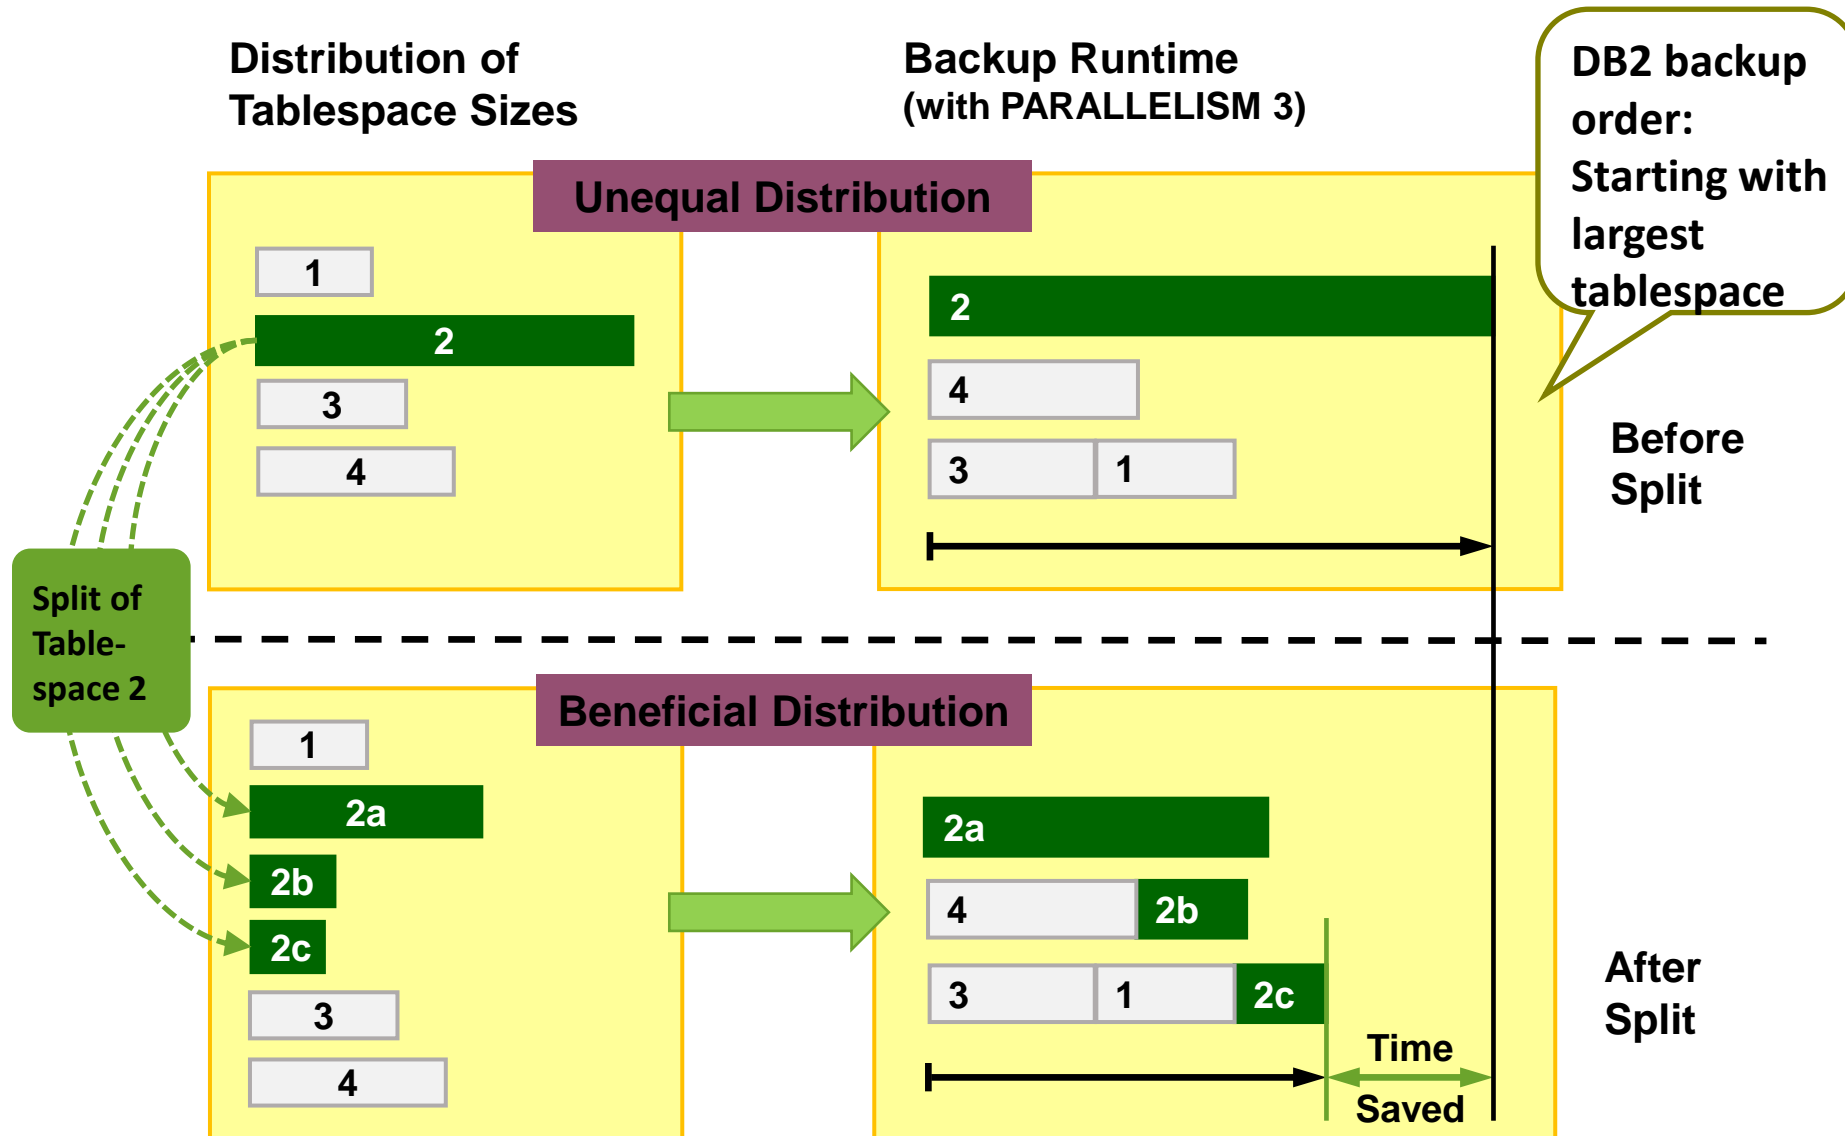
Malloc =yes

ZLib vs. Common Compression on LinuxX86

- Database size: 48.1 GB
- Test sequence:
 - db2 backup db jkl to /target compress
 - db2set DB2_BCKP_COMPRESSION=ZLIB
 - ...
 - db2 update db cfg for jkl using encr lib libdb2compr_encr.so
 - ...
 - db2 backup db jkl to /target compress
- Common Db2 Compression
 - Runtime: 474 seconds
 - Size: 21.1 GB (43% of orig size)
- ZLib Compression
 - Runtime: 305 seconds
 - Size: 12.75 GB (27% of orig size)
- Backup runtime reduced by 35%
- In general, Zlib compressed a little bit better compared to NX842; runtimes are about the same

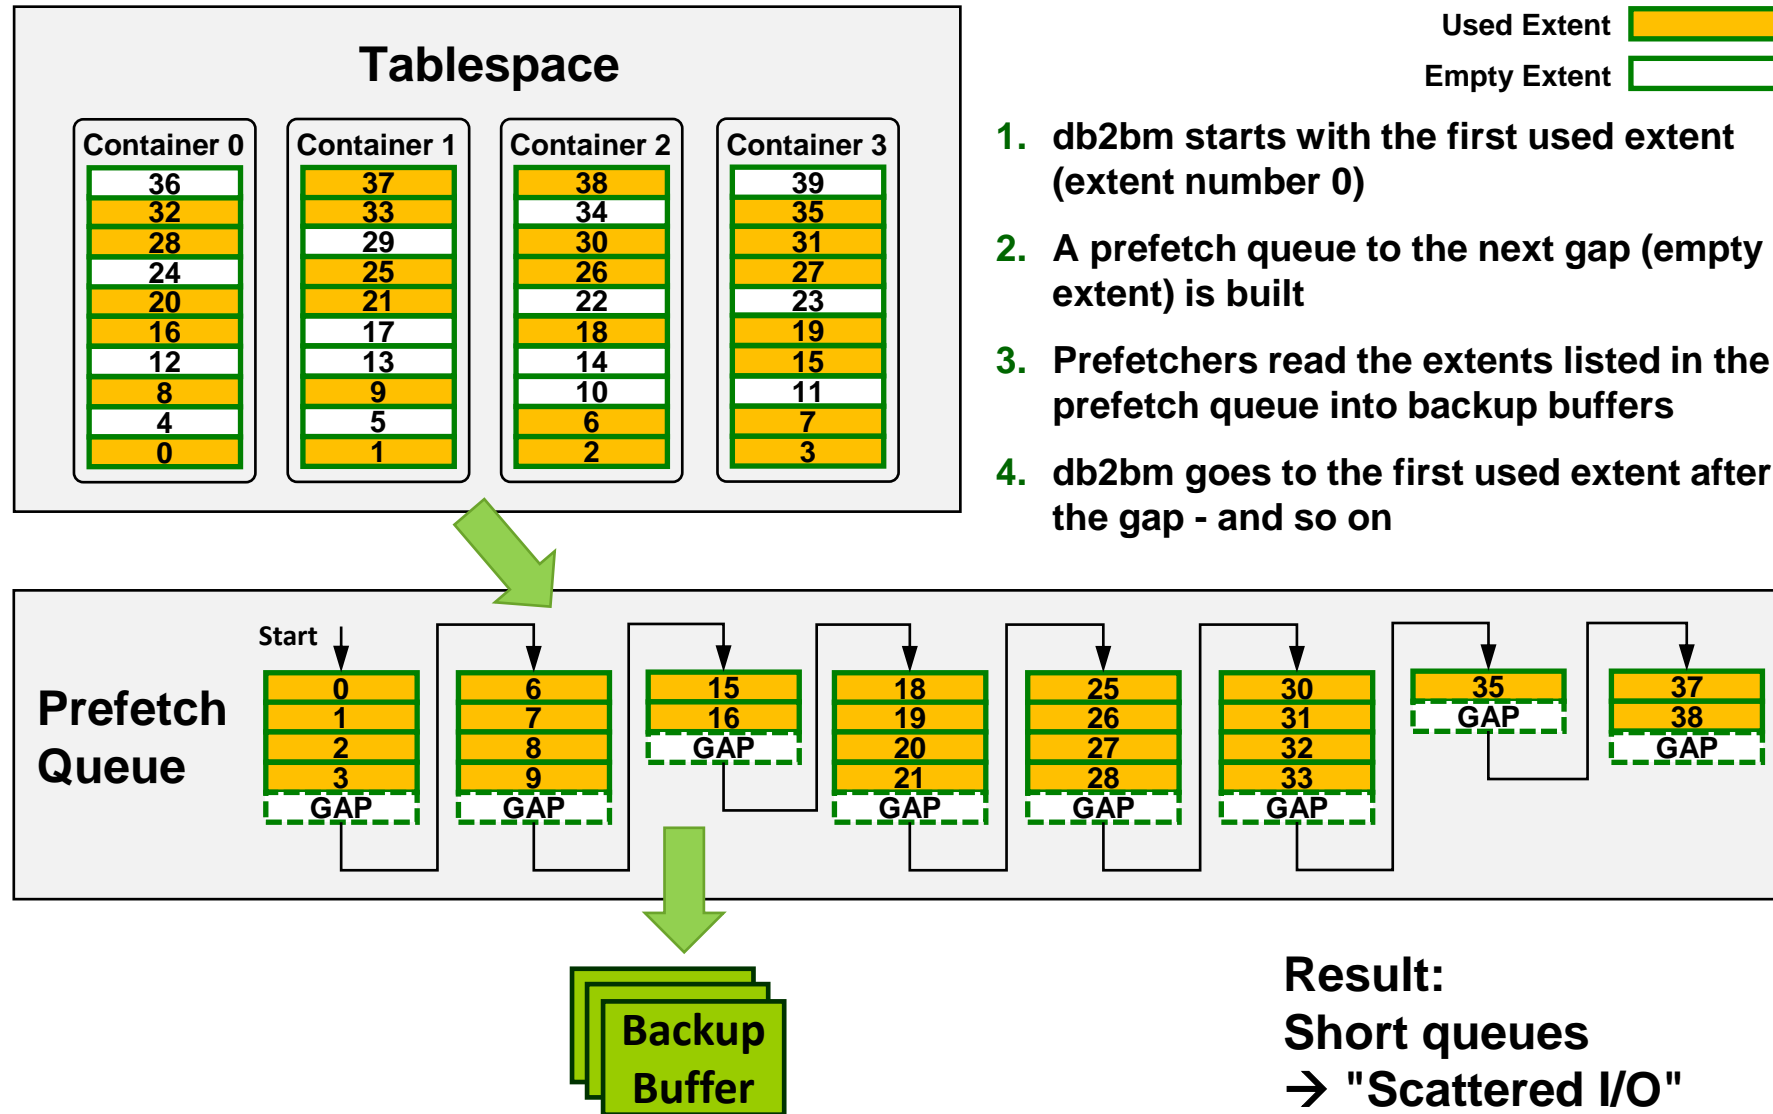
X86: intel provides a special zlib_next generation library, that uses SIMD instructions, factors 2-4 faster than traditional zlib implementation, uses Vector Units in Intel CPUs (zlib_ng)

Backup Execution: Tablespace Read Sequence



- Can be monitored using WaitQ element of BM section in the BARstats
- Addressed by `ADMIN_MOVE_TABLE SP` or `DB6CONV` (in SAP, or using tablespace pool) to redistribute tables to tablespaces

Backup Behavior in Case of Fragmentation



To address this, consider using `db2 ALTER TABLESPACE ... REDUCE MAX` command

Result:
Short queues
→ "Scattered I/O"
→ Increased backup time

Tools – db2pd, ...

- db2pd -utilities –repeat
- Or db2 list utilities
- dsmtop/db2top
- db2pd -barstats

```
db6db2:db2jkl 56> db2 list utilities
```

```
ID = 3
Type = BACKUP
Database Name = JKL
Member Number = 0
Description = offline db
Start Time = 08/23/2023 09:39:04.915061
State = Executing
Invocation Type = User
Throttling:
  Priority = Unthrottled
Progress Monitoring:
  Estimated Percentage Complete = 20
```

```
db6db2:db2jkl 55> db2pd -uti
```

```
Database Member 0 -- Active -- Up 0 days 00:49:23 -- Date 2023-08-23-09.39.08.291886
```

Utilities:

Address	ID	Type	State	Invoker	Priority	StartTime	DBName	NumPhases	CurPhase	Description
0x000000020B0EDFE0	3	BACKUP	0	0	0	Wed Aug 23 09:39:04	JKL	1	1	offline db

Progress:

Address	ID	PhaseNum	CompletedWork	TotalWork	StartTime	Description
0x000000020B0E9388	3	1	1472074924 bytes	22645651780 bytes	Wed Aug 23 09:39:04	n/a

Tools - BARstats

```
2016-08-25-12.01.57.841616-300 E15769840A2034 LEVEL: Info
PID : 10092648 TID : 39895 PROC : db2sysc 0
INSTANCE: db2vfb NODE : 000 DB : VFB
APPHDL : 0-2350 APPID: *LOCAL.db2vfb.160825165101
AUTHID : DB2VFB HOSTNAME: rubix-lp02.aus.stglabs.ibm.com
EDUID : 39895 EDUNAME: db2agent (VFB) 0
FUNCTION: DB2 UDB, database utilities, sqluxLogDataStats, probe:395
MESSAGE : Performance statistics
DATA #1 : String, 1514 bytes

Parallelism = 4
Number of buffers = 8
Buffer size = 16781312 (4097 4kB pages)

BM# Total I/O Compr MsgQ WaitQ Buffers MBytes Compr
MBytes
---
000 655.56 588.94 62.05 0.06 0.30 3406 74941 100569 25614
001 655.56 599.01 37.13 0.02 17.29 1665 54944 67462 12509
002 655.34 586.01 48.81 0.03 17.41 2523 62758 81687 18928
003 655.34 588.57 46.71 0.03 17.57 2212 66360 82893 16533
---
TOT 2621.81 2362.55 194.71 0.15 52.58 9806 259004 332613

MC# Total I/O MsgQ WaitQ Buffers MBytes
---
000 655.56 122.44 533.07 0.00 2474 39577
001 655.62 121.34 533.96 0.24 2441 39033
002 655.56 121.34 533.90 0.24 2458 39321
003 655.56 120.33 534.93 0.24 2437 38985
---
TOT 2622.31 485.46 2135.88 0.72 9810 156918

2016-08-25-12.01.57.841914-300 E15771875A477 LEVEL: Info
PID : 10092648 TID : 39895 PROC : db2sysc 0
INSTANCE: db2vfb NODE : 000 DB : VFB
APPHDL : 0-2350 APPID: *LOCAL.db2vfb.160825165101
AUTHID : DB2VFB HOSTNAME: rubix-lp02.aus.stglabs.ibm.com
EDUID : 39895 EDUNAME: db2agent (VFB) 0
FUNCTION: DB2 UDB, database utilities, sqlubcka, probe:1070
MESSAGE : Backup complete.
```

2 forms:

- 1) Monitor while backup is running: db2pd –barstats
- 2) Final BARstats in db2diag.log

- **1 BM** (buffer manipulator): Number of buffer manipulators launched
- **MC** (media controller): Number of media controllers launched
- **2 Total**: Time a BM or MC exists which equals the backup runtime
- **3 I/O**: Time spent for read and write operations
- **4 Compr**: Time spent for the compression
- **5 MsgQ**: Time waiting for an I/O buffer
- **6 WaitQ**: Time waiting for next instruction from the db2agent, can be used to tell if tablespaces are evenly distributed
- **7 Mbytes**: Amount of uncompressed and compressed data read or written by the different BM or MC. *Mbytes TOT* shows the size of the uncompressed and compressed backup.
- **8 Compr MBytes**: quantity of uncompressed data that was compressed

Checking Backup Images

- Use `db2ckbkp` to check images on disk
- Use `db2adutl` to check images in TSM
- Use the registry variable
`DB2_BCKP_PAGE_VERIFICATION` to
check page consistency during the backup



Tools – db2ckbkp

- db2ckbkp <image>
 - Checking the image
- db2ckbkp -H
 - Reading the header of the image

```
db2jkl 54> db2ckbkp
JKL.0.db2jkl.DBPART000.20230823093905.001
```

```
[1] Buffers processed:    ###...###
```

```
Image Verification Complete - successful.
```

```
=====
MEDIA HEADER REACHED:
=====
Server Database Name      -- JKL
Server Database Alias    -- JKL
Client Database Alias     -- JKL
Timestamp                 -- 20230823093905
Database Partition Number -- 0
Instance                  -- db2jkl
Database Configuration Type -- 0 (Non-shared data)
Sequence Number           -- 1
Database Member ID        -- 0
Release ID                -- 0x1500 (DB2 v11.5.8.0)
AL version                -- V:11 R:5 M:8 F:0 I:0 SB:0
Database Seed             -- 0x78861BF4
DB Comment's Codepage (Volume) -- 0
DB Comment (Volume)      -- SAP database JKL
DB Comment's Codepage (System) -- 0
DB Comment (System)      -- SAP database JKL
Authentication Value      -- 255 (Not specified)
Backup Mode               -- 0 (Offline)
Includes Logs             -- 0 (No)
Compression              -- 2 (Encrypted)
Backup Type               -- 0 (Database-level)
Backup Granularity        -- 0 (Non-incremental)
Merged Backup Image       -- 0 (No)
Status Flags              -- 0x1
                          -- Consistent across all members
System Catalogs in this image -- 1 (Yes)
Catalog Partition Number  -- 0
DB Codeset                -- UTF-8
DB Territory              -- US
LogID                     -- 1684329437
LogPath                   -- /db2/JKL/log_dir/
Backup Buffer Size         -- 16781312 (4097 4K pages)
Number of Sessions        -- 1
Platform                  -- 0x1E (Linux-x86-64)
Encrypt Info Flags        -- 0x1
                          -- Source DB was encrypted
```

```
The proper image file name would be:
JKL.0.db2jkl.DBPART000.20230823093905.001
```

```
Image header dumped -- NO VERIFICATION PERFORMED.
```


Tools – db2 history file

- db2 list history backup all for <dbname>
 - Operation: **Backup**
 - Object: **Database**
 - Type: **Full**



```
Op  Obj  Timestamp+Sequence  Type  Dev  Earliest Log  Current Log  Backup ID
-----
B   D   20230627111747004    F    N  S0000000.LOG  S0000000.LOG
-----
Contains 155 tablespace(s):

00001 SYSCATSPACE
00002 SAPTOOLS
00003 SAPEVENTMON
00004 SYSTOOLSPACE
...
00049 ZTOM#TEST@01D
00050 ZTOM#TEST@01I
00051 ZTOM#TEST@01L
...
00136 D01#ZBWTP@10D
00137 D01#ZBWTP@10I
00138 D01#ZBWTP@10L
...
-----
Comment: DB2 BACKUP D01 OFFLINE
Start Time: 20230627111747
End Time: 20230627115722
Status: A
-----
EID: 803291 Location: /backup
```

Layout & Configuration – Registry + DB CFG

- DB2_BCKP_INCLUDE_LOGS_WARNING
 - Default: False
 - SAP customers: True
- DB2_BCKP_PAGE_VERIFICATION
 - Default: False
- DB2_BCKP_COMPRESSION
 - Default: Compress
 - Values: NX842, ZLIB

db2 get db cfg for <dbname>

- Encryption Library for Backup (ENCRLIB) = libdb2encr.so
- Encryption Options for Backup (ENCROPTS) =
CIPHER=AES:MODE=CBC:KEY
LENGTH=256
- Encrypted database = YES
- Keystore type (KEYSTORE_TYPE) = PKCS12
- Keystore location (KEYSTORE_LOCATION) =
<PATH/keystore.p12

Settings influence the backup command

More Reading

- [Understanding Db2's Backup and Restore Statistics](#)
- [Db2 backup command in Db2 Knowledge Center \(Db2 11.5\)](#)
- [Hardware accelerated backup and log file compression in Db2 Knowledge Center \(Db2 11.5\)](#)
- [Power NX842 Compression for Db2](#)
- [Db2 Encryption overview in Db2 Knowledge Center \(Db2 11.5\)](#)
- [Example output for backup performance monitoring in Db2 Knowledge Center \(Db2 11.5\)](#)
- [db2ckbkp – check backup command in Db2 Knowledge Center \(Db2 11.5\)](#)



SAP on Db2 documentation on SAP Help Portal


<https://help.sap.com/viewer/p/DB6>

- New/Updated SAP / Db2 guides

SAP on Db2 in der SAP Community

<https://www.sap.com/community/topic/db2-for-linux-unix-and-windows.html>

- News about Db2 development
- Recent Blogs



Installation and Upgrade

Learn how to install and upgrade SAP systems running on IBM Db2.

Software Logistics Toolset

The software logistics toolset comprises all tools for installation, system copy and rename, dual-stack split, and system upgrade. The pages on SAP Support Portal also provide links to the relevant guides.

IBM Db2 High Availability Solution: IBM Tivoli System Automation for Multiplatforms

Learn how to set up a high-availability (HA) solution for IBM Db2 for Linux, UNIX, and Windows using IBM Tivoli System Automation for Multiplatforms (SA MP).

Running an SAP System on IBM Db2 11.1 with the Db2 pureScale Feature

Learn how to set up and run an SAP system running on IBM Db2 for Linux, UNIX, and Windows 11.1 with the IBM Db2 pureScale Feature.

Running an SAP System on IBM Db2 10.5 with the Db2 pureScale Feature

Learn how to set up and run an SAP system running on IBM Db2 for Linux, UNIX, and Windows 10.5 with the IBM Db2 pureScale Feature.

Running an SAP System on IBM Db2 10.1 with the Db2 pureScale Feature

Learn how to set up and run an SAP system running on IBM Db2 for Linux, UNIX, and Windows 10.1 with the IBM Db2 pureScale Feature.



Administration

Read how to monitor and administrate the IBM Db2 database for your SAP system.

Database Administration Guide for SAP on IBM Db2 for Linux, UNIX, and Windows

This is our main information resource about how to administrate SAP systems running on the database IBM Db2 for Linux, UNIX, and Windows.

SAP Business Warehouse on IBM Db2 10.5 and Higher for Linux, UNIX, and Windows: Administration Tasks

This is administration information especially for SAP BW systems running on IBM Db2 version 10.5 and higher.

SAP Business Warehouse on IBM Db2 10.1 and Lower for Linux, UNIX, and Windows: Administration Tasks

This is administration information especially for SAP BW systems running on IBM Db2 version 10.1 and lower.

Enabling SAP Business Warehouse to Use IBM Db2 for Linux, UNIX, and Windows as Near-Line Storage

Learn how to set up and operate a near-line storage solution on IBM Db2 for Linux, UNIX, and Windows.

DBA Cockpit (for SAP Systems Based on SAP NetWeaver 7.02 and Higher)

The DBA Cockpit is a tool for monitoring and administrating SAP systems running on IBM Db2 for Linux, UNIX, and Windows.

DBA Cockpit (for SAP NetWeaver 7.0 and 7.1 Only)

(PDF)

DBA Cockpit (for SAP NetWeaver 7.01 Only)

(PDF)

Technical Monitoring Cockpit

Learn more about the technical monitoring cockpit, our upcoming tool for monitoring SAP systems.



Database Upgrade

Read how to upgrade the IBM Db2 database for your SAP system.

Upgrading to Version 11.1 of IBM Db2 for Linux, UNIX, and Windows

Upgrading to Version 10.5 of IBM Db2 for Linux, UNIX, and Windows

Upgrading to Version 10.1 of IBM Db2 for Linux, UNIX, and Windows

Upgrading to Version 9.7 of IBM Db2 for Linux, UNIX, and Windows



More Information

SAP on IBM Db2 for Linux, UNIX, and Windows Community

Go to our community to find blogs, news and more about SAP on IBM Db2 and to ask questions.