

Db2 TAB: A look at Db2 11.1.4.4

Roger E. Sanders

IBM

Session code: TAB



Please note :


- **IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice and 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.**

A quick overview

Db2 11.1: A new beginning

- **Starting with DB2 11.1, we have begun to move towards a “continuous delivery” model**
 - Explicitly plan to deliver function throughout the active development lifetime of a version
 - New versions will still happen occasionally for various reasons
- **What does this mean to you?**
 - Introduction of Modification levels to the product signature
 - Function delivered continuously during a version’s lifecycle
 - No more “big gulps”, just a series of “little bites”

The Db2 11.1 lifecycle



Regular fix/mod pack deliveries (every 6-9 months)

- With overriding focus on stability for immediate production deployment
- Modifications will contain select functional enhancements (often off by default)

GA, FP1, FP2, FP3 Highlights: Mission Critical Workloads

Comprehensive Enterprise Security



Enterprise Encryption

- Centralized Key Managers (KMIP)
- PKCS#11 HSM support
- SSL Encryption for non-pS HADR (Linux x86, ALL platforms)

Mainframe Class Availability

pureScale “for the Masses”

- Up and running in hours
- Zero data loss DR with HADR
- Multi-switch GDPC (no SPOF)
- Seamless HADR pureScale upgrades
- Restore REBUILD
- Online Index Create
- Add/Drop CF

- Faster MCR by default
- Improved sockets performance
- Multiple hosts in maintenance mode
- Recover HADR Standby w/ tbspc restore (now for pS)
- Improved XA performance

HADR ease of use:

- Export HADR TSAMP configuration to XML file

Even Greater Availability

- Online Crash Recovery
- Faster Rollback of Multiple and Large Txn
- ADMIN_MOVE_TABLE advances
- Better tablespace state and replay window monitoring with HADR
- Per partition OLR
- Avoid Lock Escalation
- Timeout support for vendor Archive

Significant Core Database Advances



Increased Serviceability

- 16x increase in active log space capability
- Enhanced monitoring for lock and logging events

Additional Core Function

- Workload Manager (WLM) multi-tenancy extensions
- Federation simplification and integration

Simpler, Faster, More Online Upgrades

- Faster, no need for offline backup
- Streamlined HADR upgrade
- DB2 Version 9.7 direct to 11.1

Very Large Database Performance

- Higher user throughput

GA, FP1, FP2, FP3 Highlights: SQL & Warehousing Workloads

SQL and Compatibility



PDF manuals for DB2 LUW v11 Multi-Lingual SQL Advances

- BINARY, EXTENDED ROWSIZE
- CHAR(255)

BLU Restriction Removal

- Automatic Dictionary Creation improvements
- ALTER VARCHAR/VARGRAPHIC length support
- Codepage 819

Richer and Even More Compatible SQL

- JSON support
- BOOLEAN data type
- Common table expression (Db2 for z/OS compatibility)
- WITH and SELECT INTO support
- Alias for XML

Massive Scale Warehousing at In-Memory Performance



MPP BLU Scalability

- PB scale in-memory warehousing

Next Gen In-Memory Performance, Function and Workloads



Faster ELT/ETL performance More Query Workloads Optimised More Function supported

- Generated Columns
- RCAC
- User maintained temporal tables

Continued Performance Improvements

- Synopsis table enhancements
- Additional SIMD exploitation
- Aggregation enhancements, sort elimination
- INSERT from sub-select – multi-core parallelism improvements
- BLU Acceleration indexes

Just an FYI: RFEs that have been shipped so far in DB2 11.1

RFE Id	Summary	Level
113550	db2cklog does not show logs are encrypted	11.1.3.3
94883	HADR incorrect SSL protocol version (support on non-Linux)	11.1.3.3
61542	ROLLFORWARD parameter request: have option -fromdb (like -fromnode)	11.1.3.3
87417	Timeout option for log archiving sessions (UNIX only)	11.1.3.3
47816	Certification RFE for DB2 hosted with RHEL clustered using Veritas for automated failover	v11
103623	PDF manuals for DB2 LUW v11	n/a
86557	No deadlock error with implicit db2rbind	11.1.2.2
104160	The software support for the compatibility with COBOL for AIX V5.1	11.1.2.2
102317	Improve Db2 handling of DGTs DDL log records in HADR (also during rollforward process)	11.1.2.2
97080	Output of change history event monitor incomplete	11.1.2.2
96096	DB2 BLU enable data skipping for "select with wildcard"	11.1.2.2
65928	Tablespace-level Restore and Rollforward on Standby database	11.1.2.2
97727	Improve deadlock problem in DB2 10.5	11.1.2.2
99811	Reduce xhshlct contention with high concurrency applications	11.1.2.2
100427	Inconsistency concat behavior between db2 zOS and db2 luw	11.1.2.2
73583	Allowing tablespace-level restores for PRIMARY/STANDBY DBs in HADR env	11.1.2.2
102881	Redirect restore not working because of missing free disk space	11.1.2.2
80372	A restore operation with the REBUILD option is not supported in a Db2 pureScale environment	11.1.2.2
103573	Request about DB LUW support MS Visual Studio 2015	11.1.2.2
96245	TRUNC compatibility with BLU	11.1.2.2
100424	AS identifier is mandatory for xquery argument in XMLQUERY, which is inconsistent behavior with db2 z/OS	11.1.2.2

Db2 11.1.4.4


Db2 11.1.4.4






- **This is most likely the final modification level for Db2 11.1**
 - I.e. future Db2 11.1 updates after this will just contain fixes
 - Next updates would be DB2 11.1.4.5 followed by Db2 11.1.4.6
- **Key content:**
 - Extent reclaim support for pureScale
 - JSON enhancements
 - HADR Read on Standby (RoS) enhancements
 - Preview of 4K sector size support

Extent reclaim support for pureScale

- **When table objects are dropped, the storage may not be immediately available to be released from the tablespace**
 - Only space above the tablespace high water mark (HWM) can be released
- **The ALTER TABLESPACE ... REDUCE statement attempts to reduce the high water mark for the table space by moving live extents to unused ones lower in the tablespace**
 - Supported with reclaimable storage DMS table spaces (created in Db2 9.7 or newer)
- **Not *yet* supported on pureScale..... but we're working on it for Db2 11.1.4.4!**

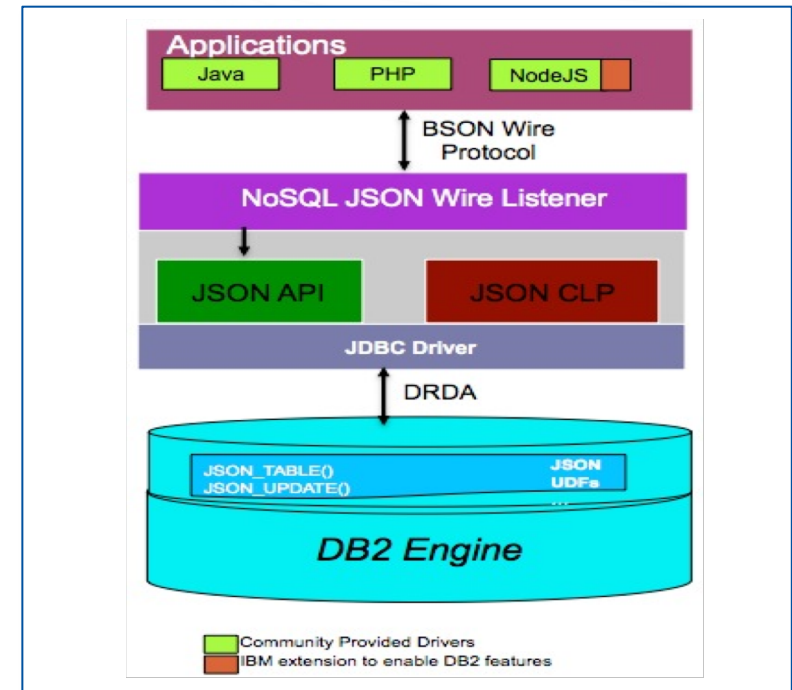
Extent reclaim



-  Internal table space metadata extents
-  Table 1
-  Table 2
-  Table 3
-  Extent that is allocated to a table space, but not to a table

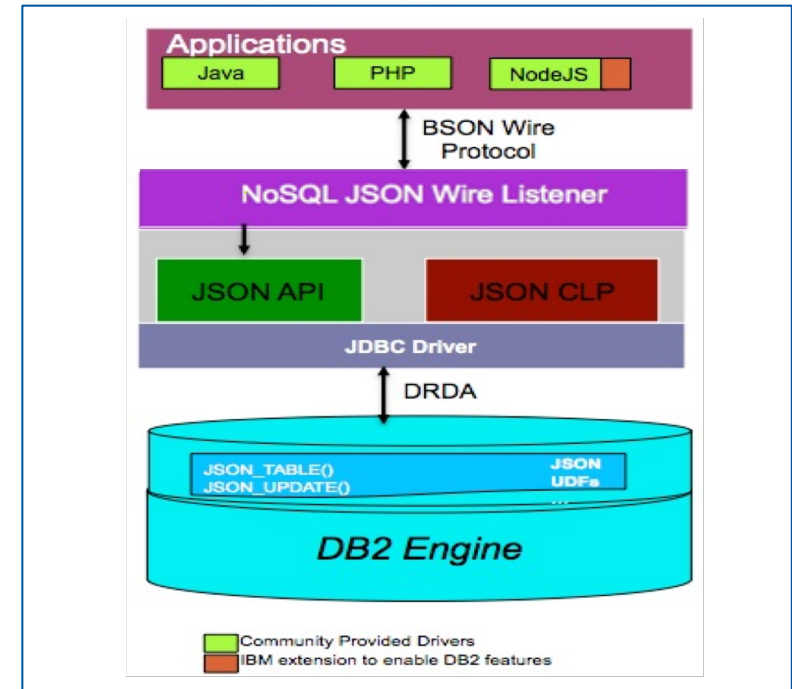
A brief history of Db2 & JSON (Part 1)

- **Db2 10.5 FP1 introduced JSON NoSQL support**
 - Focused on allowing Db2 to participate in the NoSQL paradigm
- **This support used undocumented, proprietary JSON SQL functions in Db2 server**
 - aka “SYSTOOLS” JSON SQL functions



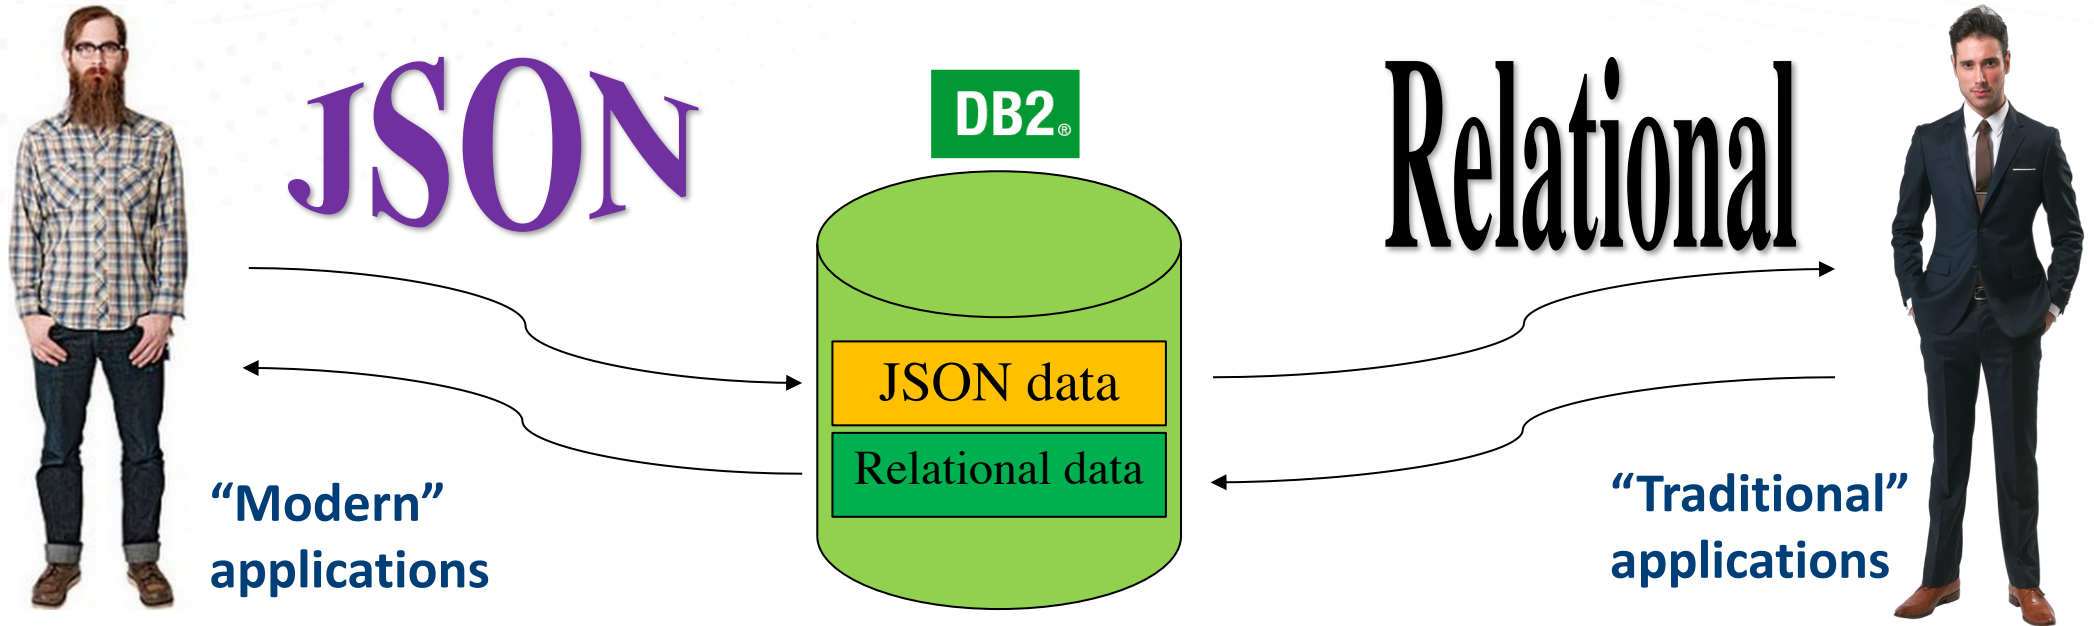
Enhancements to JSON NoSQL support

- Db2 NoSQL JSON wire listener is enhanced to support latest mongo db client (3.6.3)
- Added support for Kerberos Authentication to wire listener
- Enhanced wire listener script to enable jcc and nosql trace using command line options.



A brief history of Db2 & JSON (Part 2)

Our customers began to ask for native SQL support for JSON in Db2



As a tactical response, we revealed our proprietary "SYSTOOLS" JSON SQL functions

First wave of ISO JSON SQL functions coming in Db2 11.1.4.4!

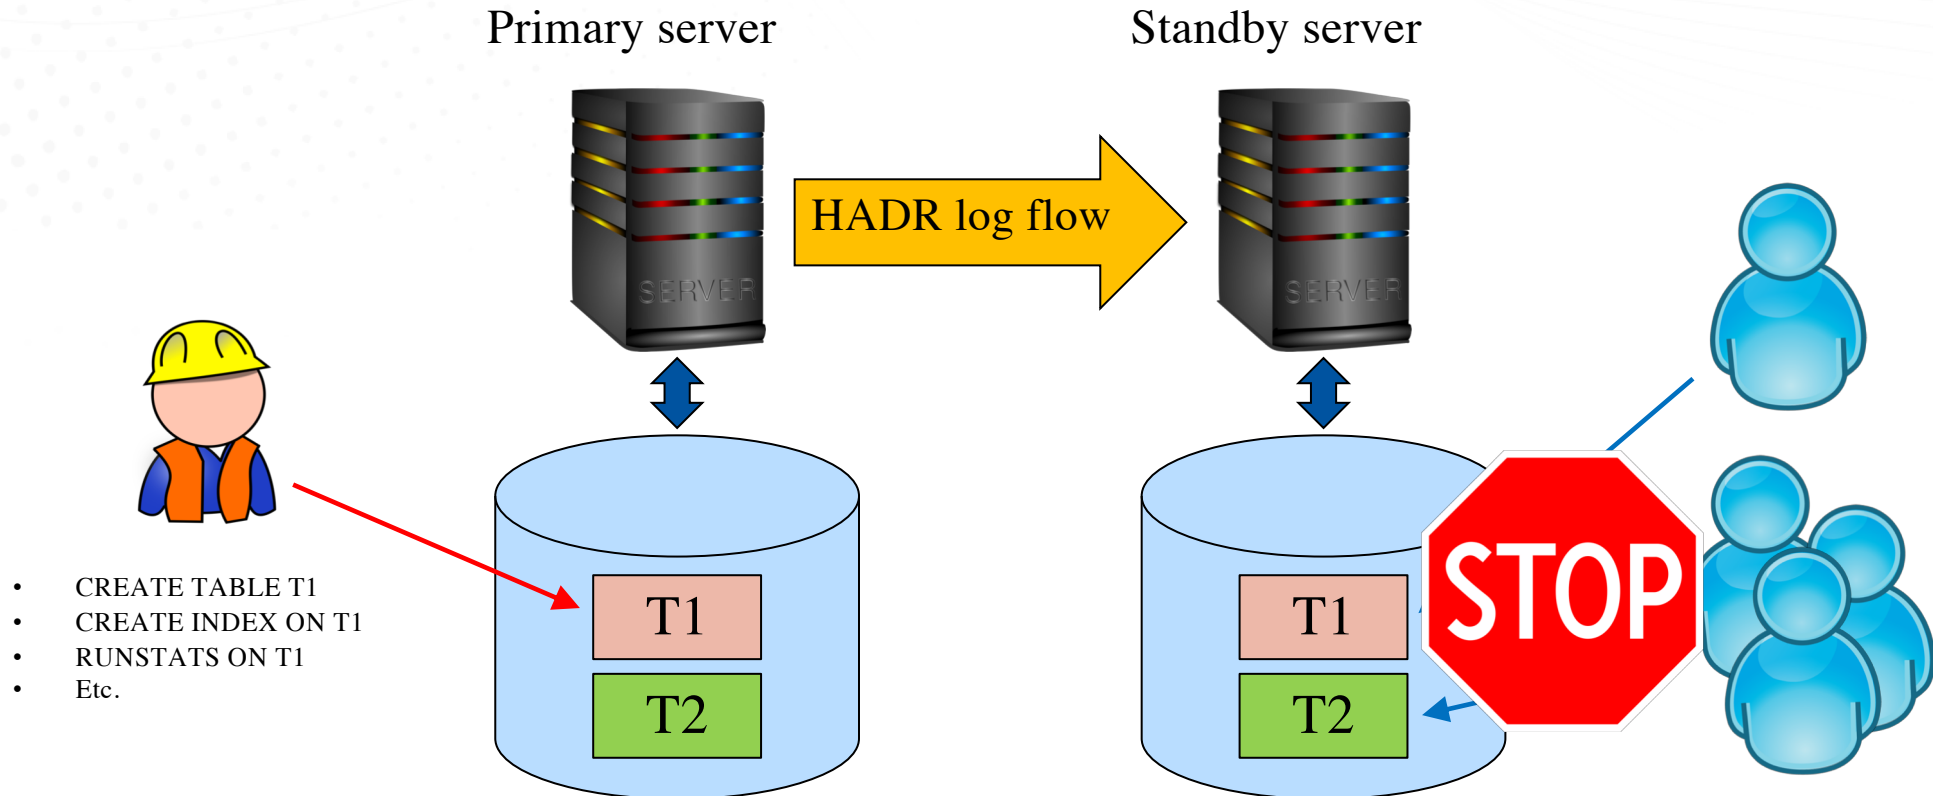
- First set of JSON SQL functions based on ISO technical report on SQL support for JavaScript Object Notation (JSON)
- The (proprietary) SYSTOOLS functions will be de-emphasized but will continue to be supported

Schema	Name	Comments
SYSIBM	BSON_TO_JSON	Convert BSON formatted document into JSON strings
SYSIBM	JSON_TO_BSON	Convert JSON strings into a BSON document format
SYSIBM	JSON_ARRAY	Creates JSON array from input key value pairs
SYSIBM	JSON_OBJECT	Creates JSON object from input key value pairs
SYSIBM	JSON_VALUE	Extract an SQL scalar value from a JSON object
<i>SYSIBM</i>	JSON_QUERY	Extract a JSON object from a JSON object
<i>SYSIBM</i>	JSON_TABLE	Creates relational output from a JSON object
SYSIBM	JSON_EXISTS	Determine whether a JSON object contains the desired JSON value

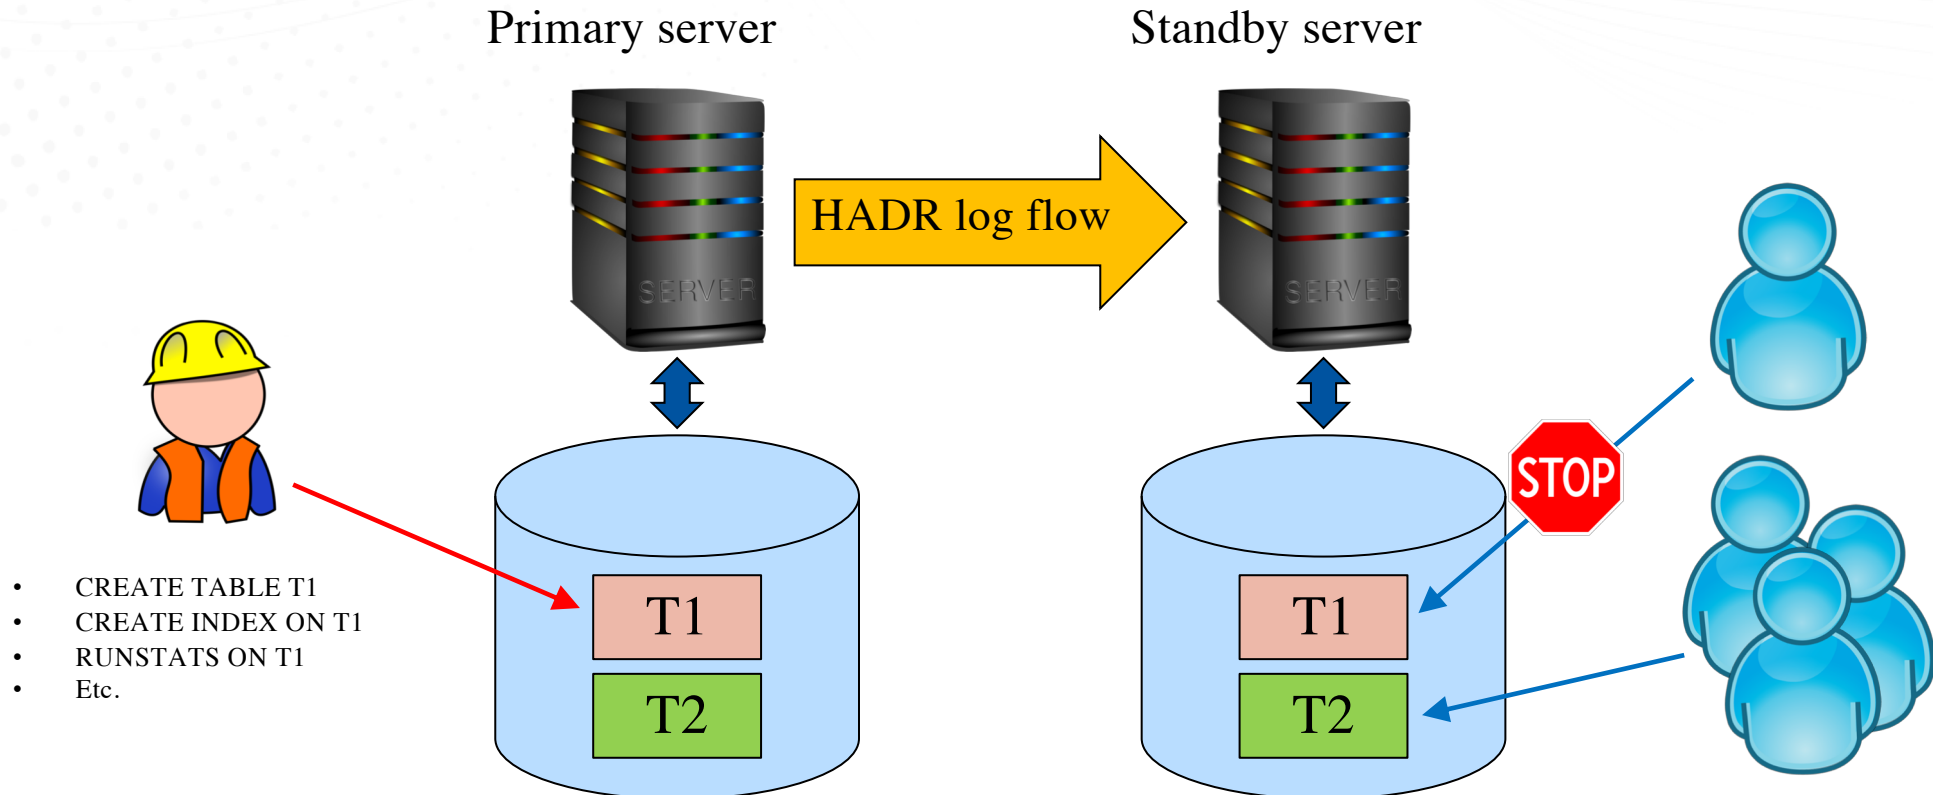
What makes these JSON SQL functions different/better?

- **Following a public “standard” (ISO technical report)**
- **Easier to use**
 - No need to qualify call to function or add SYSTOOLS to function path
 - No need to grant EXECUTE privilege
- **Simpler and more flexible storage options**
 - You choose the stored format: JSON or BSON
 - You choose the table organization: row or column
 - You choose the column data type: BLOB, CHAR, CLOB, VARBINARY, VARCHAR
- **Conversion functions are now optional not mandatory**

HADR Read on Standby (RoS) as it exists today...



HADR Read on Standby (RoS) as it exists in Db2 11.1.4.4...



Preview: Support for storage devices using 4K sector size

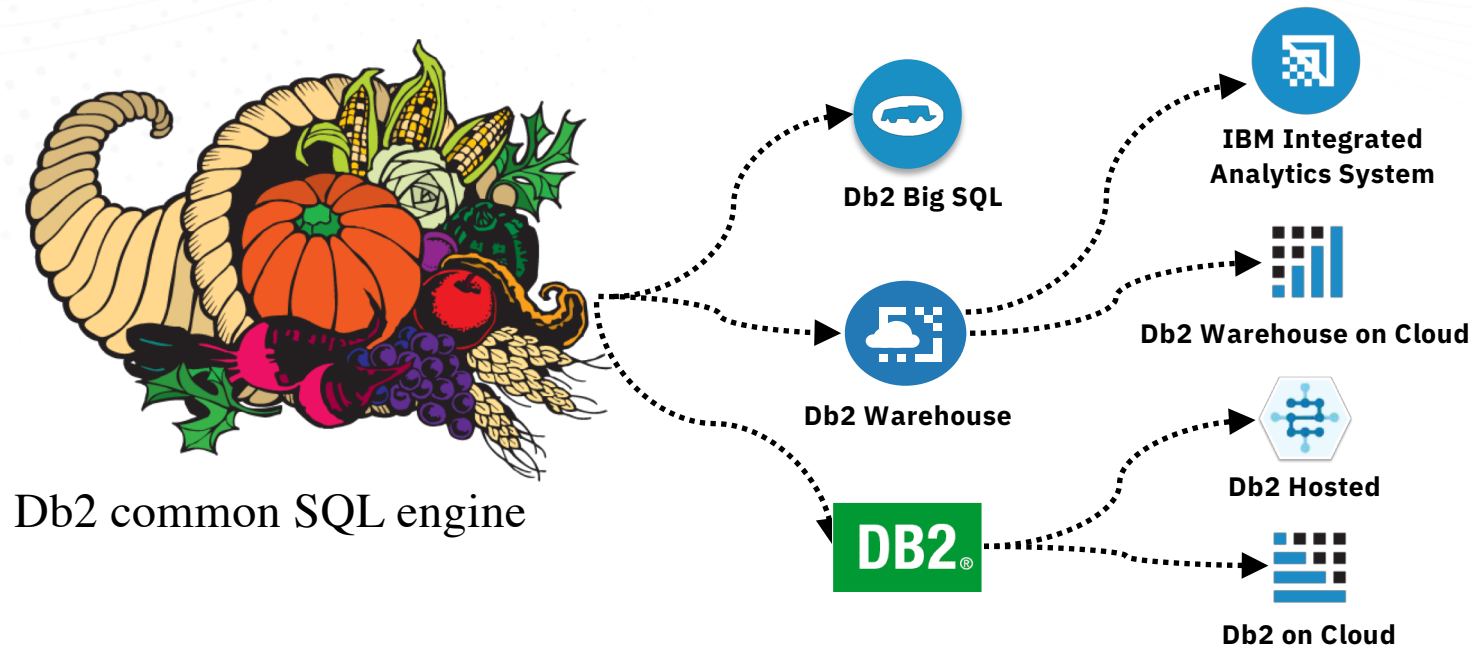
- **This support can be enabled by setting the `DB2_4K_DEVICE_SUPPORT` registry variable to `ON`**
- **This feature is not supported for production use at this time**
 - Try it out and let us know!
- **Restrictions and limitations:**
 - Not available on pureScale
 - The use of DMS Raw containers is not supported
 - Backup and load copy files will be slightly larger
 - There may be a performance penalty accessing:
 - LOB data stored on 512-byte sector storage
 - Backup or load copy files created prior to the enablement of 4K device support

Db2 “VNext”

Db2 “VNext”

- **We plan to release a new version of Db2 in 2019**
 - This will get all our Db2 and Db2 Warehouse offerings on the same level of the Db2 common SQL engine
 - Db2 11.1 will be 3 years old in 2019!!
- **We will continue to follow the “continuous delivery” paradigm started with Db2 11.1**
 - Steady release of function through GA and subsequent modification updates
 - Db2 11.1 shipped 4 modification updates since GA in 2016
- **Version/Release number for Db2 VNext not decided yet**

Db2 Common SQL Engine product lines



Db2 VNext GA

- **All users of the Db2 common SQL engine will move to the new level**
 - E.g. Db2, Db2 Warehouse, IIAS, Hosted Db2 and managed DB2 Warehouse cloud services
 - Reduced latency between feature appearance in different form factors as development will all be on common level
- **Will contain any enhancement that is “ready to go” for on-premise environments**
 - Subsequent modification updates will contain more!

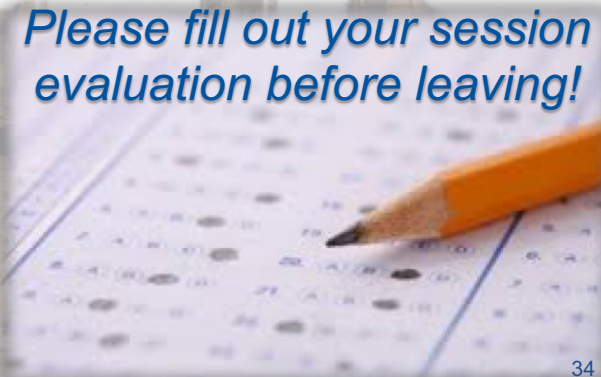
Questions?

Roger E. Sanders

IBM

roger.sanders@us.ibm.com

Session code: TAB



*Please fill out your session
evaluation before leaving!*