

DB2 (Linux Unix & Windows)



Presented by : Ramesh R Gattu
Database Engineering – DB2 (LUW) / NoSQL
Infrastructure & Operations
Nationwide Insurance, Columbus Ohio



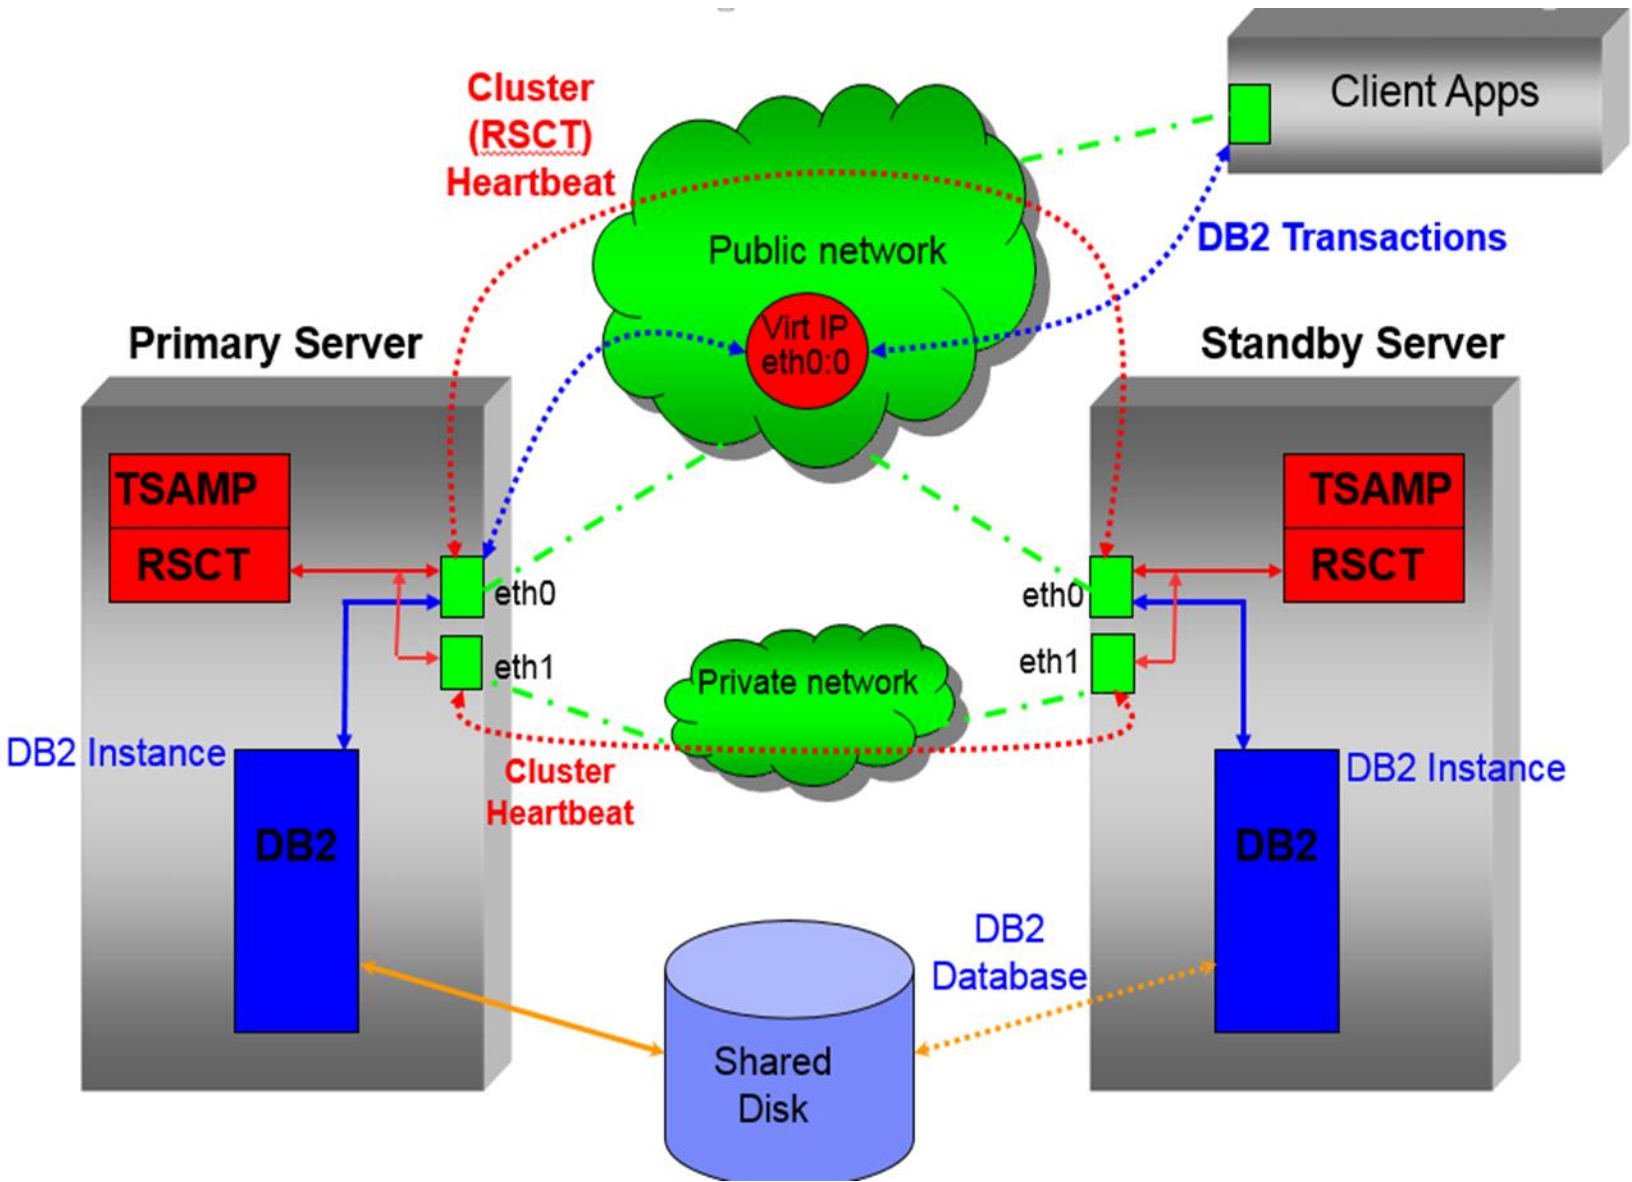
Nationwide[®]
is on your side

DB2 (Linux Unix & Windows)

- Distributed DB2 (LUW) HA Redesign
 - Shared single disk to Dual Disk conversion
- ROR (Read Only Root) binary disk usage for DB2 database software installation
- Day to Day operational enhancements
 - Server / Database restarts
 - Validation enhancement for restarts – Using shell and curl
 - Shared file system usage – Using auto mount process
 - Netbackup / Storage – Cost effectiveness for RDBMS and Big data computing
- Data / Database as a Service (Push button / Cloud)
- Third Party tools on DB2 (LUW) and best practices:
 - Tripwire (Security scanning)
 - Guardium (Data security)
 - Qualys (Vulnerability management)
 - IIQ (Identity IQ – Privilege access control and more)
 - Lieberman (Enterprise random password manager)
 - DB2 native logging & encryption

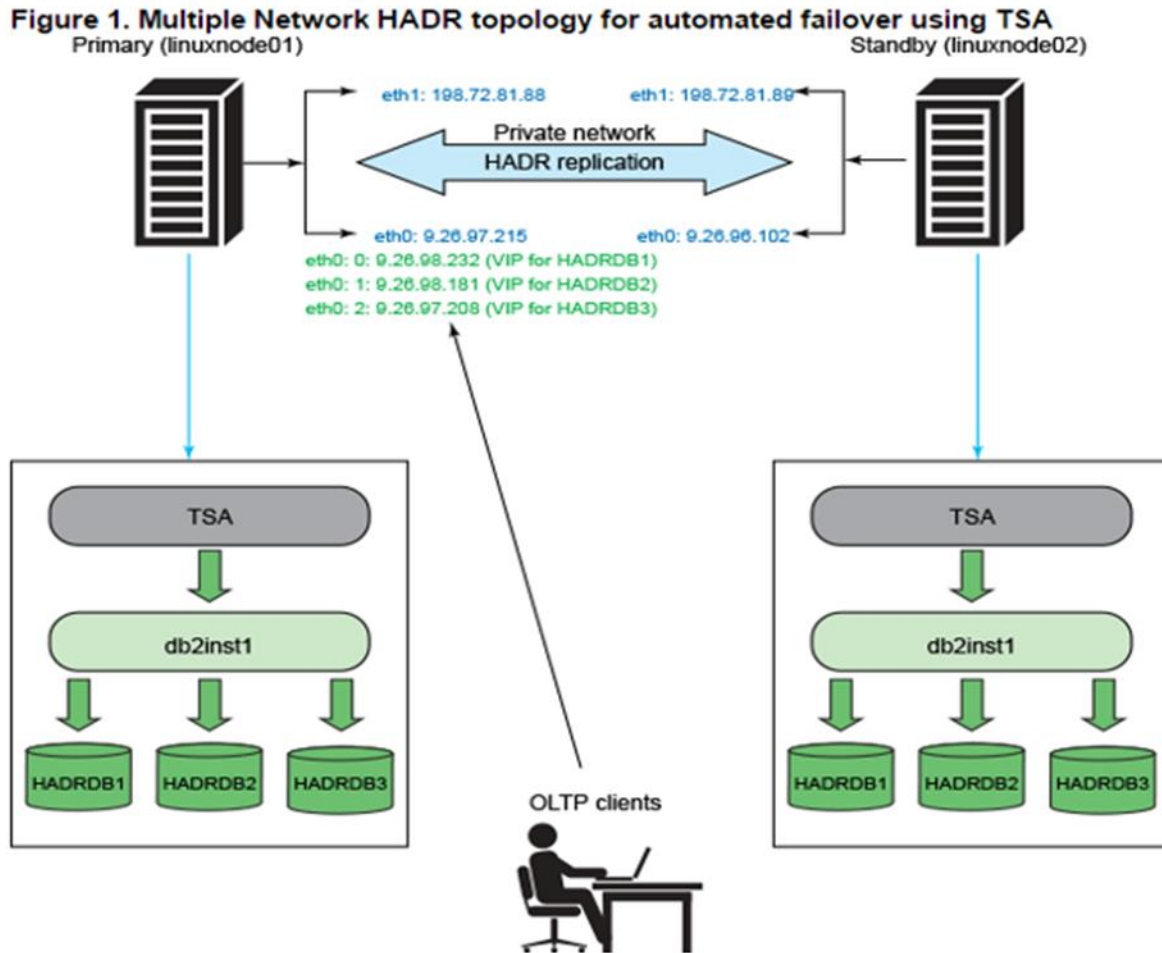
DB2 (LUW) HA Redesign – Single to Dual Disk

- Distributed DB2 (LUW) HA Redesign : Shared single to Dual Disk conversion
- Example of single disk solution – <https://www.ibm.com/developerworks/data/tutorials/dm-009db2hadr/>



DB2 (LUW) HA Redesign – Single to Dual Disk

- Distributed DB2 (LUW) HA Redesign : Shared single to Dual Disk conversion
- Example of Target dual disk solution – <https://www.ibm.com/developerworks/data/tutorials/dm-1009db2hadr/>



There are two nodes in the topology:

- linuxnode01 — the primary node that hosts the HADR primary databases HADRDB1, HADRDB2, and HADRDB3
- linuxnode02 — the standby node that hosts the respective HADR standby databases

DB2 (LUW) HA Redesign – Single to Dual Disk

- Distributed DB2 (LUW) HA Redesign : Shared single to Dual Disk conversion

Challenging but fun facts:

- Obtaining customer sign-off (explaining DUAL storage charges and redoing of entire DNS/vip)
- Downtime during conversion
- Working with IRM ,Network, and Firewall teams to create an approximate 100 plus brand new vips and associated DNS entries
- Sitting TSAMP software was older version EOL dates approaching soon
- Non recoverable load situation with batch jobs (shared file system solution)
- Guardium dependency
- Re – tailoring performance and monitoring tools like Ignite ,ESM and site scope.

High level Steps:

Stakeholders:zVM, Linux, Storage, Network, Firewall and Performance teams

- Line-up required storage ,DNS and IRM / Audit exceptions in place few weeks earlier.
- DB2 team ensures the shared file systems are mounted on the node that was tailored as the primary owner of the file systems during initial setup. If not a planned failover is conducted to perform the task
- zVM / Linux will tear down the cluster setup gracefully and ensures old dns is latched to the primary
- DB2 team breaks the existing cluster configuration between Primary and Standby

DB2 (LUW) HA Redesign – Single to Dual Disk

- Distributed DB2 (LUW) HA Redesign : Shared single to Dual Disk conversion

High level (Steps continued)

- Linux will work on the stand by server and perform tasks like creating the file systems ,allocating storage, creating local users and granting privileges etc..
- DB2 team will complete database build work on standby (this includes creating instance ,building database ,backup restore from Primary and setting up High Available Disaster Recovery
- Majority of the commands require root access but it is recommended that you work with linux team and obtain proper “SUDO” rulesjust for the main commands.
- Both virtual nodes will be prepped for required authentication as root or sudo
- DB2 team will now setup new DB2 HA Redesign using DB2HAICU tool. This will attach brand new DNS/vip to the primary server .At this stage primary server has both old and new DNS attached.
- Middleware will update JVM/JNDI entries replacing old DNS with new.
- Validation will be performed by DB2-Run team and signed off.
- Old DNS will now be cleaned up OFF primary server.

DB2 (LUW) HA Redesign – Single to Dual Disk

- Distributed DB2 (LUW) HA Redesign : Shared single to Dual Disk conversion

High level Tech notes :

Pre-work –Linux:

- Removing the Peer value from both servers parmfile, and remove the SAN entries for the inactive node.
- /etc/zfcp.conf will need to be cleared on the inactive node. Storage can then be added to the inactive node without also being added to the active
- Rename the existing luns .
- Update FSTAB such that all the core file systems are “AUTO” mounted every time server is rebooted
- Create netmon.cf file on both primary and stand by servers under /var/ct/cfg/ .The netmon.cf file should be updated with the quorum IP. This is used by RSCT to perform the required failovers
- Enable effective syslog messaging by configuring /etc/syslog.conf file. Ensure output files are rotated every day and keep last 14 files (recommended).

Things to make a note of (lessons learned) :

- We have seen some instances of server triggering a reboot. IBM advised that the issue was due to a very low default LSCOMG value .When there is a blimp in the TCP/IP protocolinstead of waiting the primary server immediately triggers a reboot to correct the issue. Problem is corrected by updating LSCOMG value to a “4”
- Reminding ourselves to re-build the cluster (db2haicu) after every major version upgrade.
- Ensure DB2 auto-start parameter is TURNED OFF
- Having more than one gateway for highly available assets.

DB2 (LUW) Process Improvements

ROR (Read Only Root) binary disk usage for DB2 database software installation

Objective :To eliminate manual process of loading DB2 software on each individual server to save time and minimize storage tickets pertaining to /opt/

- Read Only Root disk is created using secure form of DASD.
- Once created the ROR disk can be Read/Write enabled and access controlled via a single virtual server in the data center
- Required DB2 software can now be loaded and untarred on to ROR
- ROR can now be mapped on to any target server and required upgrade/install can be performed .

Things to make a note of (lessons learned) :

- Internal Audit and Risk management approvals
- Separate ROR disks (one each for Test and Prod)
- Perfect solution for TEST (lower environments)
- Production environment should be approached on a case by case basis depending on the service level agreements and application availability
- Alert mechanism in place for mass server reboots

DB2 (LUW) Process Improvements

Server / Database Restarts :

Server reboot / DB2 Auto start stand alone: Using scripts via boot sequence

Server reboot / DB2 Shared disk, Dual disk with HADR start/stop automation:

Opportunity: On a quarterly basis, we usually stop/start DB2 on our zLinux servers for various OS activities. We have to stop/start/failover DB2 in all our servers in a fixed or split schedule, at the same time. This involves a lot of manual effort/work and potential human errors if we miss any sequence of step(s) while stop/start DB2 and also failover in case of clusters.

Automated solution: Developed using shell and curl the new process helps STOP, START, failover DB2 database(s) in any given server they are run on, in an orderly fashion, thereby saving time and avoiding issues.

High level steps involved:

- Script identifies and creates workload based on stand alone, Tivoli running single disk and Tivoli running on HADR dual disk
- Puts Tivoli in maint mode ,deactivates HADR and stops DB2
- Starts DB2, activates HADR and brings Tivoli out of maint mode
- Runs a report against all the servers, generates an email if it sees any issues pertaining dns/vip, DB2 startup ,lssam and HADR state
- Finally embedding the auto –restart in to init boot sequence

DB2 (LUW) Process Improvements

Shared file system usage :

- Central file system for day to day database activities
- Can be provisioned on NAS and enabled to automount on any server in the data center on demand with regular “change directory” command .
- Minimize the usage of core file systems like FS and LOGFS for maint activities

Netbackup process

- De dupe process
 - Moving away from traditional Tape to a more “always on Disk” technology.
 - Charge backs are very economical ...slightly higher than Tape.
- Elastic Cloud Storage
 - It is intended for storing large amounts of data that is infrequently accessed.
 - Not high performing
 - Costs around 4.5 cents per gigabyte per month

IBM - DB2 RFE (Request For Enhancement)

- There is an effort underway by one of our ex-teammate to work with IBM towards un-interrupted offline reorg during Index Re-build phase
- As a possible solution a RFE (Request For Enhancement) has been submitted to IBM requesting to alter DB2 behavior and let the index build to be run in the background
- Feel free to socialize within your DBA friends circle as well
- https://www.ibm.com/developerworks/rfe/execute?use_case=viewRfe&CR_ID=7709

DB2 (LUW) Process Improvements

Data / Database as a Service (Auto build for Cloud)

Objective: Create a new relational database build with single push of a button

High level process : Linux, Unix and windows

- 1) VM build done using Ansible
 - Allocates required memory, storage and CPU
- 2) OS build done using Ansible ,Puppet or python
 - Installs required plugins,
 - Creates file systems and defines storage limits,
 - LDAP set-up and authentication
- 3) Database build done using Shell, Power shell or Ansible
 - Database software/binary installation
 - Instance and Database creation
 - Authentication, grants and privilege allocations
 - High availability enablement
- 4) Disaster Recovery abilities

Steps 1 thru 4 can be executed with dependency via a orchestration tool with required error codes .

Auto builds have been successful both in RDBMS and NoSQL

DB2 (LUW)

Third Party Tools : Best industry practices

- IRM (Internal Risk Management) & Audit tools
 - Guardium,Qualys,TripWire, Lieberman & IIQ
 - DB2 Native logging
 - Encryption
- Guardium – DB2 :
 - STAP, KTAP Preparation OS level
 - ATAP at instance level
 - Memory configuration OS and Instance, things to keep an eye on during drop and recreation of instance
 - Including Guardium with DB2 upgrades
- Tripwire – DB2 :
 - General topic on what can be done to automate things that are deemed as manual audit controls
 - Industry standard and best practices
- Qualys:
 - If using this tool how often do you get reports and nature of vulnerability
 - Open APAR'S Vs Vulnerability time –line
- Lieberman & IIQ – DB2 : Anyone using ...if so open discussion around best practices
- DB2 Native logging & Encryption
 - Anyone using ... if so what are the best practices to deal with growing data and performance related situations.

Questions ?