

A Crash Course on Db2 pureScale and Advantages

Kent Collins

BNSF Railway

Session code: D03

Monday April 30, 2018 03:20 PM - 04:20 PM

Db2 for LUW - II

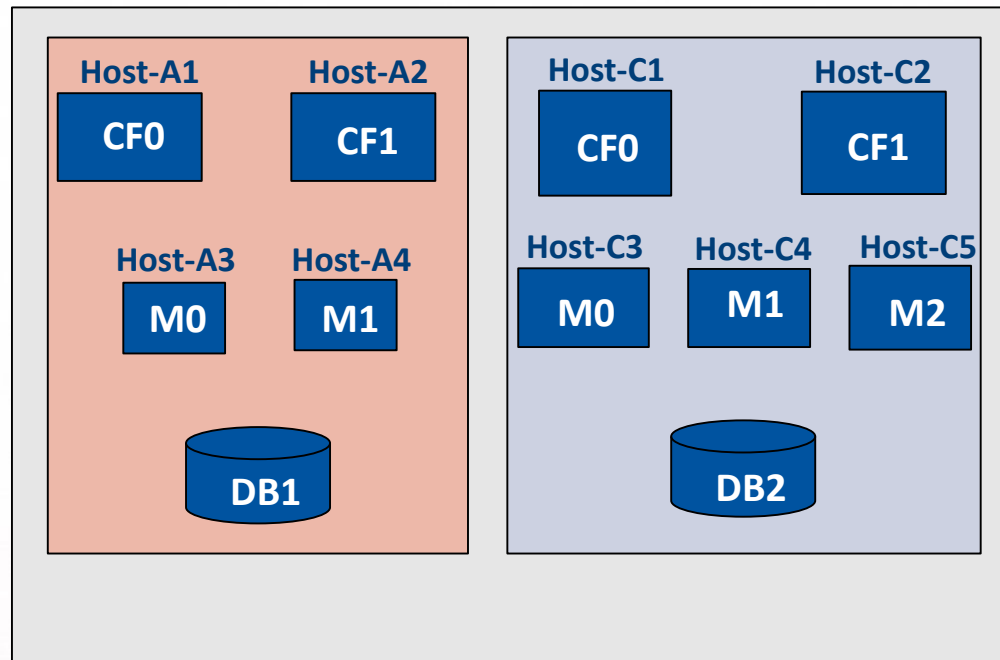
Agenda

- Db2 pureScale Landscapes
 - Virtual | Physical
- pureScale Deployment
 - Single-Tenant | Multi-Tenant
- pureScale Best Practices
 - RASP – Reliability | Availability | Serviceability | Performance
 - Dynamic Partition Management (DPM)
- pureScale HA/DR Setup/Configuration
- pureScale Workload Management

pureScale Linux Virtual Cluster Single Physical Server – Non-Prod Only

- Install and setup VMware ESX/ESXi 5.0
- Define a Host-only IP Address – Private Network between Members and CFs
- Define a Shared Host IP Address for IN/OUT Public communications
- Define a Member 0 guest VM
- Assign Virtual Disks to Member 0
- Install Linux OS on Member 0
- Clone Member 0 to create other members and two CFs. All must be on the same Physical Server
- Install PureScale and create the Db2 instances on Members and CFs
 - db2_setup or db2_install
- Create or backup and restore an existing database

PureScale Linux Virtual Cluster - On Single Physical Server Non-Prod Only




Physical Server

- 2 Network Adapters
- Local Disk or SAN Export as Virtual
- One Db2 Instance per VM
- Multi-Tenant

VM Mobility

Stop Member
Host into Maint mode
Move VM
Host into normal mode
Start Member



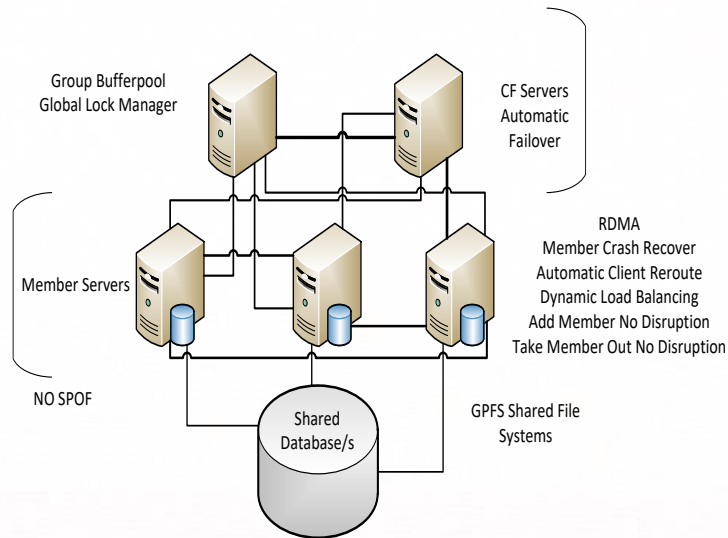
KVM or VMWare
Live Move Not Supported
but not needed for a
cluster

TCP/IP Transport
SAN Storage virtual disks
(VMFS or RDM)

Virtual NICs

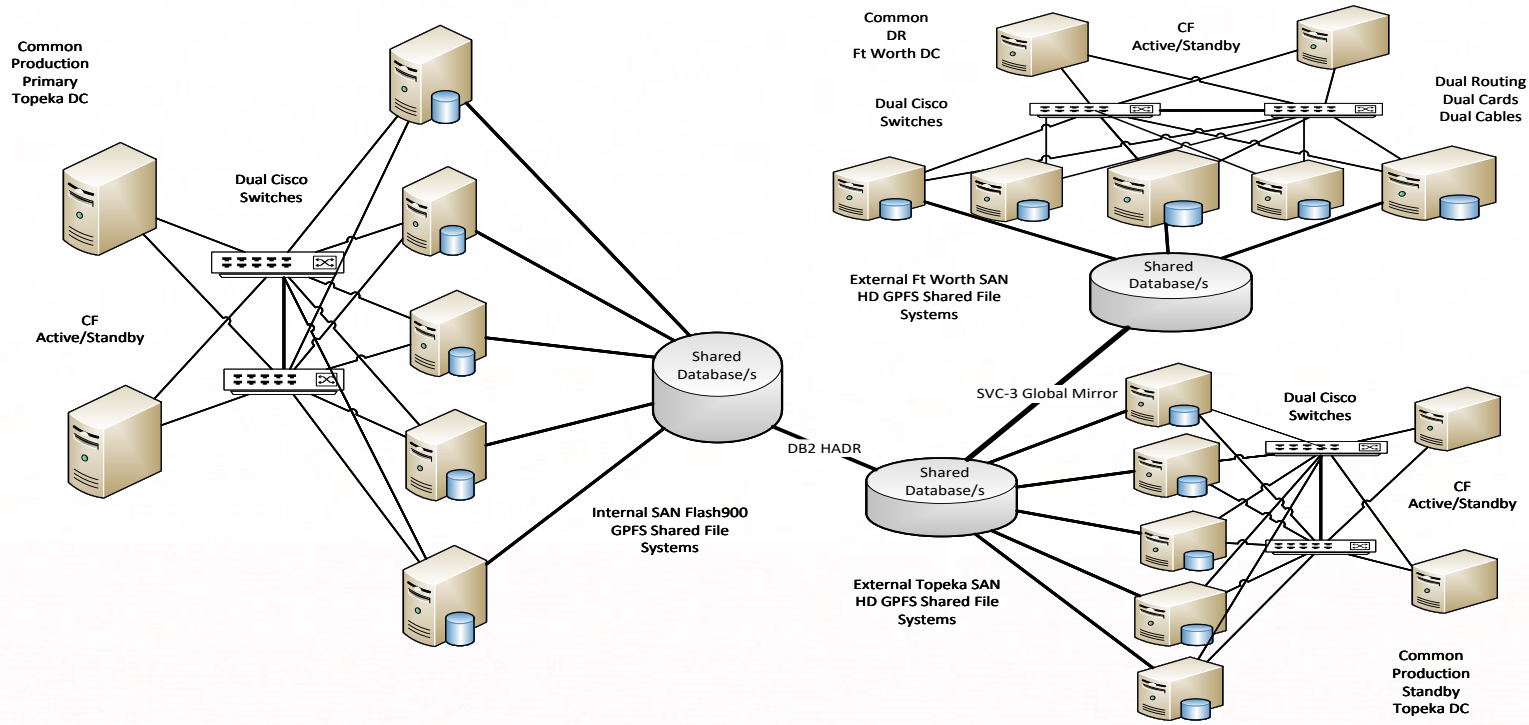
Within Hypervisor
Requirements

PureScale Physical Cluster

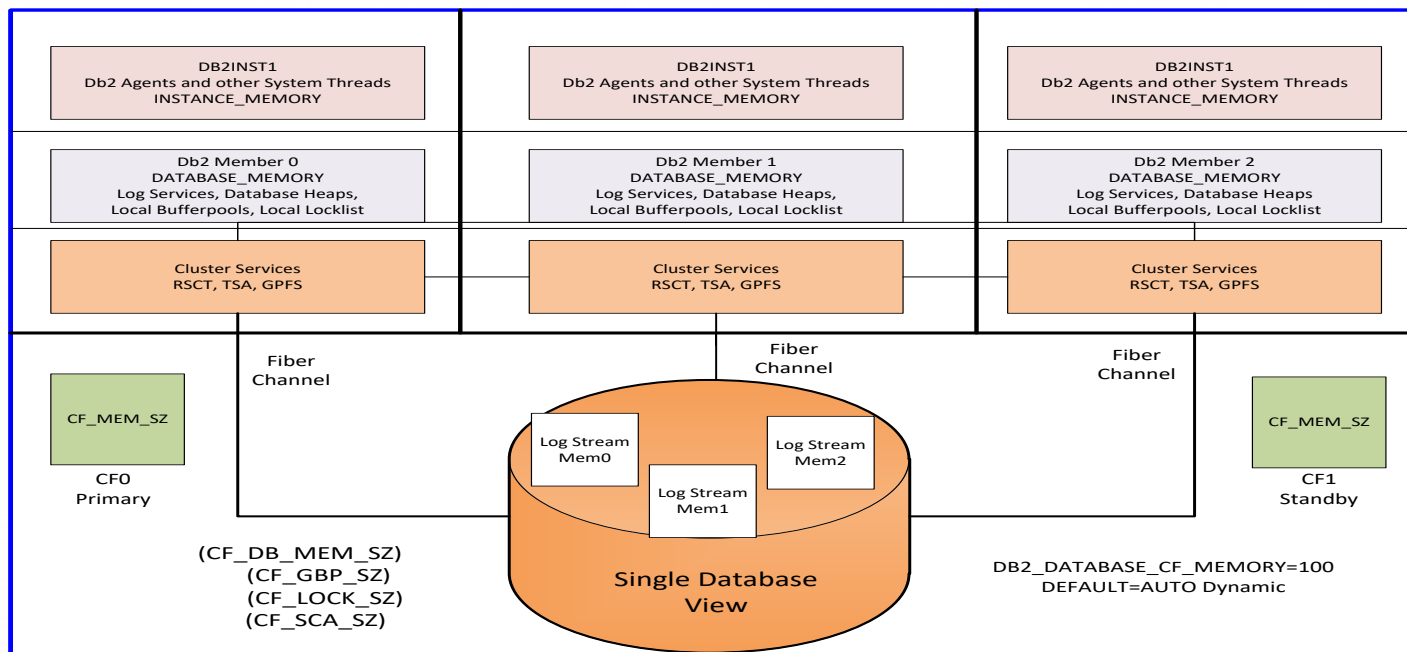


- Endless Scalability
- High Availability – Db2 Rolling FP Upgrade
- DR
- NO SPOF
- Multi-Tenant
- Compression
- Encryption
- JSON
- pureXML
- GIS
- RESTful Services
<https://www.ibm.com/developerworks/data/library/techarticle/dm-1105httprestdb2/index.html>
- LPAR Supported
- Text Search
- Hybrid Cloud Enabled
- Top Shelf Security
- Multi-Model
 - Graph, Document, Key-Value, Relational
- ACID
- Top Shelf DB on Unix or Windows
- Lowest TCO
- Data Governance
- IBM Supported: 10 years lead technology patents
- TimeSeries – Temporal Tables
- Multi-Dimensional Tables
- AutoTiering Storage Management
- Shadow Tables

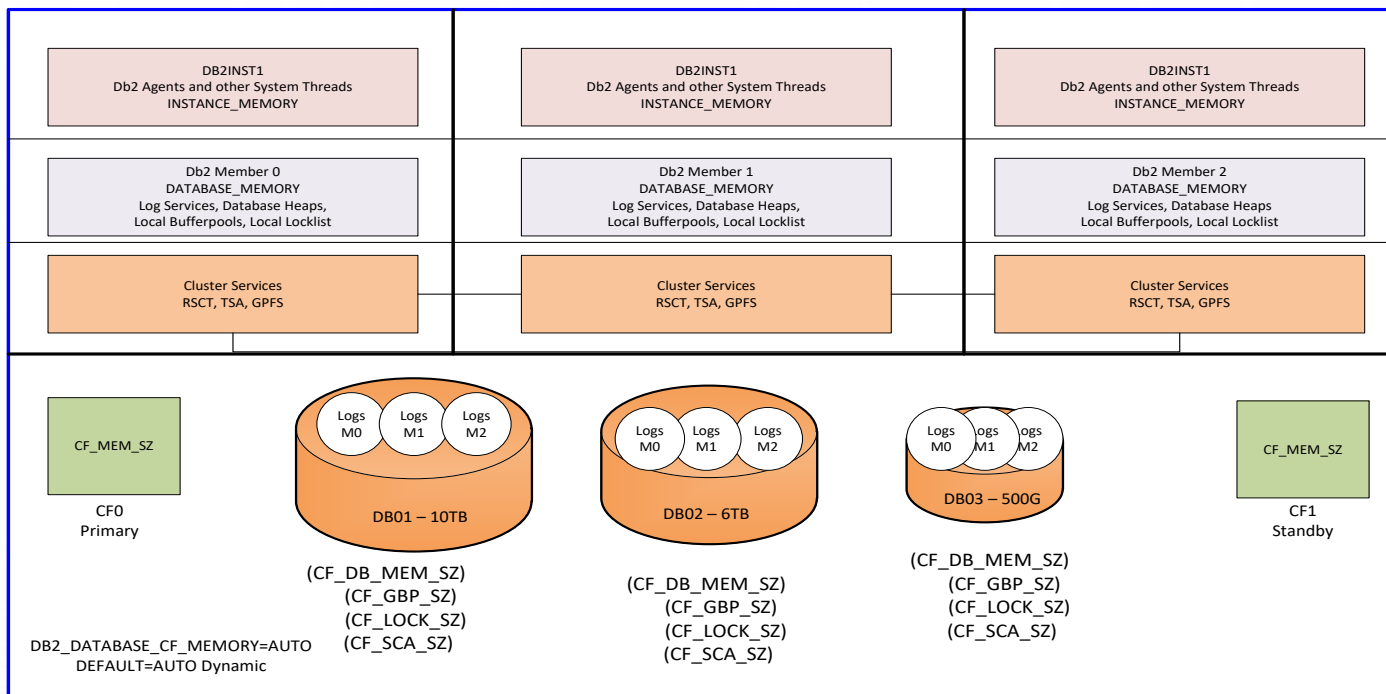
PureScale Physical Cluster – Top Shelf



pureScale Single-Tenant Database



pureScale Multi-Tenant Databases



pureScale Best Practices – Hardware

- Power8
 - OpenPower; AI, Data and Performance Workloads
 - Hypervisor is built into the card not a third party add-on
 - 4Ghz; Large L3 Cache
 - Write Heavy Loads $\text{sum}(\text{members}) / 6$
 - Read Heavy Loads $\text{sum}(\text{members}) / 12$
 - Fastest Cores and Memory (Runs in CPU loop)
 - Turn off all High CPU Monitors
- CF RAM
 - $(\text{sum}(\text{all LBP}(4\text{K pages}) \text{ across all Members and databases }) * (.25 - .40)) = (.65 * (\text{CF RAM} * .90)$
 $(1067008 * 5) * .40 = 2134016(4\text{K}) \text{ or } 8,336\text{MB}$ CF RAM should be set to 15G`
- Switch Configuration and Features
 - Very important to setup switch correctly
- Physical Memory for Members
 - Set Per Core amount based on size of databases,
 - Single or Multi-Tenant, Transaction type and R/W ratio
- Keep all Member Servers the same size and configuration
- Keep all CF Servers the same size and configuration
- IVP All Hardware before running live load
- Separate Racks and Power Supplies
- Setup Internal Server Communications so DSN resolves local /etc/hosts on all servers in the cluster

10.100.78.80	axccfp011.iss.bnr.com	axccfp011
10.100.78.81	axccfp012.iss.bnr.com	axccfp012
10.100.78.82	axdb2p031.iss.bnr.com	axdb2p031
10.100.78.83	axdb2p032.iss.bnr.com	axdb2p032
10.100.78.84	axdb2p033.iss.bnr.com	axdb2p033
10.100.78.85	axdb2p034.iss.bnr.com	axdb2p034
10.100.78.86	axdb2p035.iss.bnr.com	axdb2p035

pureScale Best Practices – CF

- DB2_DATABASE_CF_MEMORY=-1 Manage Databases to equal workloads
- Set NUMDB to something reasonable Default is 32. Who has 32 databases in one instance?

CF Size

- ❖ $GBP = (\text{sum}(\text{all LBP across all Members and databases}) * (.25 - .40))$
Based on R/W ratio with Reads lower percentage

- CF_MEM_SZ = FIXED (45822976) Set the CF size to 90% of RAM on CF Servers
- Properly size CF_DB_MEM_SZ Do not undersize or oversize as both are bad.
 (sum(all LBP(4K pages) across all Members for this DB))
 * (.25 - .40) % based on R/W ratio with Read lower = (.65 * CF_DB_MEM_SZ)
- CF_LOCK_SZ 15% of CF_DB_MEM_SZ
- CF_SCA_SZ 20% of CF_DB_MEM_SZ
- Set the settings the same across all Members
- Consider EHL but test

pureScale Best Practices – DBM/Registry Settings

- **Turn on monitor settings**
- **Fix the number of agents and connections**
- **Fix the number of initial agents**
- **Automatic the number of pool agents**
- **SHEAPTHRES=0**
- **DFTDBPATH** < Not the instance Share > **GPFS FS like /db/home**
- **DIAGPATH** on it's own GPFS File System
- **CF_DIAGPATH** same as **DIAGPATH**
- **INSTANCE_MEMORY** Fixed to **90-95%** of RAM on Member

- **DB2_SKIPDELETED=ON**
- **DB2_SKIPINSERTED=ON**
- **DB2_USE_ALTERNATE_PAGE_CLEANING=ON**
- **DB2_HISTORY_FILTER=T,U,G,L**
- **DB2_LIMIT_FENCED_GROUP=YES**
- **DB2_RESTORE_GRANT_ADMIN_AUTHORITIES=YES**

pureScale Best Practices – DB Settings

- **EXTENDED_ROW_SZ=DISABLE**
- **OPT_DIRECT_WRKLD=YES** (EHL Perhaps ???)
- **CONNECT_PROC** (To control) Perhaps
- **CF self-tuning memory OFF**
- **LOGINDEXBUILD=ON** (If HADR)
- **BLOCKNONLOGGED=YES** (If HADR)
- **BLK_LOG_DSK_FUL=YES**
- **MAX_LOG < FIXED >**
- **MAXAPPLS=<FIXED> and Monitored**
- **AVG_APPLS=1**
- **MAXFILOP=61440**
- **TRACKMOD=YES**
- **SELF_TUNING_MEM=ON**
- **DATABASE_MEMORY < FIXED SIZE If Multi-Tenant > (75% of INSTANCE_MEMORY / SUM(DATABASE_MEMORY) for Multi-Tenant and 80%+AUTOMATIC for Single Tenant**
- **All other allowable settings AUTO**
- **LOGFILSIZ, LOGPRIMARY, LOGSECOND** (Proper Size)
- **AUTO_DEL_REC_OBJ=ON**
- **NUM_DB_BACKUPS** < what is agreed number of full backups to keep >
- **REC_HIS_RETENTN** < what is agreed to retain history of backups >

pureScale Best Practices – Disk Storage

- Unmount Used by TSA to protect cluster from corruption
- Lsof (list open files)
- df
- smitty fs

- **GPFS**
 - mmfsck
 - mmlsmount /dev/ts01
File system ts01 is mounted on 7 nodes.
 - mmlsnsd - List disk name/File system and status
ts01 gpfs93nsd (directly attached)
 - Mmlsconfig Configuration details
- Tiebreaker disk
- No Thin Provisioning
- No SAN Compression
- No SAN Auto-Tiering

- **TSA**
 - /usr/sbin/rsct/bin/samversion

- **RSCT**
 - /usr/sbin/rsct/install/bin/ctversion -b

Configuration data for cluster
db2cluster_20161014115623.iss.bnr.com:

```
-----
clusterName db2cluster_20161014115623.iss.bnr.com
clusterId 17443359407541915194
dmapiFileHandleSize 32
minReleaseLevel 4.1.1.4
maxFilesToCache 10000
verifyGpfsReady yes
worker1Threads 512
sharedMemLimit 2047M
usePersistentReserve yes
failureDetectionTime 35
leaseRecoveryWait 35
totalPingTimeout 45
tiebreakerDisks gpfs1nsd
pagepool 512M
autoload yes
adminMode allToAll
```

rsa41svcs003f 4.1.0.3 Apr 26 2017 11:58:22

```
RSCT_Build_Name=rootas001a 3.2.1.11
RSCT_Build_Time=16237.18:05:29
RSCT_Build_Context=power
```


pureScale Best Practices – Storage Considerations

- db2cluster -cfs -list -tiebreaker

The current quorum device is of type Disk with the following specifics: /dev/hdisk3

- db2cluster -cfs –mount –filesystem <fsname>
- db2cluster -cfs –unmount –filesystem <fsname>
- db2cluster -cfs –rebalance –filesystem <fsname>

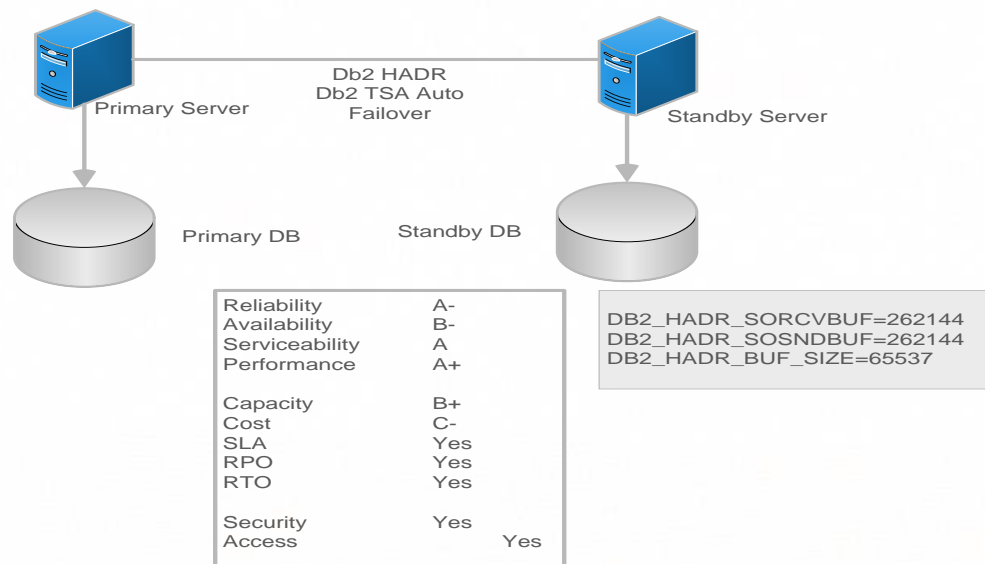
- Active log file system write performance is important – Place on separate/isolated drives – RAID0 or RAID1+0. Solid State or Flash a good option whether local or external storage
- Container file systems equal to the number of Members
- Large File Systems over many disks better than small file system on few disks
- No File System Cache
- Set MAXSIZE
- Set appropriate INITIAL and INCREASE size
- Tell IBM to provide MAX(REDUCE) ASAP
- Large Extents better than small ones
- Monitor CF Waits
- Monitor CPU Usage
- Monitor Read efficiency
- Table Partitioning
 - No Global Indexes

STATEMENT	ROWS_READ	PCT_TOTAL	ROWS_RETURNED	READ_EFFICIENCY	NUM_EXECUTIONS
SELECT A.AUTH_CLNT_ID, B.TAUTH_TOKEN_CLNT_ID, B.TAUTH_SESN_ID, B.SESN_TOKEN_LST_VRFY_TS, B.CLNT_AUTH_DT, B.CLNT_AUTH_TM, B.CLNT_AUTH_TS FROM DA_AUTH_TOKEN.TAUTH_TOKEN_CLNT A, DA_AUTH_TOKEN.TAUTH_SESN B	6517372	0.26	10	651737.20	10
SELECT A.HBD_ALARM_ID, A.EQP_INIT, A.EQP_NBR, A.ALARM_HDLG_UID, A.HBD_EQP_TYP_CD, A.HBD_ALARM_TS, A.HBD_ALARM_CD, A.HBD_CFD_LVL, A.WBES_BRULE_NME, A.TRN_CONCAT_ID, A.BO_STN_333, A.BO_STN_ST, A.MESIBO_CD, A.AT_CMNT, A.TRN_SCH_DPT_DT, A.HBD_AXLE_SEQ, A.TRN_SIDE, A.AXLE_SIDE, E.EQP_HBD_SEQ, A.CAR_ORNTN_CD, A.CRE_ALARM_ID, A.EQP_AXLE_NBR, A.AT_CMNT, A.TWD_NBR, A.TWD_LOC_NME, M.TRN_DIR_CD, M.HBD_MSG_ID, M.AXLE_CNT, M.HBD_RDNG_TS, M.HBD_CFD_LVL, M.TRN_SCH_DPT_DT FROM ME.THBD_MSG_MSTR M, ME.THBD_ALARM A LEFT OUTER JOIN ME.THBD_EQP E ON M.HBD_MSG_ID = A.HBD_MSG_ID AND E.EQP_INIT = A.EQP_INIT AND E.EQP_NBR = A.EQP_NBR	2079246	0.05	4	519811.50	26
SELECT DISTINCT TRN.EDR_MSG_ID, TRN.VNDR_EDR_NME, TRN.VNDR_EDR_TS, TRN.TRN_CONCAT_ID, TRN.TRN_DIR_CD, EQP.EQP_INIT, EQP.EQP_NBR, BRNG.EQP_AXLE_NBR, BRNG.EQP_SEQ_NBR, BRNG.EQP_SIDE, BRNG.FALT_DESC_1, ABD.CRE_ABD_ALARM_CFD_PCT, ABD.CRE_ABD_SVR_LVL_NBR FROM ME.TEDR_SS_TRN_INFO TRN, ME.VEDR_SS_VEH_INFO EQP, ME.VEDR_SS_BRNG_MEAS BRNG, ME.TCRE_ABD_MPNG ABD	1786313	0	3900	458.02	38
select A.SVD_SWPLN_ID, A.STN_333, A.STN_ST, A.SWPLN_NME, A.SWPLN_DESC, A.TRN_TYPE, A.TRN_SYM, A.TRN_SECT, A.TRN_DAY, A.TRN_PRTY, A.TRN_SCH_DPT_DT, A.SWPLN_FIN_CD, A.CRET_TS, A.CRET_USER_ID, A.LST_MAINT_TS, A.LST_MAINT_ID, A.SWPLN_STAT_CD, A.SWPLN_FIN_TS, A.SWPLN_FIN_USER_ID, B.SVD_SWPLN_CLOB_SEQ_NBR, B.SWPLN_LST_DTL_TXT, B.CRET_TS as CLOB_CRET_TS, B.CRET_USER_ID as CLOB_CRET_USER_ID, C.SVD_SWPLN_DTL_SEQ_NBR, C.SWPLN_SRC_CMNT, C.SWPLN_MTY_CAR_CNT, C.SWPLN_ID_CAR_CNT, C.SWPLN_JOB_CNT, C.CRET_TS as DTL_CRET_TS, C.CRET_USER_ID as DTL_CRET_USER_ID from yi.TSVD_SWPLN A, yi.TSVD_SWPLN_CLOB B, yi.TSVD_SWPLN_DTL C where A.STN_333 = ? AND A.STN_ST = ? and A.SVD_SWPLN_ID = B.SVD_SWPLN_ID and B.SVD_SWPLN_CLOB_SEQ_NBR = (SELECT MAX(SVD_SWPLN_CLOB_SEQ_NBR) FROM yi.TSVD_SWPLN_CLOB	497198	0.01	28	17757.07	65

pureScale Best Practices – Application Development Focus

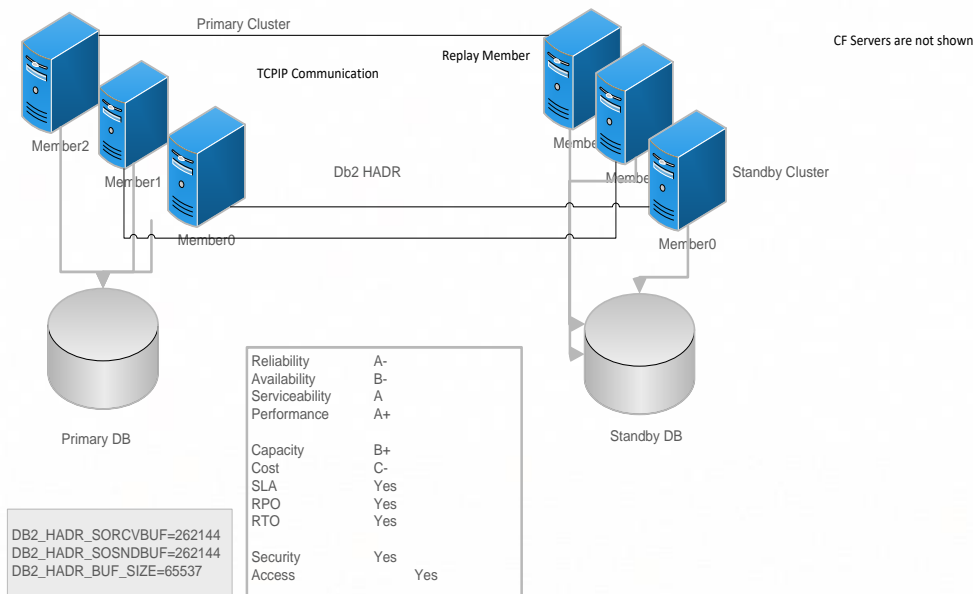
- Unqualified objects
- Bounded transactions
- Standard Error Handle
- Connection Retry Logic (up to 5 mins)
- Required Standard Testing Methodology
- ApplicationName required on all connections
- Planned Service Disruptions under 15 minutes with rolling code deployments
- Database Moves using Replication (Minimal Down)
- Connection Managers are all Sysplex and ACR aware
- Schema backup and restore
- No cross-schema tablespaces
- No cross-schema defined RI
- Eliminate or reduce Federation through database consolidation

ESE verses pureScale using Db2/HADR



Active Log Stream
/ds/activelog/db2inst1/sample/NODE0000/LOGSTREAM0000/

```
HADR database role           = PRIMARY
HADR local host name         (HADR_LOCAL_HOST) = server1
HADR local service name      (HADR_LOCAL_SVC) = HADR_db2inst1_sample
HADR remote host name        (HADR_REMOTE_HOST) = server1
HADR remote service name     (HADR_REMOTE_SVC) = HADR_db2inst1_sample
HADR instance name of remote server (HADR_REMOTE_INST) = db2inst1
HADR timeout value           (HADR_TIMEOUT) = 120
HADR target list              (HADR_TARGET_LIST) =
HADR log write synchronization mode (HADR_SYNCMODE) = NEARSYNC
HADR spool log data limit (4KB) (HADR_SPOOL_LIMIT) = 0
HADR log replay delay (seconds) (HADR_REPLAY_DELAY) = 0
HADR peer window duration (seconds) (HADR_PEER_WINDOW) = 300
```



Active Log Stream
/ds/activelog/db2inst1/sample/NODE0000/LOGSTREAM0000/

```
==> db2 get db cfg for sample | grep -i HADR
HADR database role           = PRIMARY
HADR local host name         (HADR_LOCAL_HOST) = axdb2p031
HADR local service name      (HADR_LOCAL_SVC) = 56002
HADR remote host name        (HADR_REMOTE_HOST) =
{axdb2p051:56002|axdb2p052:56002|axdb2p053:56002}
HADR remote service name     (HADR_REMOTE_SVC) =
{axdb2p051:56002|axdb2p052:56002|axdb2p053:56002}
HADR instance name of remote server (HADR_REMOTE_INST) = db2inst1
HADR timeout value           (HADR_TIMEOUT) = 120
HADR target list              (HADR_TARGET_LIST) =
{axdb2p051:56002|axdb2p052:56002|axdb2p053:56002}
HADR log write synchronization mode (HADR_SYNCMODE) = ASYNC
HADR spool log data limit (4KB) (HADR_SPOOL_LIMIT) = AUTOMATIC(30801920)
HADR log replay delay (seconds) (HADR_REPLAY_DELAY) = 0
HADR peer window duration (seconds) (HADR_PEER_WINDOW) = 0
```

```
==> db2 get db cfg for sample | grep -i hadr
HADR database role           = STANDBY
HADR local host name         (HADR_LOCAL_HOST) = axdb2p051
HADR local service name      (HADR_LOCAL_SVC) = 56002
HADR remote host name        (HADR_REMOTE_HOST) =
{axdb2p031:56002|axdb2p032:56002|axdb2p033:56002}
HADR remote service name     (HADR_REMOTE_SVC) =
{axdb2p031:56002|axdb2p032:56002|axdb2p033:56002}
HADR instance name of remote server (HADR_REMOTE_INST) = db2inst1
HADR timeout value           (HADR_TIMEOUT) = 120
HADR target list              (HADR_TARGET_LIST) =
{axdb2p031:56002|axdb2p032:56002|axdb2p033:56002}
HADR log write synchronization mode (HADR_SYNCMODE) = ASYNC
HADR spool log data limit (4KB) (HADR_SPOOL_LIMIT) = AUTOMATIC(30801920)
HADR log replay delay (seconds) (HADR_REPLAY_DELAY) = 0
HADR peer window duration (seconds) (HADR_PEER_WINDOW) = 0
```

Active Log Streams

Primary Cluster

server1.iss.bnr.com: Path to log files
server2.iss.bnr.com: Path to log files
server3.iss.bnr.com: Path to log files
server4.iss.bnr.com: Path to log files
server5.iss.bnr.com: Path to log files

= /ds/activelog/db2inst1/sample/NODE0000/LOGSTREAM0000/
= /ds/activelog/db2inst1/sample/NODE0000/LOGSTREAM0001/
= /ds/activelog/db2inst1/sample/NODE0000/LOGSTREAM0002/
= /ds/activelog/db2inst1/sample/NODE0000/LOGSTREAM0003/
= /ds/activelog/db2inst1/sample/NODE0000/LOGSTREAM0004/

Standby Cluster

serverb1.iss.bnr.com: Path to log files
serverb2.iss.bnr.com: Path to log files
serverb3.iss.bnr.com: Path to log files
serverb4.iss.bnr.com: Path to log files
serverb5.iss.bnr.com: Path to log files

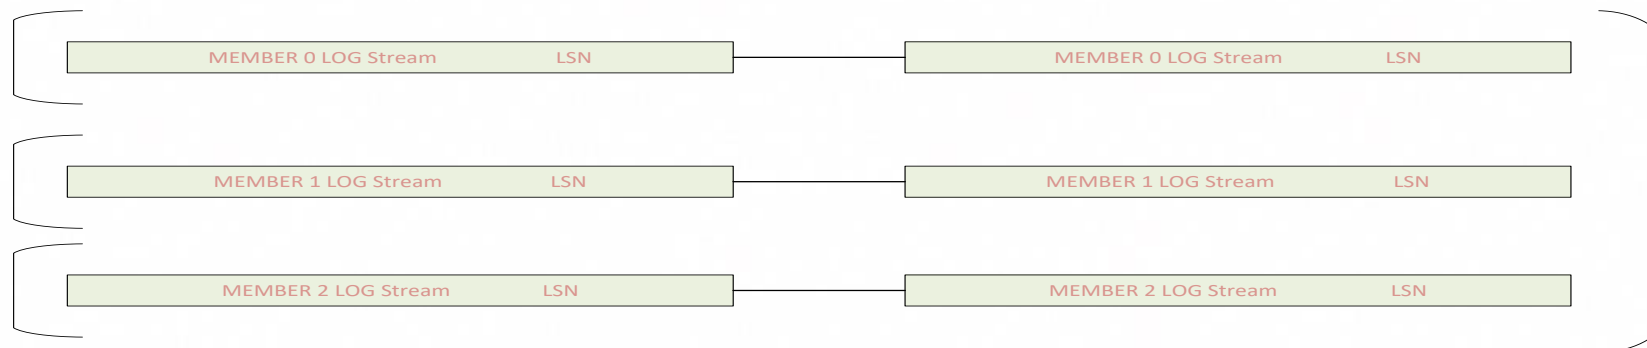
= /ds/activelog/db2inst1/sample/NODE0000/LOGSTREAM0000/
= /ds/activelog/db2inst1/sample/NODE0000/LOGSTREAM0001/
= /ds/activelog/db2inst1/sample/NODE0000/LOGSTREAM0002/
= /ds/activelog/db2inst1/sample/NODE0000/LOGSTREAM0003/
= /ds/activelog/db2inst1/sample/NODE0000/LOGSTREAM0004/

Command (db2_all ";db2 get db cfg for SAMPLE | grep -i 'Path to log files'")

Active Log Streams

Primary Cluster

Standby Cluster



Only 1 Preferred Replay Member on Standby
All other active Members can be selected as standby
On failure replay automatically moves to a active member



Preferred Replay
Member 0

Merge Replay Log of all Members LRI

On Standby

7(*)	N/A	N/A	N/A Forward[*] : 5.9T	Rollforward	6508387646282
------	-----	-----	-----------------------	-------------	---------------

```
==> db2pd -hadr -db mep05000 -allmembers | grep -i LOG_TIME
PRIMARY_LOG_TIME = 04/30/2018 12:12:32.000000 (1525108352)
STANDBY_LOG_TIME = 04/30/2018 12:12:32.000000 (1525108352)
STANDBY_REPLAY_LOG_TIME = 04/30/2018 12:12:32.000000 (1525108352)
PRIMARY_LOG_TIME = 04/30/2018 12:12:32.000000 (1525108352)
STANDBY_LOG_TIME = 04/30/2018 12:12:32.000000 (1525108352)
STANDBY_REPLAY_LOG_TIME = 04/30/2018 12:12:32.000000 (1525108352)
PRIMARY_LOG_TIME = 04/30/2018 12:12:32.000000 (1525108352)
STANDBY_LOG_TIME = 04/30/2018 12:12:32.000000 (1525108352)
STANDBY_REPLAY_LOG_TIME = 04/30/2018 12:12:32.000000 (1525108352)
PRIMARY_LOG_TIME = 04/30/2018 12:12:32.000000 (1525108352)
STANDBY_LOG_TIME = 04/30/2018 12:12:32.000000 (1525108352)
STANDBY_REPLAY_LOG_TIME = 04/30/2018 12:12:32.000000 (1525108352)
PRIMARY_LOG_TIME = 04/30/2018 12:12:32.000000 (1525108352)
STANDBY_LOG_TIME = 04/30/2018 12:12:32.000000 (1525108352)
STANDBY_REPLAY_LOG_TIME = 04/30/2018 12:12:32.000000 (1525108352)
```

pureScale Workload Management - Planning

- Cluster
- Database
- Application
- Transaction

Levels

- Action
- Minimize Outage
- Prevent the Unexpected
- Identify Workload Changes
- Be in Control
- Decisions are Business determined

Objectives

pureScale Workload Management - Goals

- SLO
 - Detector Messages < 5 minutes
- RPO/RTO
 - No Data Loss
 - < 20 minute recovery
- 99.997% Operational Efficiency
- Business Impact
 - Critical
 - Equipment Operations (24/7)
 - Shipment (24/7)
 - Serious
 - Payroll (End of Month)
 - Exec Reports (M-F 8-9AM)
 - Normal
 - Everything else

pureScale Workload Management - Cluster

- No Affinization
 - All Members are available to run all workloads
- Workloads Find Resources
 - ACR - Db2 Client V9.5 FP4 or Higher
 - WLB - Connection and Transaction Load Balancing
 - db2dsdriver.cfg
- Workloads must meet specific deployment requirements
 - Auto-Reconnect
 - Error Handling
 - Testing
 - Security

pureScale Workload Management - Database

- Availability
 - HA/DR
- Security
 - Audit
 - SOX
- Logging Capacity
 - Active logging demand
 - Blocknonlogged=Yes
 - LogIndexBuild=ON
- Federation
 - Keeping objects within the DB to avoid federation
- Schema Controls
 - Minimize Cross-Schema References
 - Schema Transport (Future)
 - Private Tablespaces
- Service Levels
 - Scheduled Outages
 - Critical Resource Management
- Recovery Importance
 - Order of Recovery

pureScale Workload Management - Application

- ApplicationName
 - Core verses Everything else
 - JDBC Client Property
 - MicroServices gone wild
 - db2jcc_applica (worthless)
 - DataFlowEngine (Good)
- Connection ID
 - Assigned Application ID
 - Application Tier
 - "ss1usr" Store Sales Tier 1 User
 - "sched2adm" Schedule Tier 2 Admin
- CONNECT_PROC
 - Set session requirements for WLM
- Standard Error Handling
- Coded for HA/DR if required
- Standard Release/Deployment Policies

pureScale Workload Management - Transaction

- **Identification**

- Application Name
- Connection ID
- Combination of other characteristics
- Workloads/Work Classes/Data Tags

- **Management**

- Service Classes
- Workloads
- Work Action Set
- Thresholds
- Move work between service classes
- CPU Resources by Service Class

- **Monitoring**

- Report/No Action
 - Threshold Assessment
- Periodic Reporting
- Alert Notifications
- Preventative Actions
- Periodic Re-Assessment/Alteration

pureScale WLM – Strategy

Protect the important business workloads

Guarantee Cluster Operation

Guarantee Instance and Database Operation

Monitor resource utilization and provide reports to AppDev

Order Capacity with enough lead time

Audit unplanned resource consumption

Forecast Projections based on Actual usage

Keep it simple and serviceable



IDUG
Leading the DB2 User
Community since 1988

IDUG Db2 Tech Conference NA
Philadelphia, PA | April 29 - May 3, 2018

 **#IDUGDb2**

A Crash Course on Db2 pureScale and Advantages **Kent Collins**

robert.collins@bnsf.com

Mohan Saraswatupura

Mohankumara.Saraswatipura@BNSF.com

www.EnterpriseDBA.com

Session code: D03

*Please fill out your session
evaluation before leaving!* session evaluation
leaving!

