

Getting Started with DB2 pureScale

CODUG – May 19th, 2015



Kelly Schlamb

Executive IT Specialist, Worldwide Analytics Platform Technical Sales
IBM Canada Ltd.

Email: kschlamb@ca.ibm.com **Twitter:** @KSchlamb

Disclaimer

© Copyright IBM Corporation 2015. All rights reserved.

U.S. Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED “AS IS” WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS AND/OR SOFTWARE.

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.

IBM, the IBM logo, ibm.com, Information Management, DB2, DB2 Connect, DB2 OLAP Server, pureScale, System Z, Cognos, Informix, Optim, InfoSphere, and z/OS are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at www.ibm.com/legal/copytrade.shtml

Other company, product, or service names may be trademarks or service marks of others.

The Next Generation Data Platform

Transact

- Provide high levels of service without the high price
- Seamlessly expand or contract as needed, paying only for what you use when you use it
- Enhance customer experience by delivering data when and where it's needed

Analyze

- Leverage next generation in-memory technology
- Get instant insight into operational and warehouse data without compromising performance of either
- Deploy analytics solutions faster with “load and go” and operational simplicity

Innovate

- Free your applications from database complexity for faster delivery
- Transform your ability to make business decisions with 35x to 73x faster analytics, with some queries running more than 1400x faster^{1,2}
- Optimize IT resources and utilization with built-in simplicity and autonomies

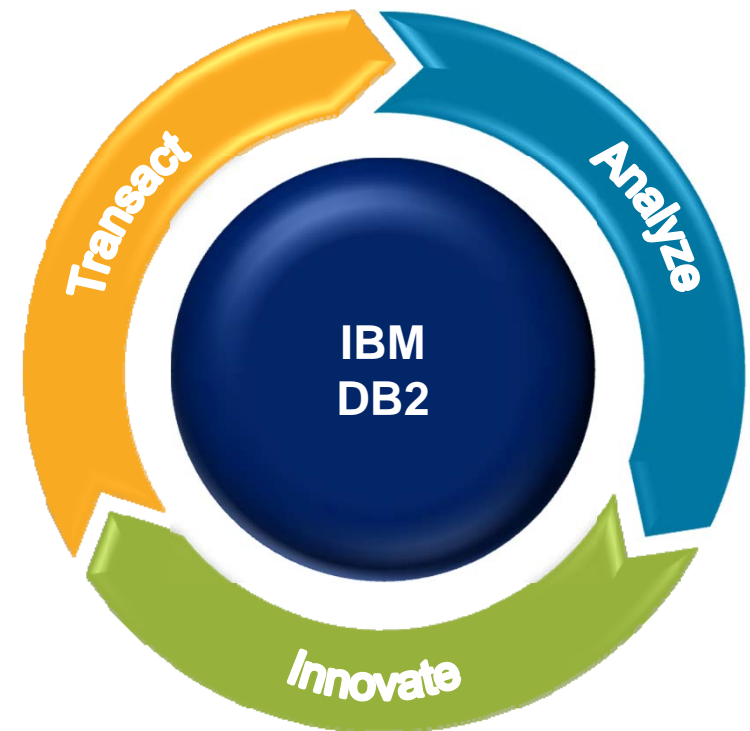


Introducing DB2 "Cancun Release"

New

Get the answers you need in the moment

- **Get real-time answers to ALL questions with instant insight into historical and “as it happens” data**
- **Change the economics of continuous availability** with broad infrastructure choices at every price point
- **Improve performance of transactions and analytics** while reducing complexity and overhead
- **Get fast time-to-value** using skills you already have for Oracle database
- **Superior performance at lower cost** for SAP environments



The Benefits of DB2 10.5 for Business Operations

Fast Transactions You Can Depend On

Deliver superior service at a price you can afford

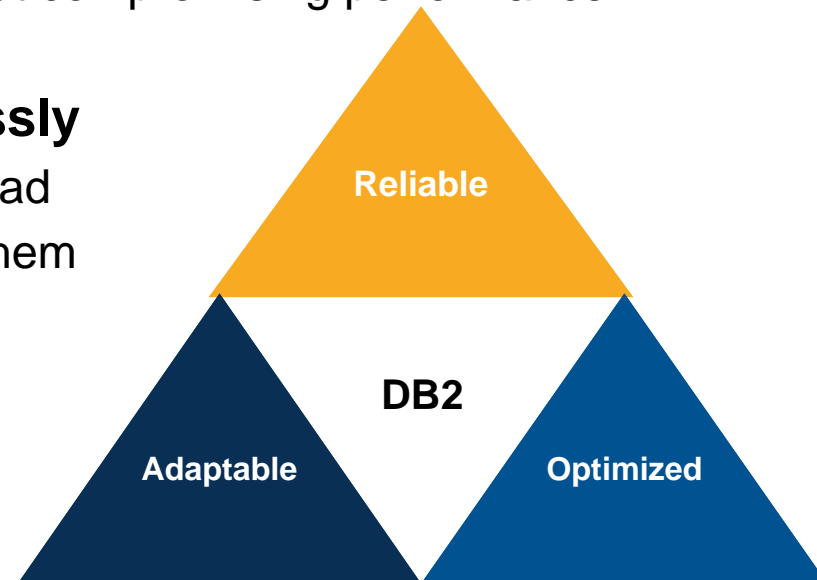
- Protect your business from planned and unplanned outages*
- **Change the economics of continuous availability** with options at every price point
- **Get real-time insight into operational data** without compromising performance

Adapt to changing business needs, seamlessly

- Expand and contract based on transactional workload
- Pay only for the resources you use when you use them
- **Leverage broader infrastructure choices**

Streamline operations

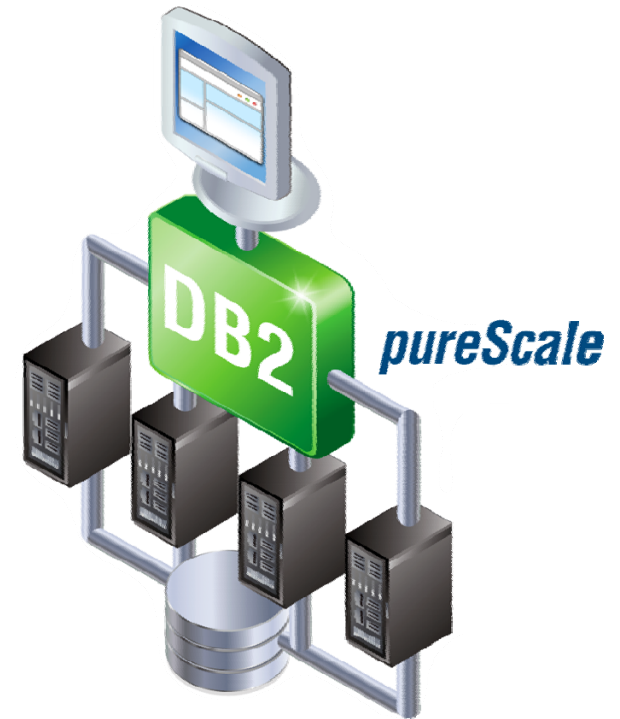
- Free your applications from database complexity
- Improve IT productivity with autonomies
- Modernize your data platform with existing skills and IT resources



* Based on IBM design for normal operation with rolling maintenance updates of DB2 server software on a pureScale cluster. Individual results will vary depending on individual workloads, configurations and conditions, network availability and bandwidth.

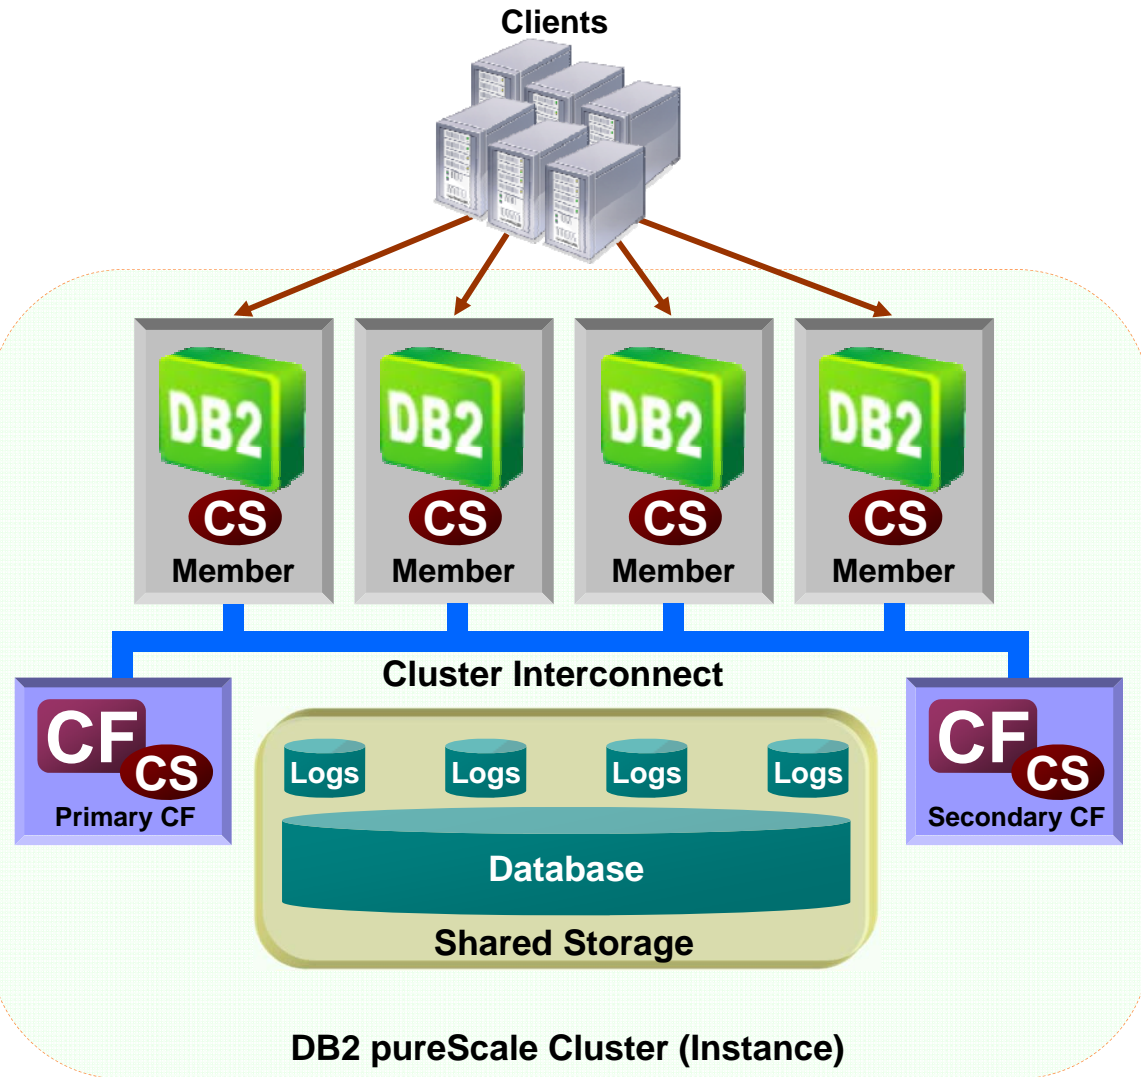
Getting Started with DB2 pureScale – Agenda

- A quick overview of the DB2 pureScale architecture
- Hardware and software requirements
- Client experiences and typical adoption paths
- Moving to DB2 pureScale
- Interesting things to test and try
- Additional reading material and getting started




DB2 pureScale Architecture

Leveraging IBM's System z Sysplex Experience and Know-How





- Multiple **DB2 members** for scalable and available database environment
- Client application connects into any **DB2 member** to execute transactions
 - Automatic workload balancing
- Shared storage for database data and transaction logs
- **Cluster caching facilities (CF)** provide centralized global locking and page cache management for highest levels of availability and scalability
 - Duplexed, for no single point of failure
- High speed, low latency **interconnect** for efficient and scalable communication between members and CFs
- **DB2 Cluster Services** provides integrated failure detection, recovery automation and the clustered file system

DB2 vs. DPF (Shared Nothing) vs. pureScale (Shared Data)



DB2 with Database Partitioning Feature

Ideal for warehousing and OLAP scale out and massively parallel query processing



DB2 pureScale Data Sharing

Ideal for active/active OLTP/ERP scale out

Machine Deployment Examples


- **Highly flexible topologies due to logical nature of member and CF**
 - A member and CF can share the same machine
 - For AIX, separate members and CFs in different LPARs
 - Virtualized environments via VMware and KVM

- **Dedicated cores for CFs**
 - Optimizes response time


- **No pureScale licenses required for CF hosts**
 - You only need to license the CPUs for hosts on which members are running

- **pureScale included in**
 - DB2 Advanced Workgroup Server Edition
 - DB2 Advanced Enterprise Server Edition
 - DB2 Developer Edition
 - DB2 Business Application Continuity Offering


Two Machines



Four Machines




Ten Machines



pureScale Client Configuration


▪ Workload Balancing (WLB)

- Application requests balanced across all members or subsets of members
- Takes server load of members into consideration
- Connection-level or transaction-level balancing




▪ Client Affinity

- Direct different groups of clients or workloads to specific members in the cluster
- Consolidate separate workloads/applications on same database infrastructure
- Define list of members for failover purposes




▪ Automatic Client Reroute (ACR)

- Client automatically connected to healthy member in case of member failure
- May be seamless in that no error messages returned to client
- Application may have to re-execute the transaction

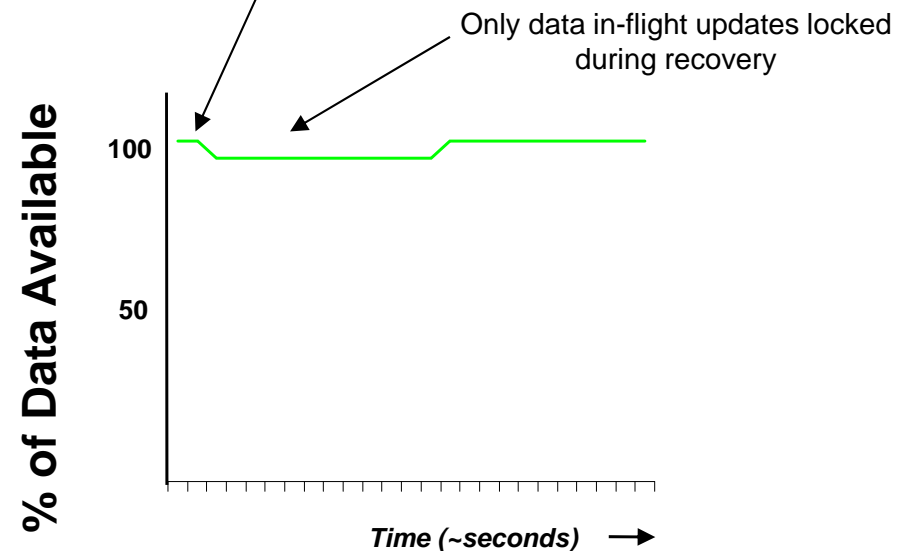


Online Recovery from Failures

- DB2 pureScale design point is to **maximize availability during failure** recovery processing
- When a database member fails, **only in-flight data remains locked** until member recovery completes
 - In-flight = data being updated on the failed member at the time it failed
- **Target time to availability** of rows associated with in-flight updates on failed member **in seconds**



Database member failure



"We pulled cards, we powered off systems, we uninstalled devices, we did everything we could do to make the cluster go out of service, and we couldn't make it happen."


-- Robert M. Collins Jr. (Kent), Database Engineer, BNSF Railway Inc.

Scale with Ease

- **Scale up or out... without changing your applications**
 - Efficient coherency protocols designed to scale without application changes
 - Applications automatically and transparently workload balanced across members
 - Up to 128 members

- **Without impacting availability**
 - Members can be added while cluster remains online

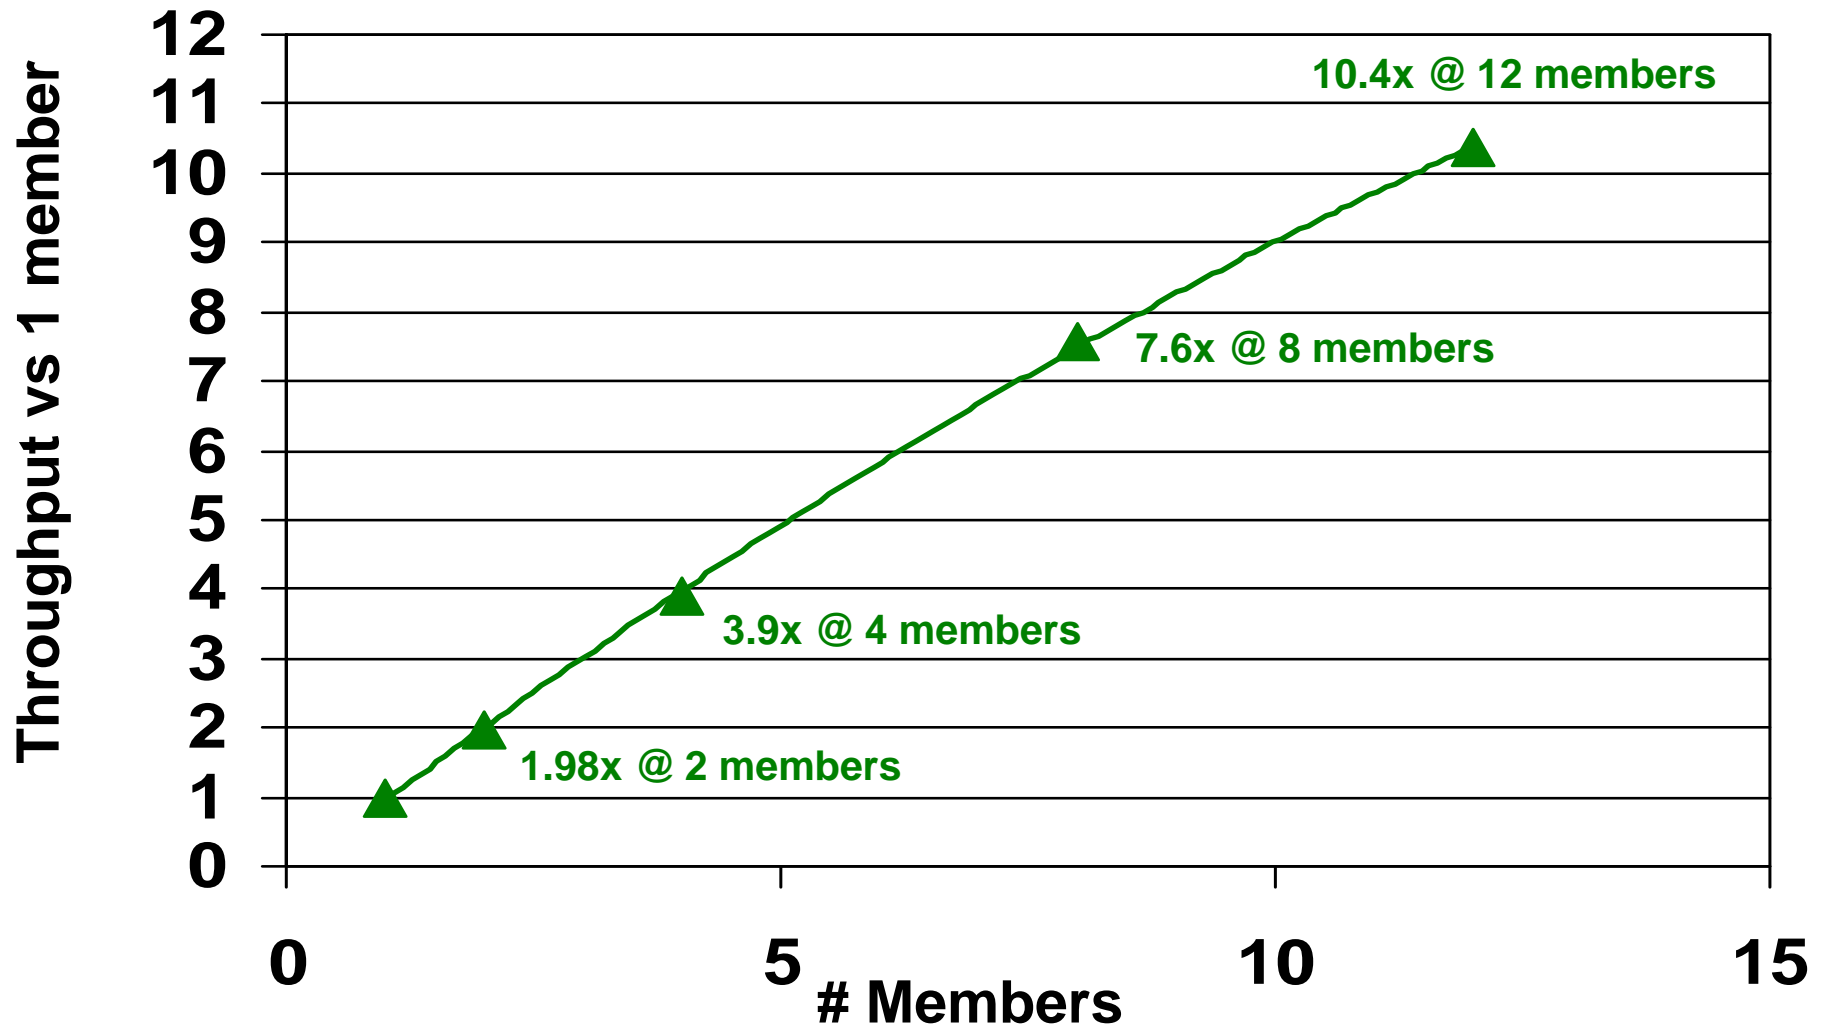
- **Without administrative complexity**
 - No data redistribution required



*“DB2 pureScale is the only solution we found that **provided near linear scalability... It scales 100 percent**, which means when I add servers and resources to the cluster, I get 100 percent of the benefit. Before, we had to ‘oversize’ our servers, and used only 50 - 60 percent of the available capacity so we could scale them when we needed.”*

-- Robert M. Collins Jr. (Kent), Database Engineer, BNSF Railway Inc.

Scalability Demonstration



OLTP 80/20 R/W workload
No affinity

12 8-core p550 members
64 GB, 5 GHz each


Duplexed CFs
on 2 additional 8-core p550s
64 GB, 5 GHz each

20Gb/s IB HCAs
7874-024 IB Switch

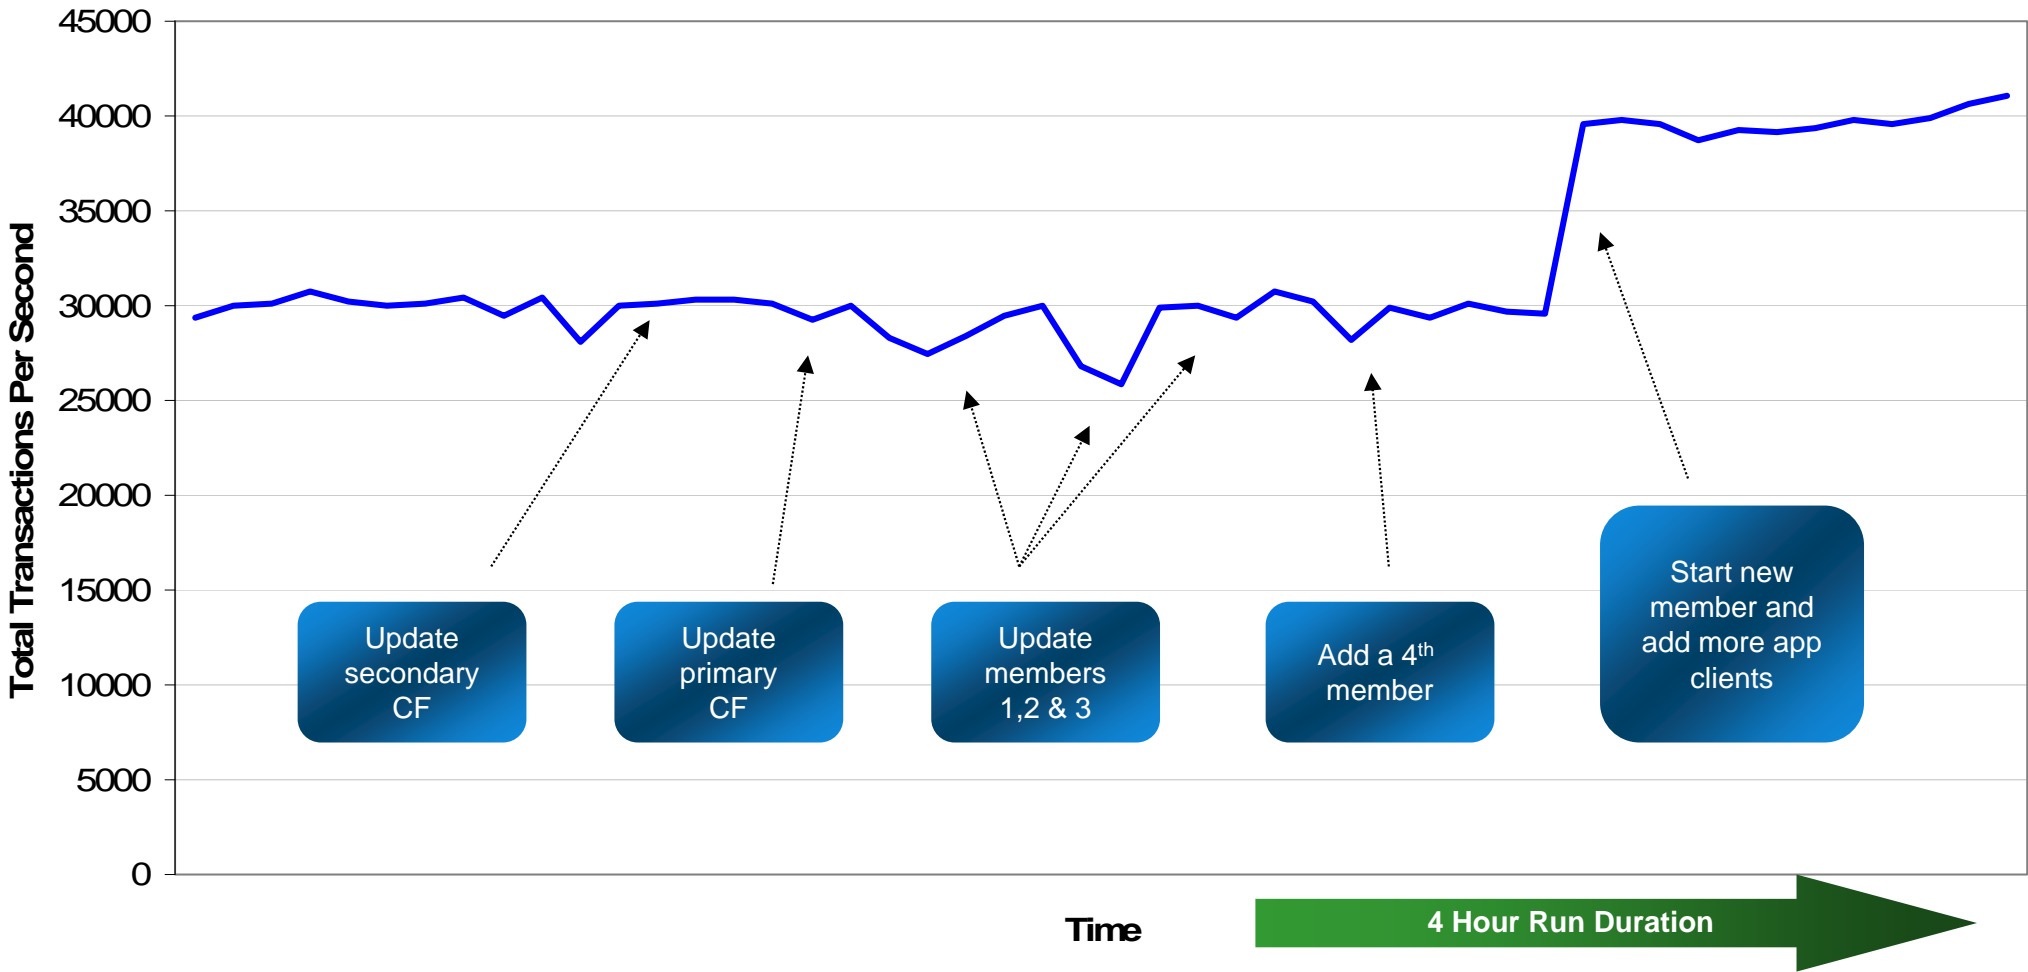
DS8300 storage
576 15K disks
Two 4Gb FC Switches

Online System and Database Maintenance

- **Transparently perform maintenance to the cluster in an online rolling fashion**
 - DB2 pureScale fix packs
 - System updates such as operating system fixes, firmware updates, etc.
- **No outage experienced by applications**
- **DB2 fix pack install involves a single `installFixPack` command to be run on each member/CF**
 - Quiesces member
 - Existing transactions allowed to finish
 - New transactions sent to other members
 - Installs binaries
 - Updates instance
 - Member still behaves as if running on previous fix pack level
 - Unquiesces member
- **Followed up by final cluster-wide `installFixPack` command to complete and commit updates across cluster**
 - Instance now running at new fix pack level



Continuous Availability During Maintenance and System Growth



Database servers

- SUSE Linux Enterprise Server 11 SP 1
- 6 - IBM x3950 X5s (Intel XEON X7560 @ 2.27GHz (4s/32c/64t))
- Mellanox ConnectX-2 IB Card
- 128GB system memory



Storage server

- 1 - IBM Storwize v7000
- 8 - SSD drives (2TB usable capacity)
4 - for data, 4 - for logs

Main thing to remember... pureScale is DB2




- **A pureScale environment looks and feels very much like a "regular" DB2 environment**
 - Same code base shared by DB2, DPF, and pureScale
 - In DB2 10.1 and 10.5, pureScale is just an installable feature of DB2

- **Immediate productivity from DBAs and application developers**
 - Single system view for utilities
 - Act and behave exactly like they do in non-pureScale
 - Backup, restore, rollforward, reorg, load, ...
 - Applications don't need to know about or care about the fact that are multiple members
 - In general, can run SQL statements or command on any member
 - SQL, data access methods, and isolation levels are the same
 - Backup/recovery processes are the same
 - Security is managed in the same way
 - Environment (even the CFs) still managed by database manager and database configuration parameters

What in pureScale is new to a DB2 DBA?

- **CFs provide centralized locking and cache management**
 - Global Lock Manager and Group Buffer Pool
 - Automatic self-tuning memory management for CF
- **Other components automatically installed and configured as part of pureScale install**
 - RSCT provides heart beating, domain management
 - TSA defines resources and dependencies, drives recovery
 - GPFS provides a clustered file system on which shared data resides
- **Online, automated member/CF recovery**
- **Default Automated workload balancing across all members**
- **Multiple transaction log streams are maintained, one per member**
 - Automatically merged when necessary (e.g. rollforward, group crash recovery)
- **Storage management primarily done through DB2 interfaces to GPFS**
- **Cluster management done through DB2 interfaces to RSCT/TSA**
- **Performance and status monitoring includes both members and CFs**



DB2 pureScale Supported Hardware and OS



**POWER6
POWER7/7+
POWER8
Flex**



OR



**x86 Intel
Compatible Servers**

Flex



**BladeCenter
H22/HS23**



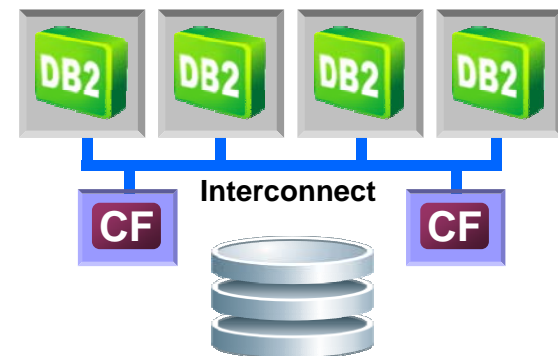
TCP/IP sockets
interconnect or high
speed, low latency
RDMA-based interconnect
(InfiniBand, 10 GE (RoCE))



GPFS compatible storage
*(ideally storage that supports
SCSI-3 PR fast I/O fencing)*

Cluster Interconnect Options

- **RDMA-capable interconnect for best performance and scalability**
 - Requires specialized network adapter cards
 - InfiniBand
 - 10 Gigabit Ethernet RoCE (RDMA over Converge Ethernet)
- **TCP/IP sockets interconnect for faster cluster setup and lower cost deployments using commodity network hardware**
 - 10 Gigabit Ethernet (10GE) strongly recommended for production installations
 - Appropriate for smaller clusters with moderate data sharing workloads where availability is the primary motivator for pureScale
- **No compromise in availability as both options provide exactly the same levels of high availability**
- **Choice of interconnect based on your performance and scalability requirements**



Virtualized Deployments of DB2 pureScale

▪ Virtualized environment options include

- RDMA-capable interconnect
 - AIX LPARs, with dedicated RDMA network adapters per partition
 - KVM with RHEL, with dedicated 10 GE RoCE network adapters per partition
- TCP/IP sockets interconnect
 - AIX LPARs
 - VMware (ESXi, vSphere) with RHEL or SLES
 - KVM with RHEL

▪ Virtualized environments provide a **lower cost of entry** and are perfect for

- Development
- QA and testing
- Production environments with moderate workloads
- Getting hands-on experience with pureScale



Virtualized Deployments: Supported Configurations

Operating System	Virtualization Technology	InfiniBand Supported?	10GE RoCE Supported?	TCP/IP Sockets Supported?
AIX, SLES, RHEL	No virtualization (bare metal)	Yes *	Yes *	Yes
AIX	PowerVM (LPARs)	Yes *	Yes *	Yes
SLES	VMware	No	No	Yes
RHEL	VMware	No	No	Yes
	KVM	No	Yes *	Yes

* Dedicated interconnect adapter(s) per host/partition

- **VMware supported with**
 - Any x64 system that is supported by both the VM and DB2 pureScale
 - Any Linux distribution that is supported by both the VM and DB2 pureScale


- **KVM supported with**
 - Any x64 system that is supported by both RHEL 6.2 and DB2 pureScale
 - RHEL 6.2 and higher

DB2 pureScale Supported Storage

- **Full explanation in “Shared Storage Considerations” section of the Information Center**
 - <http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.qb.server.doc/doc/c0059360.html>

- **pureScale supports all storage area network (SAN) and directly attached shared block storage, referenced as a logical unit number (LUN)**

- **pureScale exploits two specific features in storage when available for best results with recovery**
 - Fast I/O fencing
 - Tie-breaker support
 - Storage that supports both features is considered "category 1"



Current Category 1 Storage List




Table 1. Category 1 storage device and multipath I/O driver combinations

Storage Devices	Multipath I/O drivers required for AIX systems	Multipath I/O drivers required for Linux systems	Protocol
IBM Storwize® V7000 (6.4.0.1 or higher)	SDDPCM»», MPIO««	DM-MP	Fibre Channel
IBM SAN Volume Controller (6.4.0.1 or higher)	SDDPCM»», MPIO««	DM-MP	Fibre Channel
IBM System Storage® DS8000® series	SDDPCM driver provided by IBM (driver file <code>devices.fcp.disk.ibm.mpio.rte</code> and <code>devices.sddpcm<Your specific AIX version>.rte</code>)	DM-MP	Fibre Channel
IBM System Storage DS5000 series	MPIO driver provided by IBM (driver file <code>devices.fcp.disk.ibm.mpio.rte</code>)	DM-MP or RDAC	Fibre Channel
IBM System Storage DS4000® series	MPIO driver provided by IBM (driver file <code>devices.fcp.disk.ibm.mpio.rte</code>)	DM-MP or RDAC	Fibre Channel
IBM System Storage DS3000 series	MPIO driver provided by IBM (driver file <code>devices.fcp.disk.ibm.mpio.rte</code>)	DM-MP or RDAC	Fibre Channel
EMC VMAX/Symmetrix family ¹	»EMC PowerPath 5.5 P04 B003 ¹⁰ and EMC AIX ODM Package 5.3.0.6« MPIO driver provided by EMC (driver file <code>EMC.Symmetrix.fcp.MPIO.rte</code>)	DM-MP	Fibre Channel
NetApp FAS filers	MPIO driver provided by NetApp	DM-MP	iSCSI
Virtual I/O Server (VIOS) ²	MPIO or SDDPCM	»Not Applicable«	Fibre Channel
Hitachi Virtual Storage Platform (VSP) ⁵	MPIO driver provided by IBM or HDLM	»not supported as category 1 in Linux«	Fibre Channel
Hitachi Universal Storage (HUS) 100 Series ⁵	»MPIO driver provided by IBM or HDLM«	DM-MP	Fibre Channel
Hitachi Universal Storage (HUS) VM ⁷	MPIO driver provided by IBM or HDLM	»not supported as category 1 in Linux«	Fibre Channel
»Hitachi VSP G1000« ⁷	»MPIO driver provided by IBM or HDLM«	»not supported as category 1 in Linux«	»Fibre Channel«
IBM XIV® Storage System	MPIO	DM-MP ⁸	Fibre channel
»IBM FlashSystem™ 840« ⁹	»MPIO driver provided by IBM (driver file <code>devices.fcp.disk.ibm.mpio.rte</code>)«	»DMMP«	»Fibre Channel«
»IBM Storwize V5000 (7.3.0 or higher)« ⁹	»SDDPCM«	»DM-MP«	»Fibre Channel«

See "Shared Storage Considerations" section of Information Center for latest information:

<http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.qb.server.doc/doc/c0059360.html>

Managing and Monitoring pureScale Using DB2 Tooling

Task	pureScale Support	Product
Database Administration	<ul style="list-style-type: none"> Ability to perform common administration tasks across members and CF Integrated navigation through shared data instances 	Data Studio, Data Server Manager
	<ul style="list-style-type: none"> Recover database objects safely, precisely, and quickly 	DB2 Recovery Expert
	<ul style="list-style-type: none"> Support for high speed unload utility 	DB2 High Performance Unload
	<ul style="list-style-type: none"> Merge incremental backups into a full backup 	DB2 Merge Backup
System Monitoring	<ul style="list-style-type: none"> Integrated alerting and notification Seamless view of status and statistics across all members and CFs 	Data Studio Web Console, Optim Performance Manager, Data Server Manager
Configuration Tracking and Client Management	<ul style="list-style-type: none"> Full support for tracking and reporting of configuration changes across clients on servers 	Optim Configuration Manager
Application Development	<ul style="list-style-type: none"> Full support for developing Java, C, and .NET applications against a DB2 pureScale environment 	Data Studio
Query Tuning	<ul style="list-style-type: none"> Full support for query, statistics, and tuning advice for applications on pureScale systems 	Optim Query Workload Tuner

Typical pureScale Adoption Paths

- **Consideration is given as to what needs are met by pureScale**
 - Tier-1 databases that require the highest levels of availability?
 - Workload requires ability to scale out to multiple servers?
 - Implementation of a platform for consolidation / multi-tenancy?

- **Some clients have gone through the process of implementing an internal pilot or test system**
 - With a particular production database environment in mind
 - Or just to bring pureScale in to get some hands-on experience with it

- **A pilot allows you to go through full set of administrative tasks and see how it fits into your environment and current operational processes**
 - Also allows you to go through the full range of test scenarios
 - Failures and failovers
 - Online system and database maintenance
 - Performance and scalability testing (if test system is representative of a production environment)

Advice from Current pureScale Users

- Educate yourself on hardware and software prerequisites to be sure that you're ready to move forward without delay
- While much of the RSCT/TSA/GPFS interactions are abstracted away through DB2 commands, its useful for DBAs to have a basic understanding of concepts and know how to do basic PD/PSI
- If you're planning on using an RDMA interconnect, some of the networking aspects are likely to be new to you, so get to know your network administrator and talk with them about this
- If pureScale is likely in the future, but not just yet, create your databases so that they'll be compatible (for example, all automatic storage table spaces and perhaps on GPFS)

DB2 10.5 Trial Software Available

<http://www-01.ibm.com/software/data/db2/linux-unix-windows/downloads.html>

IBM Software > Information Management > Data Management > DB2 Product Family >

DB2 for Linux, UNIX and Windows

Accelerate insight. Improve resiliency. Maintain flexibility.

Overview Features and benefits **Download DB2**

Download DB2 10.5 to experience the world-class functionality that DB2 has to offer.

Use DB2 core database functions, warehousing capabilities and other advanced features for up to 90 days

DB2 Express-C is a no-cost community edition of DB2; available for download, deployment, and redistribution.

Download and use DB2 for 90 days

[Download DB2 Trial version ↓](#)

Download free DB2 Express-C database software

[Download DB2 Express-C ↓](#)



Offering

DB2 Data Server Trial
Version 10.5

Languages:
All Lang per ESD/PA Media

[Continue](#)

- 90 day trial
- Includes the pureScale feature
- With TCP/IP sockets support, no new hardware investment required
- VMware installs also possible

DB2 Business Application Continuity (BAC) Offering

■ Two member, active/active DB2 pureScale configuration

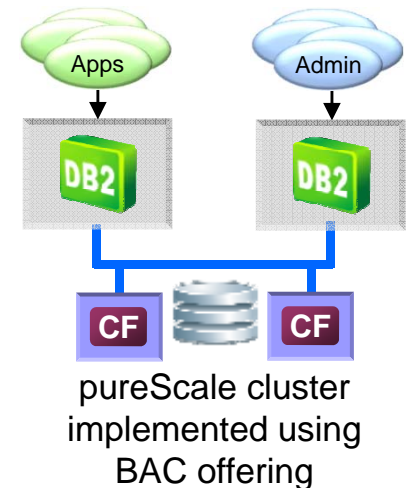
- All application workloads are directed to one primary active member
- Utilities and admin tasks allowed on the secondary admin member
- Application workloads quickly failover to secondary member during planned or unplanned outages

■ Low cost active/passive licensing model

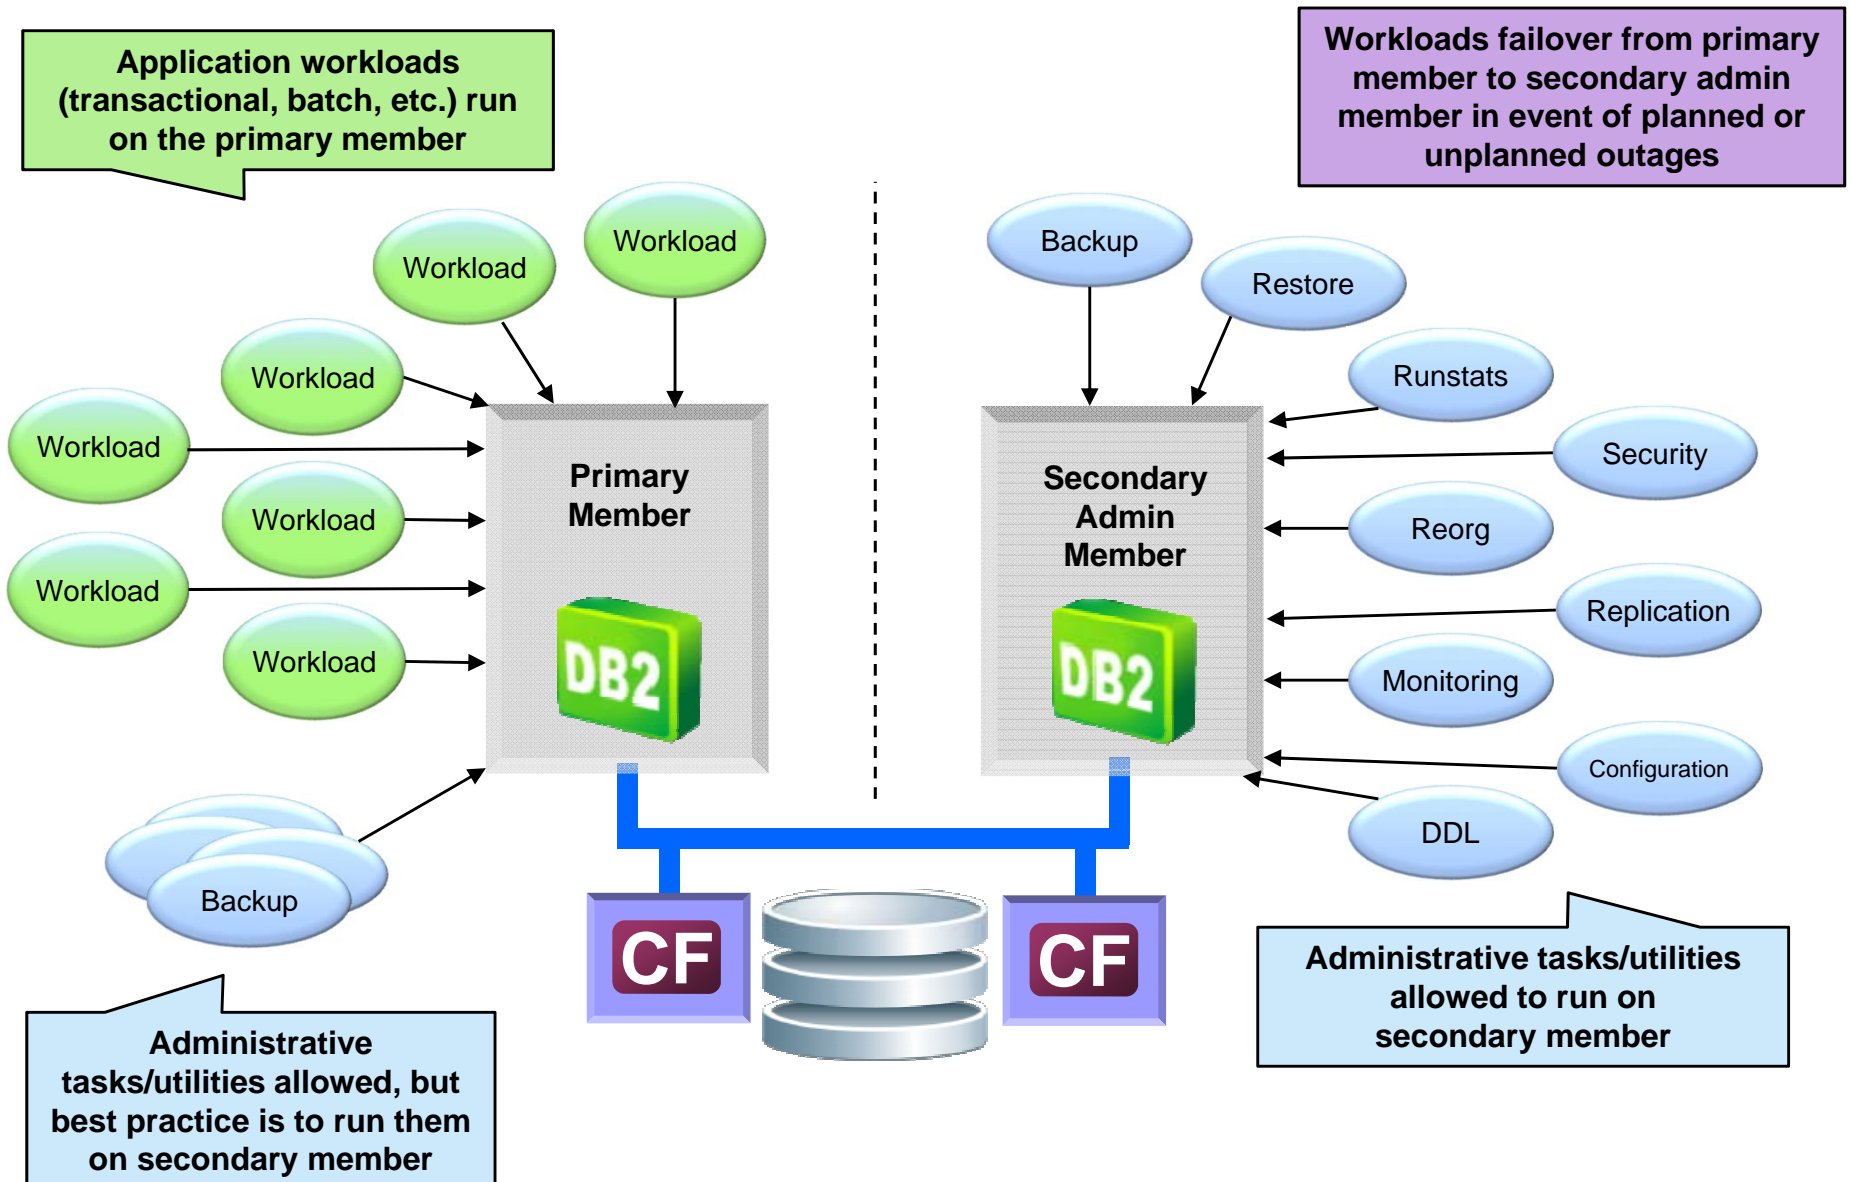
- Primary member fully licensed
- Secondary admin member licensed as idle/warm standby
- Available for DB2 Workgroup Server Edition (WSE) and DB2 Enterprise Server Edition (ESE)

■ Administrative activities allowed on secondary member includes

- Backup and restore
- Reorg and runstats
- Monitoring
- Usage of Data Definition Language (DDL)
- Database Manager configuration
- Database configuration
- Log based capture utilities for the purpose of data capture
- Security administration and setup

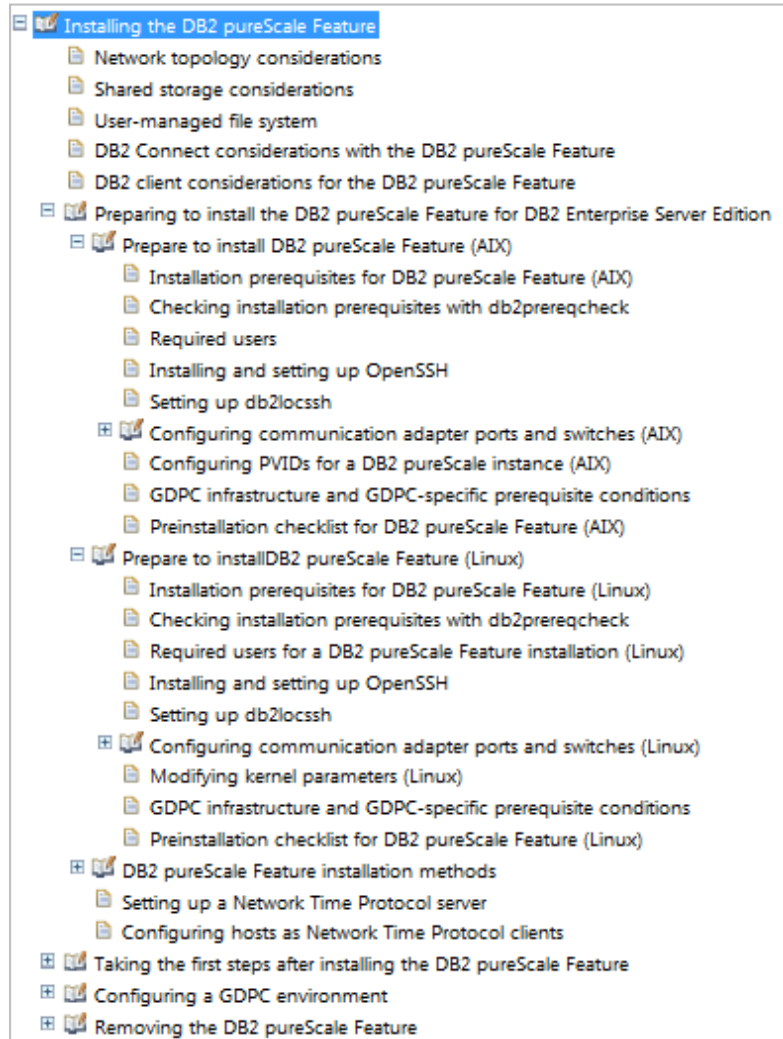


DB2 Business Application Continuity Offering Configuration



Setting Up Your DB2 pureScale Environment

Like the Boy Scouts always say... "Be Prepared"




- **The DB2 Information Center provides plenty of information including**
 - Pre-requisites for hardware and software
 - Preparation and planning steps
 - Pre-installation "cheat sheets"
 - Installation instructions

Installing DB2 pureScale

▪ Components installed

- Database members
- Cluster caching facilities (CFs)
- DB2 cluster services, which includes TSA, RSCT, and GPFS



▪ Methods

- DB2 Setup Wizard (`db2setup`): User friendly GUI for interactive deployment
 - Allows setup of a DB2 pureScale environment across multiple hosts
 - Ensures that all necessary components are installed and configured on all hosts according to best practices
- `db2_install` command: Command line install process (*deprecated*)
- Response file: Automated install

▪ Install and instance creation includes

- Installing binaries
 - Occurs across all hosts – does not require installation separately on each one
 - All components are installed and configured as part of the single install process
- Cluster domain created
- `sqlllib_shared` file system created and mounted on all hosts

Check Prerequisites using db2prereqcheck

- **Checks whether your system meets prerequisites for the installation of a specific version of DB2 before installing it**
 - Including the pureScale prerequisites
 - Includes OS level, Linux distribution, AIX technology level, C library, and uDAPL
- **Prerequisites contained within an XML file**
 - Contains prerequisites for DB2 9.8, 10.1, and 10.5
 - Command and file located in `<installPath>/cfg/DB2prereqs.xml`

▪ Examples

```
– db2prereqcheck -i -p
```

Check pureScale requirements for latest version of DB2 described in XML file

```
– db2prereqcheck -p -v 10.5.0.0
```

Check pureScale requirements for DB2 10.5.0.0

```
– db2prereqcheck -u -v 10.5.0.0
```

Check only pureScale uDAPL requirements for DB2 10.5.0.0

Pre-Installation Cheat Sheets

Preinstallation cheat sheet

Enter the appropriate required item value in the "Your Value" field.

Table 1. Preinstallation cheat sheet

Required Item	Your Value	Example
Instance owner/group name		<code>db2sdini/db2iadm1</code>
Fenced user/group name		<code>db2sdfe1/db2fadm1</code>
Installation directory name		<code>/opt/IBM/db2/V10.1</code>
Shared file system disk		<code>/dev/hdisk12</code>
Hosts to include		<code>db2_host01 - db2_host04.</code>
Netname interconnect for each member and CF		InfiniBand network example:

Table 2. Preinstallation cheat sheet - optional items

Optional Item	Your Value	Example
Tiebreaker disk		On AIX: <code>/dev/hdisk13</code> On Linux: <code>/dev/dm-0</code> Or <code>/dev/sdc</code>
FCM port range		<code>60000 - 60004</code>
cluster caching facilities port range		<code>56000 - 56001</code>
DB2 communication port		<code>50001</code>
Hosts to set up as cluster caching facilities		<code>db2_host03</code> and <code>db2host04</code>
On InfiniBand the		

AIX: <http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.qb.server.doc/doc/r0056077.html>
Linux: <http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.qb.server.doc/doc/r0057204.html>

Converting an Existing DB2 Instance to pureScale

- **As root, update the instance to a DB2 pureScale instance using one of `db2iupdt` (command line) or `db2issetup` (GUI)**

– For example:

```
db2iupdt -d -m coralpib154 -mnet coralpib154-ib0
        -cf coralpib153 -cfnet coralpib153-ib0
        -instance_shared_dev /dev/hdisk2
        -tbdev /dev/hdisk4 demoin1
```

```
> cat sqllib/db2nodes.cfg
0 coralpib154.torolab.ibm.com 0 coralpib154-ib0 - MEMBER
128 coralpib153.torolab.ibm.com 0 coralpib153-ib0 - CF
```

- **If using `db2iupdt`, instance will start with one CF and one member**

– Add additional members and a second CF using `db2iupdt -add`

Moving Existing Databases to pureScale

▪ DB2 pureScale database prerequisites

- All table spaces must be managed by Automatic Storage
- Database and logs must reside on GPFS
 - Can setup file systems in advance and move database to it
 - If pureScale not yet installed, use `db2cluster_prepare` command to create DB2 managed GPFS file systems



▪ Run `db2checkSD` to verify that a database is ready to be moved to pureScale

- Verifies that unsupported features are not used, all table spaces are automatic storage, etc.
- For example: `db2checkSD SAMPLE -l db2checkSD.log`

▪ Can backup DB2 10.5 non-pureScale database and restore into DB2 10.5 pureScale instance

Migrating Databases to Automatic Storage

- **Enable database for automatic storage by creating a storage group**
 - Existing table spaces are not impacted at this point

- **Drop temporary table spaces and recreate as AS**

- **Convert DMS table spaces to AS using one of**
 - ALTER TABLESPACE
 - Redirected restore of table spaces

- **SMS cannot be converted to AS**
 - Fundamentally different architectures
 - Must rebuild the table space (e.g. ADMIN_MOVE_TABLE)

- **If SYSCATSPACE is SMS then two options exist**
 - Rebuild the database (could use db2look and db2move)
 - Might be able to use Transportable Schema feature to move all existing table spaces into a new automatic database (with AS SYSCATSPACE)

The System is Setup... Now It's Time to Play

- **Types of tests to highlight "always on" availability and scalability**
 - Workload balancing across multiple members
 - Graceful shutdown cases to show automatic client reroute is working and applications aren't impacted
 - Non-graceful shutdown cases that show failure recovery occurs as expected and workload gets rebalanced
 - Scalability tests to show that applications changes are not necessary as members are added to the cluster
 - Silently perform system and DB2 fix pack maintenance without shutting the cluster down and impacting application availability

Sample Tests

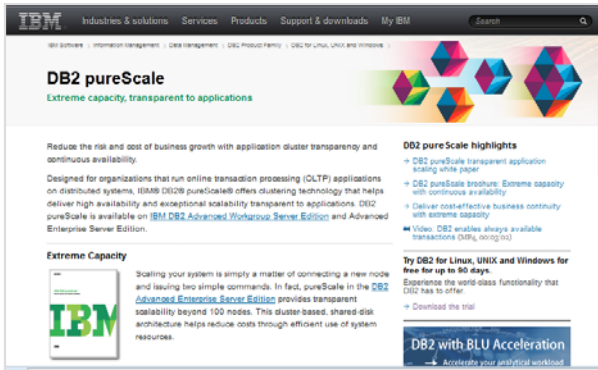
- **When running tests, start application and ensure that workload balancing (or client affinity) is behaving as expected**

Test	Expected Results
Graceful quiesce/shut down of member for system maintenance	No impact to the application. Connections should move to the other members as the transactions commit. Application work will be routed to member after it is started again.
Shut down of member for system maintenance, forcing applications off	Transactions running on member will be forced off and rolled back. Connections will automatically be reestablished on other members. Transactions may be seamlessly re-executed in some conditions or application may have to re-execute transaction statements. Application work will be routed to member after it is started again.
Killing of DB2 member process	DB2 automatically restarted on same host. Member recovery performed to rollback in-flight transactions and release update locks. In meantime, connections reestablished elsewhere. Transactions may be seamlessly re-executed in some conditions or application may have to re-execute transaction statements. Application work will be routed to member again after recovery is complete.
Non-graceful shutdown of DB2 member host	Restart light processing occurs on a different, healthy host to rollback transactions and release update locks. In meantime, connections reestablished elsewhere. Transactions may be seamlessly re-executed in some conditions or application may have to re-execute transaction statements. Application work will be routed to member again after host issue resolved and DB2 is automatically restarted on it.

Sample Tests (cont.)

Test	Expected Results
Graceful shutdown of primary CF for maintenance	Secondary CF takes over as primary. No errors returned to the application. When maintenance completes, CF process will be restarted on host, going through catchup processing and moving into peer (secondary) state.
Killing of primary CF process	Secondary CF takes over as primary. No errors returned to the application. CF process will be automatically restarted on host where primary CF was originally killed, going through catchup processing and moving into peer (secondary) state.
Non-graceful shutdown of primary CF host	Secondary CF takes over as primary. No errors returned to the application. When issue resolved with failed host, CF process will automatically be restarted on it, going through catchup processing and moving into peer (secondary) state.
Same tests as above for secondary CF (in peer state)	Primary CF remains as-is with no impact to applications. When stopped/failed CF is restarted it will go through catchup processing and move back into peer (secondary) state.
Test scalability of application by starting a stopped member or adding a new member online	If the application is not currently consuming the resources on the existing members then starting/adding new members will result in the distribution of the same load over more members (lowering per-member activity). However, if the application is driving enough load to fully consume the existing members then the starting/adding of additional members should increase transactions-per-second (tps) metric accordingly.
Perform rolling fix pack maintenance	No impact to the application. If system was running at full capacity then quiescing of each member as it is updated will result in less overall work being done, with a reduction in transactions-per-second.

Additional DB2 pureScale Reading Material



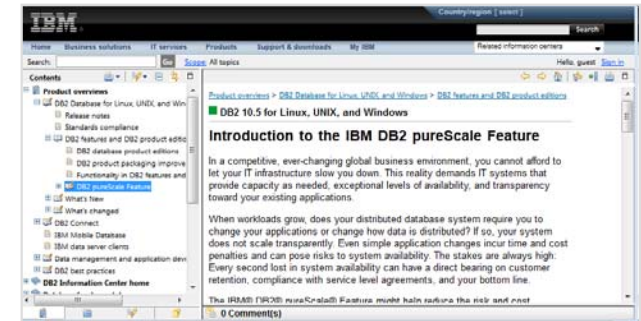
IBM DB2 pureScale Product Information

<http://www-01.ibm.com/software/data/db2/linux-unix-windows/purescale/>



IDUG Presentations

<http://www.idug.org/p/se/in/q=purescale&sb=1>




DB2 Information Center

http://www-01.ibm.com/support/knowledgecenter/SSEPGG_10.5.0/com.ibm.db2.luw.licensing.doc/doc/c0057442.html



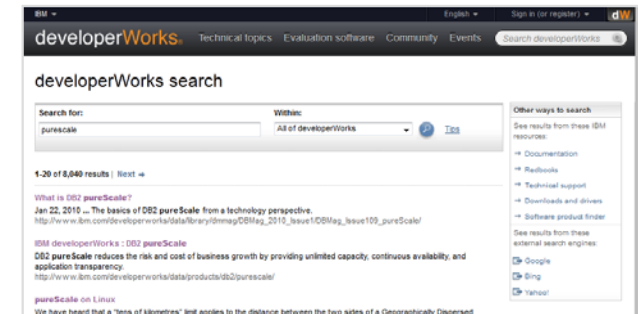
DB2 pureScale Redbook

<http://www.redbooks.ibm.com/abstracts/sg248018.html>



DB2 pureScale Book

<http://public.dhe.ibm.com/common/ssi/ecm/en/imm14079usen/IMM14079USEN.PDF>



developerWorks Articles

<http://www.ibm.com/search/csass/search/?q=purescale&sn=dw&dws=dw>

Next Steps

- **Contact your IBM representative**
 - pureScale deep dive, Proof of Technology (POT), Proof of Concept (POC)
 - Implement a pilot/test system in-house

- **Kick the tires by downloading and using DB2 10.5 (including pureScale) for 90 days**
 - <http://www-01.ibm.com/software/data/db2/linux-unix-windows/downloads.html>

- **If the availability of pureScale seems appealing but scalability isn't currently a concern, consider the Business Application Continuity Offering**

- **Take a course: "*DB2 pureScale Implementation and Control for DB2 LUW Administrators*"**
 - http://www-03.ibm.com/services/learning/ites.wss/us/en?pageType=course_description&courseCode=CL800

THANK
YOU

