


# 고성능, 고가용성 데이터베이스 구축을 위한 DB2 PureScale 및 Appliance 기술

강성희 차장, IMPE ASIA



1

DB2 PureScale 소개

2

Appliance 기반의 DB2 PureScale

3

Value of PureData System for Transaction

4

Q & A

# DB2 PureScale 소개

PureScale 은 OLTP workload 에 최적화되어있고, 24x7 운영이 필요한 고가용업무, 업무 증가에 따른 시스템 확장요구를 지원하도록 설계되어 있습니다.

### *Continuous 24x7 availability*

In the presence of planned and unplanned events

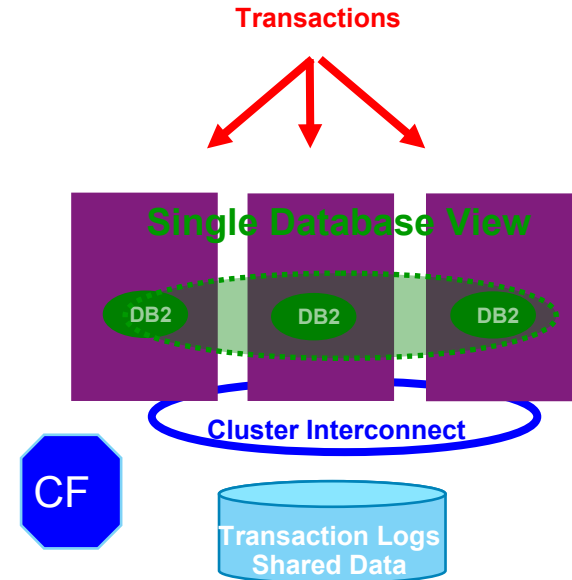
### *Simple growth*

Without significant application change  
Without administrative complexity

### *Rapid response to workload changes*

Through, for example, dynamic workload balancing and resource/machine additions and/or removal

## Low Cost of Management



## ***Closest solution to the Z “gold standard” in the distributed world***

Modeled after the Z Sysplex; but using COTS components  
Differentiated from competition via superior HA and scaling

노드 수가 증가하더라도 리소스 공유에 따른 오버헤드가 증가하지 않는 중앙 집중형 데이터 공유 아키텍처로 안정적인 대형 데이터베이스 구축이 가능합니다.

## Efficient Centralized Locking and Caching

As the cluster grows, DB2 maintains one place to go for locking information and shared pages

Optimized for very high speed access

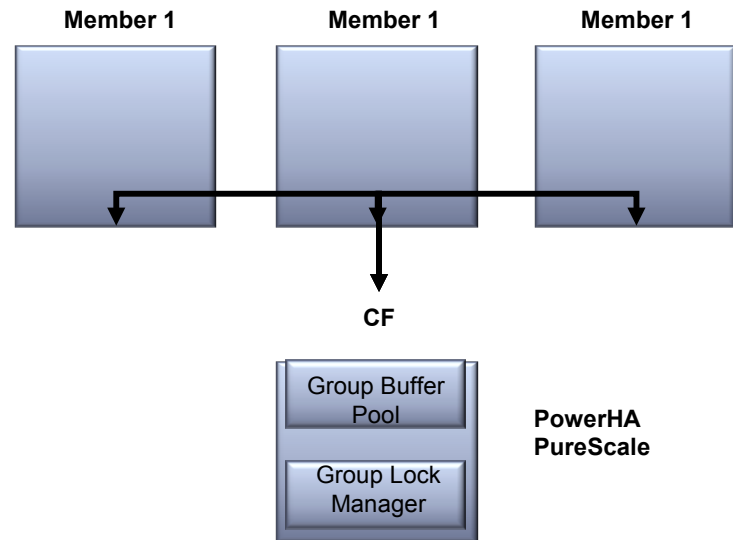
- DB2 PureScale uses Remote Direct Memory Access (RDMA) to communicate with the powerHA PureScale server
- No IP socket calls, no interrupts, no context switching

## Results

Near Linear Scalability to large numbers of servers

Constant awareness of what each member is doing

- If one member fails, no need to block I/O from other members
- Recovery runs at memory speeds



## PureScale 구성시 시스템 측면 주요 고려 사항입니다.

- Virtualization of the interconnect layer is currently no longer possible
  - DDR IB is older technology (these are the only adapters that could be shared by LPARs)
  - QDR IB is what you'll find with most systems today
  - 10 GigE is \*familiar\* to many
  
- Not all Power hardware supports all the adapters you need
  - P750 – DDR IB adapters only, and you can only have 1 if you have at least 2 processor boards
  - Check how many adapters can go into the box, and count your number of LPARs, they need to be 1:1 unless you have DDR
  
- If you happen to have DDR IB adapters, key points to remember
  - One adapter could be shared with up to 8 LPARs
  - You can not overcommit the adapter
  - Keep the sum of all member allocations equal to the CF allocation
  
- Keep your switches and adapters the same type
  - QDR and DDR do not play nice together

## PureScale 구성시 시스템 측면 주요 고려 사항입니다.

### ▪Processing Power

- Highly dependent on workload
- Number of members
  - The more members, the less cores per member, however, you may need more cores per CF
- Failure tolerances
  - How many concurrent member failures do you need to tolerate?

### ▪General Rule of thumb

- 1 core per CF for every 6 member cores when it's write heavy application
- 1 core per CF for every 10 member cores when it's a read heavy application

NOTE: In other words, add up the total member cores, divide by 6 or 10 and that's your starting point for the number of cores needed PER CF

## PureScale 구성시 메모리 사용 주요 고려 사항입니다.

### ▪Memory

- The more the better...
- Prevent swapping where possible

### ▪General rules of thumb for members

- 4 to 8 GB per core
- AUTOMATIC for the memory settings in DB2 is recommended with a few exceptions
- LOCKLIST should be set to 3 to 6% of the LBP size.

### ▪For the CF

- Dependent on read/write ratio - the higher the writes, the higher the memory requirements on the CF
- Minimum of 25% of all LPB should exist in the GBP
- 35% to 40% recommended for an average 70% read workload with 30% writes
- With 2 members only, 40% to 50% is recommended



## PureScale 구성시에서 Storage 주요 고려 사항입니다.

### ▪ Shared Storage Considerations for DB2 PureScale

- Latest list available at the DB2 information center
- Category 1 provides the fastest recovery times (roughly, in the 10 to 30 seconds range) when a host fails
- Category 2 provides reasonable recovery times (roughly, in the 1 to 3 minute range) when a host fails
- Category 3 is everything else that has not been officially tested with DB2 PureScale.

NOTE: failure times can depend on the number of file systems present and the type and number of concurrent failures occurring.

### ▪ Key points

- Make sure your driver matches what's in the table, and a lack of a driver listed, implies not supported.
- Keep an eye on the NOTES under the tables to get more information about a particular storage subsystem
- VIO must use N\_Port ID Virtualization to stay in Category 1, otherwise, you fall into Category 2

# PureScale 구성사에서 Storage 주요 고려 사항입니다.

Home Business solutions IT services Products Support & downloads My IBM Related information centers

Search:  Go Scope: All topics Hello, guest

**Contents**

- DB2 Information Center home
- Product overviews
- Database fundamentals
  - Installing
    - Requirements for DB2 products
    - Installing DB2 database servers
    - Installing the DB2 pureScale Feature
      - Network topology considerations
      - Shared storage considerations
      - User-managed file system
      - DB2 Connect considerations with the DB2 pureScale Feature
      - DB2 client considerations for the DB2 pureScale Feature
    - Preparing to install the DB2 pureScale Feature for DB2 Enterprise Edition
      - Installing DB2 pureScale Feature (AIX)
      - Installing DB2 pureScale Feature (Linux)
      - DB2 pureScale Feature installation methods
        - Setting up a Network Time Protocol server
        - Configuring hosts as Network Time Protocol clients
      - Taking the first steps after installing the DB2 pureScale Feature
      - Configuring a GDPC environment
      - Removing the DB2 pureScale Feature
    - Installing the DB2 Partitioned Database Environment
    - Response files
    - Installing the DB2 Information Center
    - Converting instances to a new DB2 database product
    - Converting instances to the IBM DB2 pureScale Feature
    - Applying fix packs
    - Uninstalling DB2 database products
    - Installing the integrated cluster manager
    - Installing IBM Data Studio
    - Installing IBM Data Server drivers and clients

## Category 1 storage device and multipath I/O driver combinations

Storage devices and multipath I/O driver combinations listed in this category can successfully support both the DB2 cluster services tiebreaker and fast I/O fencing. Category 1 devices have been validated with the DB2 pureScale Feature and result in the highest resiliency and fastest recovery times.

*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	DM-MP	Fibre Channel
IBM SAN Volume Controller (6.4.0.1 or higher)	SDDPCM	DM-MP	Fibre Channel
IBM System Storage® DS8000® series	SDDPCM	DM-MP	Fibre Channel
IBM System Storage DS5000 series	MPIO	DM-MP or RDAC	Fibre Channel
IBM System Storage DS4000® series	MPIO	DM-MP or RDAC	Fibre Channel
IBM System Storage DS3000 series	MPIO	DM-MP or RDAC	Fibre Channel
EMC VMAX/Symmetrix family <sup>1</sup>	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) <sup>2</sup>	MPIO or SDDPCM		Fibre Channel
Hitachi Virtual Storage	MPIO driver provided by IBM or HDLM		Fibre

# Appliance 기반의 DB2 PureScale

# PureData System for Transaction Hardware View



### Storage System

- Utilizes Storewize v7000 chassis and expansion units.
- Holds up to 192 disks in large configuration.
- 1:3 workload balanced SSD to HDD ratio

### Balanced HDD & SSD config.

- 48-disk module (12 SSD + 36 HDD) to optimize for performance & cost.
- RAID 10 for extreme storage reliability
- Up to 9.5 TB SSD and 64 TB HDD capacity

### Pure System™ Manager (PSM)

- 2 per rack for redundancy
- Integrate management for all system resources



### External Network Connectivity

- Dual 10Gb Ethernet Switches for external and rack-to-rack communication.



### Flex System Chassis

- Up to 2 per rack (Large).
- Holds up to 14 compute nodes each
- Back plane with fully redundant Network (En) and Storage (SAN) connectivity



### Flex System™ Compute Node

- up to 24 per rack
- Serves as CF and Member nodes for PureScale instances

## PureData System for Transactions – Available Configurations



Configurations	Upgrade		
	Small ¼ Rack	Medium ½ Rack	Large Full Rack
Flex System Chassis	1	1	2
Flex System Compute Nodes	6	12	24
CPU Cores	96	192	384
Memory	1.5 TB	3.1 TB	6.1 TB
Storage Unit (600GB drives)	1	2	4
Storage Expansion (600GB drives)	1	2	4
User Capacity	18.6 TB	37.2 TB	74.4 TB
Raw SSD Storage	4.8 TB	9.6 TB	19.2 TB
Raw HDD Storage	32.0 TB	64.0 TB	128.0 TB

## PureData System for Transactions – Cluster Elasticity



**Flexibility:** Predefined deployment pattern to handle different transaction requirements:

Number of Flex System Compute Nodes	CF	Member
6	2	4
4	2	2
2	*2	*2

\* CFs and Members are co-located

**Optimizing utilizations:** PureData system will support multiple PureScale instances in a single system:

Up to 3 PureScale instances for the small configuration

Up to 6 for medium and 12 instances for large

**Capacity on Demand:** Grow from 2 to 4 to 6 PureScale members as the workload demand increases

## TDA Checklist

How do I upgrade a PureData System for Transactions once capacity limit is reached?

**Answer:** PureData System for Transactions is designed for modular upgrade. When capacity limit is reached, you can seamlessly upgrade from Small to Medium to Large configuration. Simply contact your IBM sales representative to order and a PureData specialist will perform the upgrade on site.

How do I update the software on my PureData system?

**Answer:** For simplicity, complete software update packs will be released regularly to keep your system performing at optimal level. Once you download the update pack, you can choose to apply some or all component updates via a graphical management console.

Can I mount other servers to my PureData System rack?

**Answer:** PureData systems are designed to be highly integrated for optimal performance. Other non-PureData servers should not be added to the same rack since power and network requirements of these systems may negatively impact the performance of your PureData System.

Are there any hidden cost?

**Answer:** No. Your PureData system comes fully loaded with full feature DB2.

## The basic process to create your database in <4 hours.

**Plug in system and power on.**

**Complete the System Console configuration (“SGEN phase”)**

**Use the Workload console to:**

- Create the instance, specifying your cluster configuration (2, 4 or 6 nodes)
- Configure administrative access to the instance
- Create or deploy a database
  - Default workload
  - Customized workload
  - Clone

**Use the Database Operations Console to:**

- Configure access control
- Optional: schedule automatic database backups
- Add more disk storage to the database



**DONE!**

***After your instance and database have been created you can***

- *Begin loading objects into the database from a remote workstation using a supported tool such as IBM Data Studio*
- *Configure remote applications to access the database*



## Application deployment scenarios

### New deployment of application

- Database deployment through PureData console
- Application installation routine connects to PureData DB and creates database objects
  - Optional:* Use PureApplication application pattern which connects to database and creates structure in database
- Application connects to PureData just like a regular DB2 database

### Migration of existing application

- Database deployment through PureData console
- Migrate existing database using Database Conversion Workbench (DCW) to new PureData database
- Configure DB2 driver to connect to PureData

# Application deployment scenarios

## Database tools download

The screenshot shows the IBM PureData System for Transactions web interface. The top navigation bar includes 'IBM PureData System for Transactions', 'Workload Console', and 'System Console'. The user is logged in as 'admin'. The main menu has 'Welcome', 'Database', 'Cloud', and 'System'. The 'Database Tools' section is expanded, showing a list of tools with expand/collapse icons. The tools listed are Data Studio client, Data Tools Runtime Client, pureQuery Runtime, InfoSphere Data Architect, and Database Conversion Workbench. Each tool has a brief description.

**IBM PureData System for Transactions** | Workload Console | System Console | admin | Help | Logout

Welcome | Database | Cloud | System

**Database Tools** | Expand All | Collapse All

Use these database tools on your workstation, database client, or application server to develop database applications, administer, and monitor your databases throughout the data lifecycle.

- Data Studio client**
- Data Tools Runtime Client**
- pureQuery Runtime**
- InfoSphere Data Architect**

[InfoSphere Data Architect](#)  
InfoSphere Data Architect is a comprehensive development environment for data modeling, relating and integrating data assets, and developing database applications.

---

**Database Conversion Workbench**

[Database Conversion Workbench](#)  
A tool to migrate existing databases to the IBM PureData System for Transactions.

## Systems Management and Monitoring

**All systems management is handled through the System Management Console and Workload Management Console**

- System Management
- System Monitoring including Tivoli
- User Interface for System Monitoring
- User Interface for Workload Monitoring
- Troubleshooting
- Command line access for Power User to perform advance administration tasks  
(SSH login for root access)

## Backup and Restore Technology

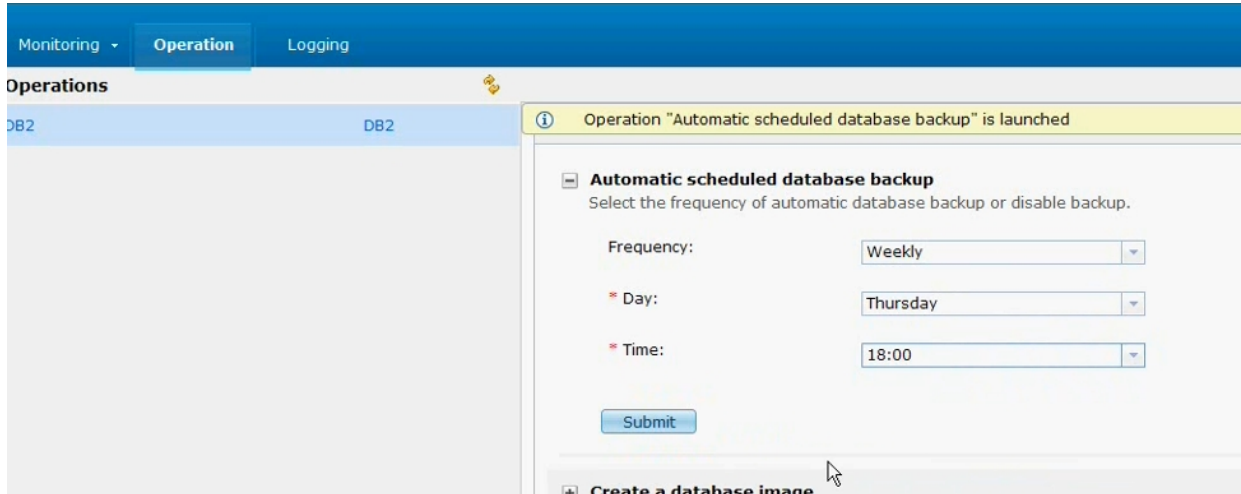
- DB2 Backup and Restore Technology is tightly integrated into the Workload Management Console in the PureData system for transaction.
- Identical in functionality to the Backup and Restore used in all other DB2 configurations.
- Backup and Restore Technology supported by the PureData System:
  - DB2 Backup and Restore to Tivoli Storage Manager (TSM) (recommended)
  - 3<sup>rd</sup> party Backup and Restore Products
  - External storage/ SAN for backup images

## Backup and Restore Technology

### From the Database Service Console you can easily schedule backups

- Granularity can be daily/weekly at a time of your choosing
- Backup images are available as sources for cloning

**Other centralized schedulers can be used to schedule backups like any other DB2 database, however calling the REST API will register the backup image with the console so it can be used for cloning.**



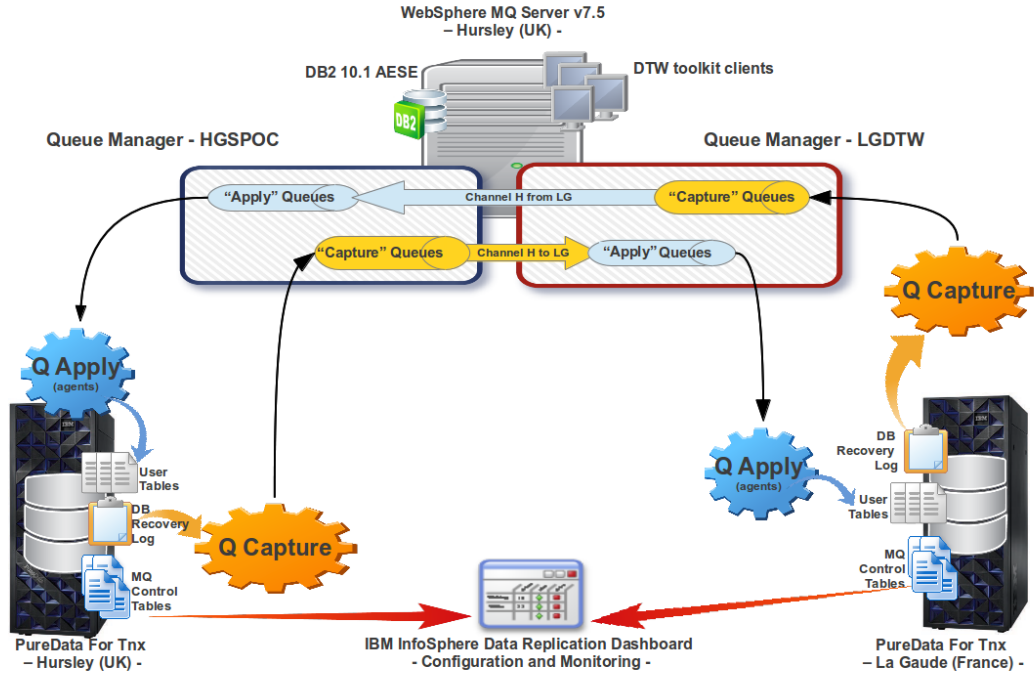
The screenshot displays the Database Service Console interface. At the top, there are navigation tabs: 'Monitoring', 'Operation' (selected), and 'Logging'. Below the tabs, the 'Operations' section is visible, showing a table with two columns labeled 'DB2'. A yellow notification banner at the top of the configuration panel states: 'Operation "Automatic scheduled database backup" is launched'. The configuration panel for 'Automatic scheduled database backup' includes the following fields:

- Frequency: Weekly
- \* Day: Thursday
- \* Time: 18:00

A 'Submit' button is located below these fields. At the bottom of the panel, there is a link labeled '+ Create a database image'.

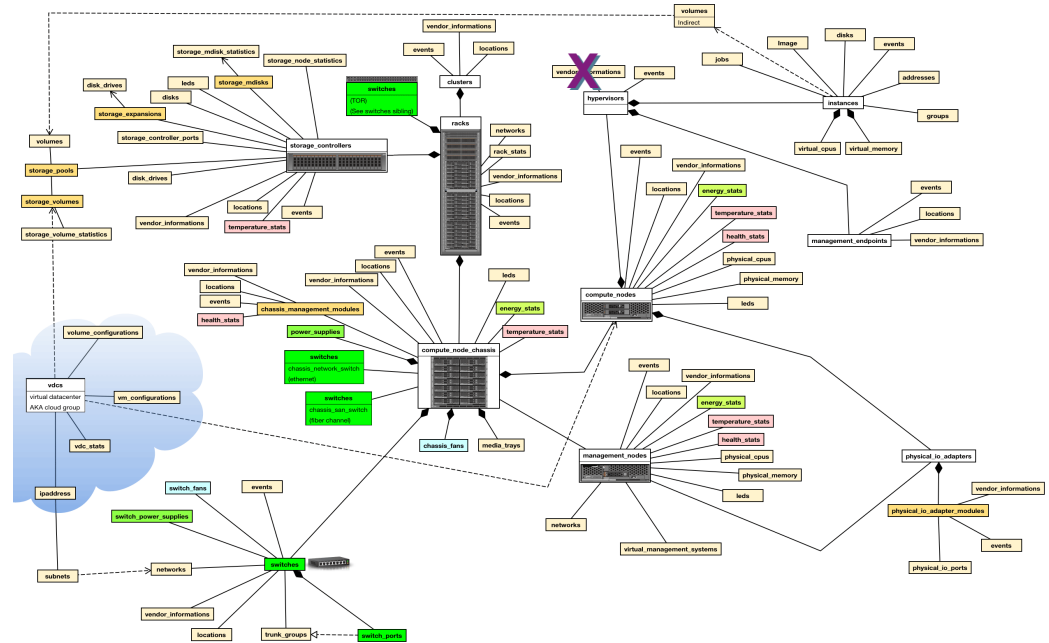
# DR

## Q Replication between PureData Systems for Transactions for Disaster Recovery



# Monitoring Schema

- The System Management Console, the Workload Management Console and the Database Operations Console can be used to monitor all aspects of the PureData System.
- The figure to the right shows all of the system, workload and database objects that can be monitored.



# Monitoring Schema

IBM PureData System for Transactions - Database Performance Monitor | 181-b740-e8d1099aa51d | Log Out | Help

Databases Health Performance Configuration Job Manager

Health Summary

Recent 60 minutes | Last refresh: Sun Nov 11 2012 09:39:34 AM | Add Database... | Configure Alerts

Alert Severity: All, Critical Alerts (5), Warning Alerts (3), Critical and Warning Alerts (5)

Data Source	Critical	Warning	Data Server Status	Monitoring Status	Operations	Member Status
db2inst2.dtw.d_83165e85_0a...	6	0			--	--
db2inst2.mydb.d_83165e85_0...	2	2			--	--
db2s55.s55.d_886ef6a8_693d...	6	0			--	--
db2s56.S56.d_d5130f9a_b45...	6	2			--	--
db2s57.S57.d_7011067f_2ba...	43	7			--	--

Health Summary | SQL Statements

View: Historical Data | End Time: 11/11/12 21:28 | Duration: - All - | Automatic Refresh: 26 sec | Baseline: 11/09/12 17:28 - All -

SQL Statements Dashboard: db2inst2.dtw.d\_83165e85\_0a7f\_4dc7\_a0c5\_682c1c78c106 [All Members] | db2inst2.dtw.d\_... | Disconnect

Learn about tuning SQL statements, stopping SQL statements, and forcing applications.

Top Individual Executions: Execution Summary

Statement Text	Statement Category	Execution Elapsed	Number of Executions	Number of Accesses	CPU Time	Rows Read	I/O Time	Loc
CALL SYSPROC.A...	CALL	02:25:24.664000	2	1	0.002000	0	0.000000	
select * from db2in...	DML_Select (block...	4.482000	1	1	0.000000	0	0.000000	
select count(*) from...	DML_Select (block...	5.845000	1	1	0.000000	0	0.000000	
CALL SYSPROC.A...	CALL	01:03:26.548000	1	1	0.019000	0	0.000000	

SQL Statement Details

Overview | Server Execution Times | Row Activity | I/O | Locking and Communication

Statement	Most Recent Identification	Most Recent Compilation
CALL SYSPROC.ADMIN_CMD(?)	Statement identifier: 01000000000000000a00100...	Compilation time: 0.000000
	Number of Accessed Members: 1	Number of compilations: 2
	Package name: --	Isolation level: CS
	Consistency token: --	Cost estimation (timersons): 1
	Section number: --	Special Registers for Compilation
	Cache insert time: 11/08 20:49:26	CURRENT SCHEMA: --
	Last Execution time: 11/08 23:20:31	CURRENT REFRESH AGE: 0
		CURRENT PATH: --



# PureFlex System 기반의 PureScale, PureData System for Transaction 은 H/W와 S/W 의 최적화가 설계 단계부터 고려되었습니다.


*Uninterrupted access to data with consistent performance*

**Traditional systems**  
- build it yourself

**Over several days/weeks:**

1. Define High Availability topology
2. Configure HW/SW/Network
3. Set up storage pools
4. Install multiple operating systems
5. Install database instances
6. Set up primary and secondary management systems
7. Set up database members
8. Set up backup processes
9. Test, tune, reconfigure

6-node database cluster



**PureData System**  
- built-in expertise

**In minutes,**

1. Just specify database, description and topology pattern

**PureFlex System 기반의 PureScale, PureData System for Transaction 의 내재된 전문성의 의 도움을 받고 또한 사용자 경험을 패턴으로 자산화 할 수 있습니다.**

*Deploy topology and databases in minutes using patterns*


PureScale™ Instances

**Topology patterns**

Automatically creates, configures and deploys a database system topology with built-in redundancy and high performance

**Database patterns**

Automatically creates, configures and deploys IBM or client-specified databases optimized for transactional workloads



## FAQ

### How do I upgrade a PureData System for Transactions once capacity limit is reached?

**Answer:** PureData System for Transactions is designed for modular upgrade. When capacity limit is reached, you can seamlessly upgrade from Small to Medium to Large configuration. Simply contact your IBM sales representative to order and a PureData specialist will perform the upgrade on site.

### How do I update the software on my PureData system?

**Answer:** For simplicity, complete software update packs will be released regularly to keep your system performing at optimal level. Once you download the update pack, you can choose to apply some or all component updates via a graphical management console.

### Can I mount other servers to my PureData System rack?

**Answer:** PureData systems are designed to be highly integrated for optimal performance. Other non-PureData servers should not be added to the same rack since power and network requirements of these systems may negatively impact the performance of your PureData System.

### Are there any hidden cost?

**Answer:** No. Your PureData system comes fully loaded with full feature DB2.

# Value of PureData System for Transaction

## IBM Pure Systems Family

### PureFlex



인프라

인프라 서비스  
Delivering Infrastructure  
Services

### PureApplication



애플리케이션 플랫폼

플랫폼 서비스  
Delivering Platform Services

### PureData



데이터 플랫폼

데이터 서비스  
Delivering Data Services

## IBM PureData Systems > PureData system for Transaction

# Meeting Big Data Challenges – Fast and Easy!



## PureData

*System for Transactions*

고가용성, 확장성, 통합성을 보유,  
대용량 트랜잭션을 처리하는 데이터베이스 서비스

## PureData

*System for Analytics*  
Netezza 기술 기반

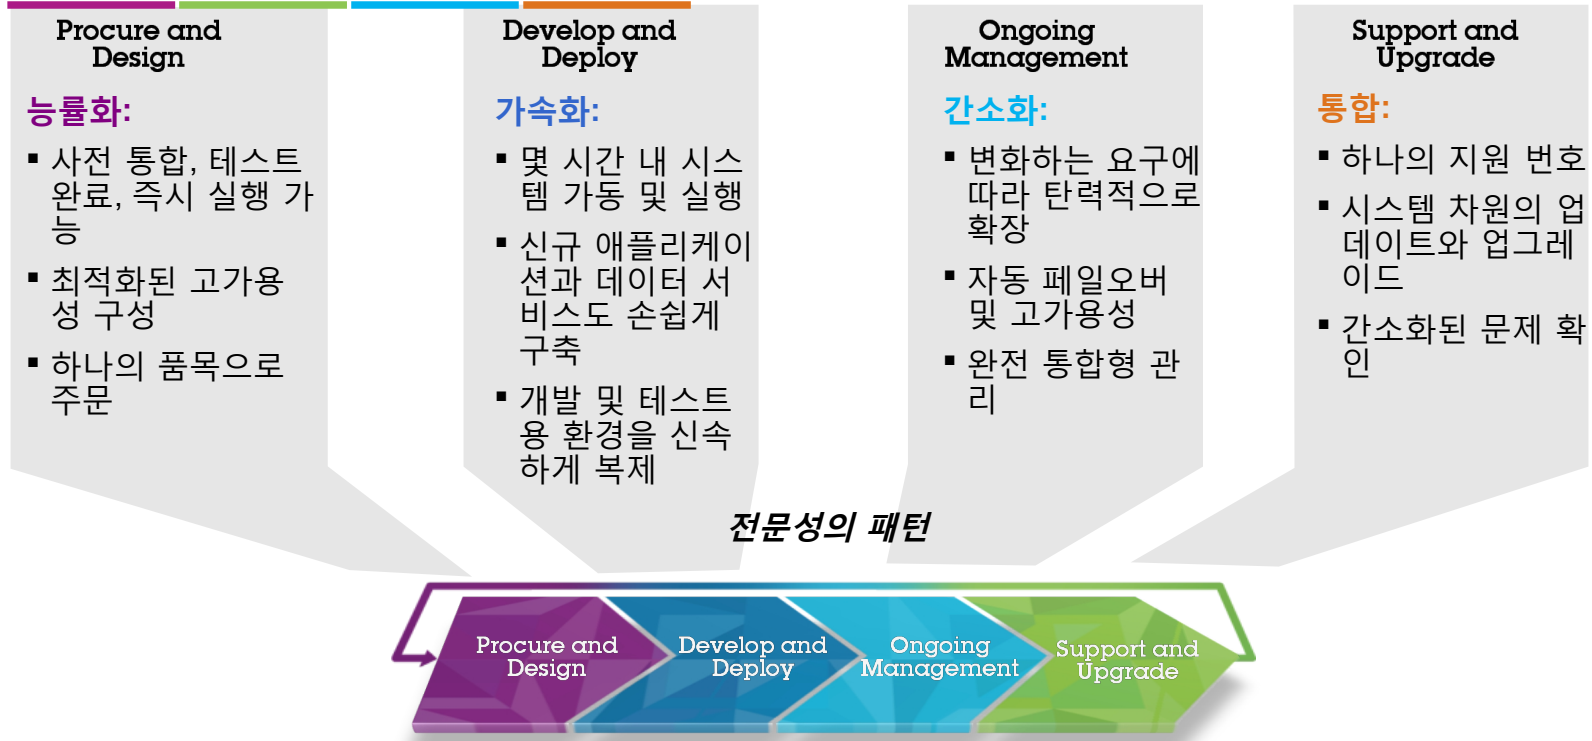
최소한의 관리로 petabytes 데이터를  
손쉽게 로드, 복잡한 분석 및 보고서 실행

## PureData

*System for  
Operational Analytics*

다양한 형태의 분석업무와 트랜잭션 처리가 혼재된 업무를 위한  
다수의 사용자 환경 운영을 위한 데이터웨어하우스 서비스

## IBM PureData Systems > PureData System for Transaction



**Q & A**



