

“New Generation of PDA”

**고성능 분석 및 머신러닝을 위한 맞춤형 솔루션
IBM Integrated Analytics System**

한국 IBM
Hybrid Data Management
IBM Cloud
2018.03

Do data science faster !!





고객들이 말하는 PDA 선택이유

IBM PureData System for Analytics를 선택한 고객의 50%가량은 기존 Oracle기반의 DW시스템의 성능 한계 및 고비용을 이유로 들고있습니다.

Exadata/TeraData/Sybase를 운영하기 위해 훨씬 많은 시간을 아키텍처 설계, 데이터 모델과 인덱스 생성, 시스템 구성, 테스트 및 기타 다른 작업에 보내야 했습니다.
또한 광범위한 성능튜닝은 업무추가에서는 늘상 수행해야 하는 일이었습니다.

타 DW시스템과 비교하여 PDA의 성능 튜닝 작업은 사실상 “존재하지 않는다”

PureData System for Analytics을 운영하고 있는 한 회사는 "우리는 어떠한 사전 교육 없이 6개월만에 시스템을 오픈하여 운영하고 있습니다" 라고 보고했습니다.

PureData System for Analytics를 운영하면서 우리는 인덱스를 만들 필요가 없었습니다. 대다수의 사용자는 “원하는 데이터를 얻기위해 DBA를 통해 요약/집계 테이블을 요청할 필요가 없게 되었습니다.”

Source: ITG: Comparing Costs and Time to Value with Teradata Data Warehouse Appliance, April 2013, ITG: Comparing Costs and Time to Value with Oracle Exadata Database Machine X3, August 2013.

새로운 DW 시스템 : IBM Integrated Analytics System

More speed, More flexible, and More reliable system



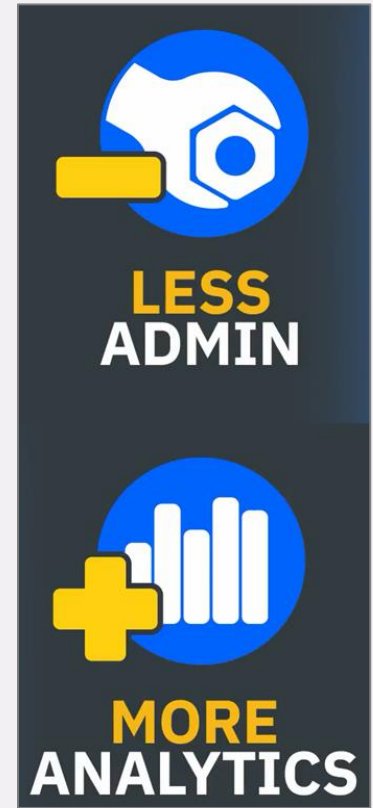
세계 최초의 DW Appliance
Netezza



높은 성능과 극도의 단순성, 낮은 TCO
IBM PureData System Analytics



Beyond PDA
“Do Data Science Faster”
IBM Integrated Analytics System





Alyse Daghelian

Vice President, Global Analytics Sales



간단하게 말해 IIAS 같은
제품은 없다

THERE IS SIMPLY
NOTHING LIKE IIAS

- ◇ Capabilities Around Netezza
- ◇ Speed >> 2x to 5x faster
- ◇ Ease-of-Use
- ◇ System for NOW + FUTURE
- ◇ Built-In Machine Learning
- ◇ Public / Private Portability
- ◇ Write Once - Run Anywhere



Rob Thomas

General Manager of Analytics Platform

TWO **MAJOR** **PROBLEMS**

IT담당자가 직면하는
2가지 중요한 문제



No. 1 Challenge

Clients are trying to apply machine learning and data science at scale.

**And they're realizing
It's really hard.**

대량의 데이터를 이용하여 머신러닝
기법과 데이터 사이언스를
적용해보려고 하나 현실적으로
매우 매우 힘들다

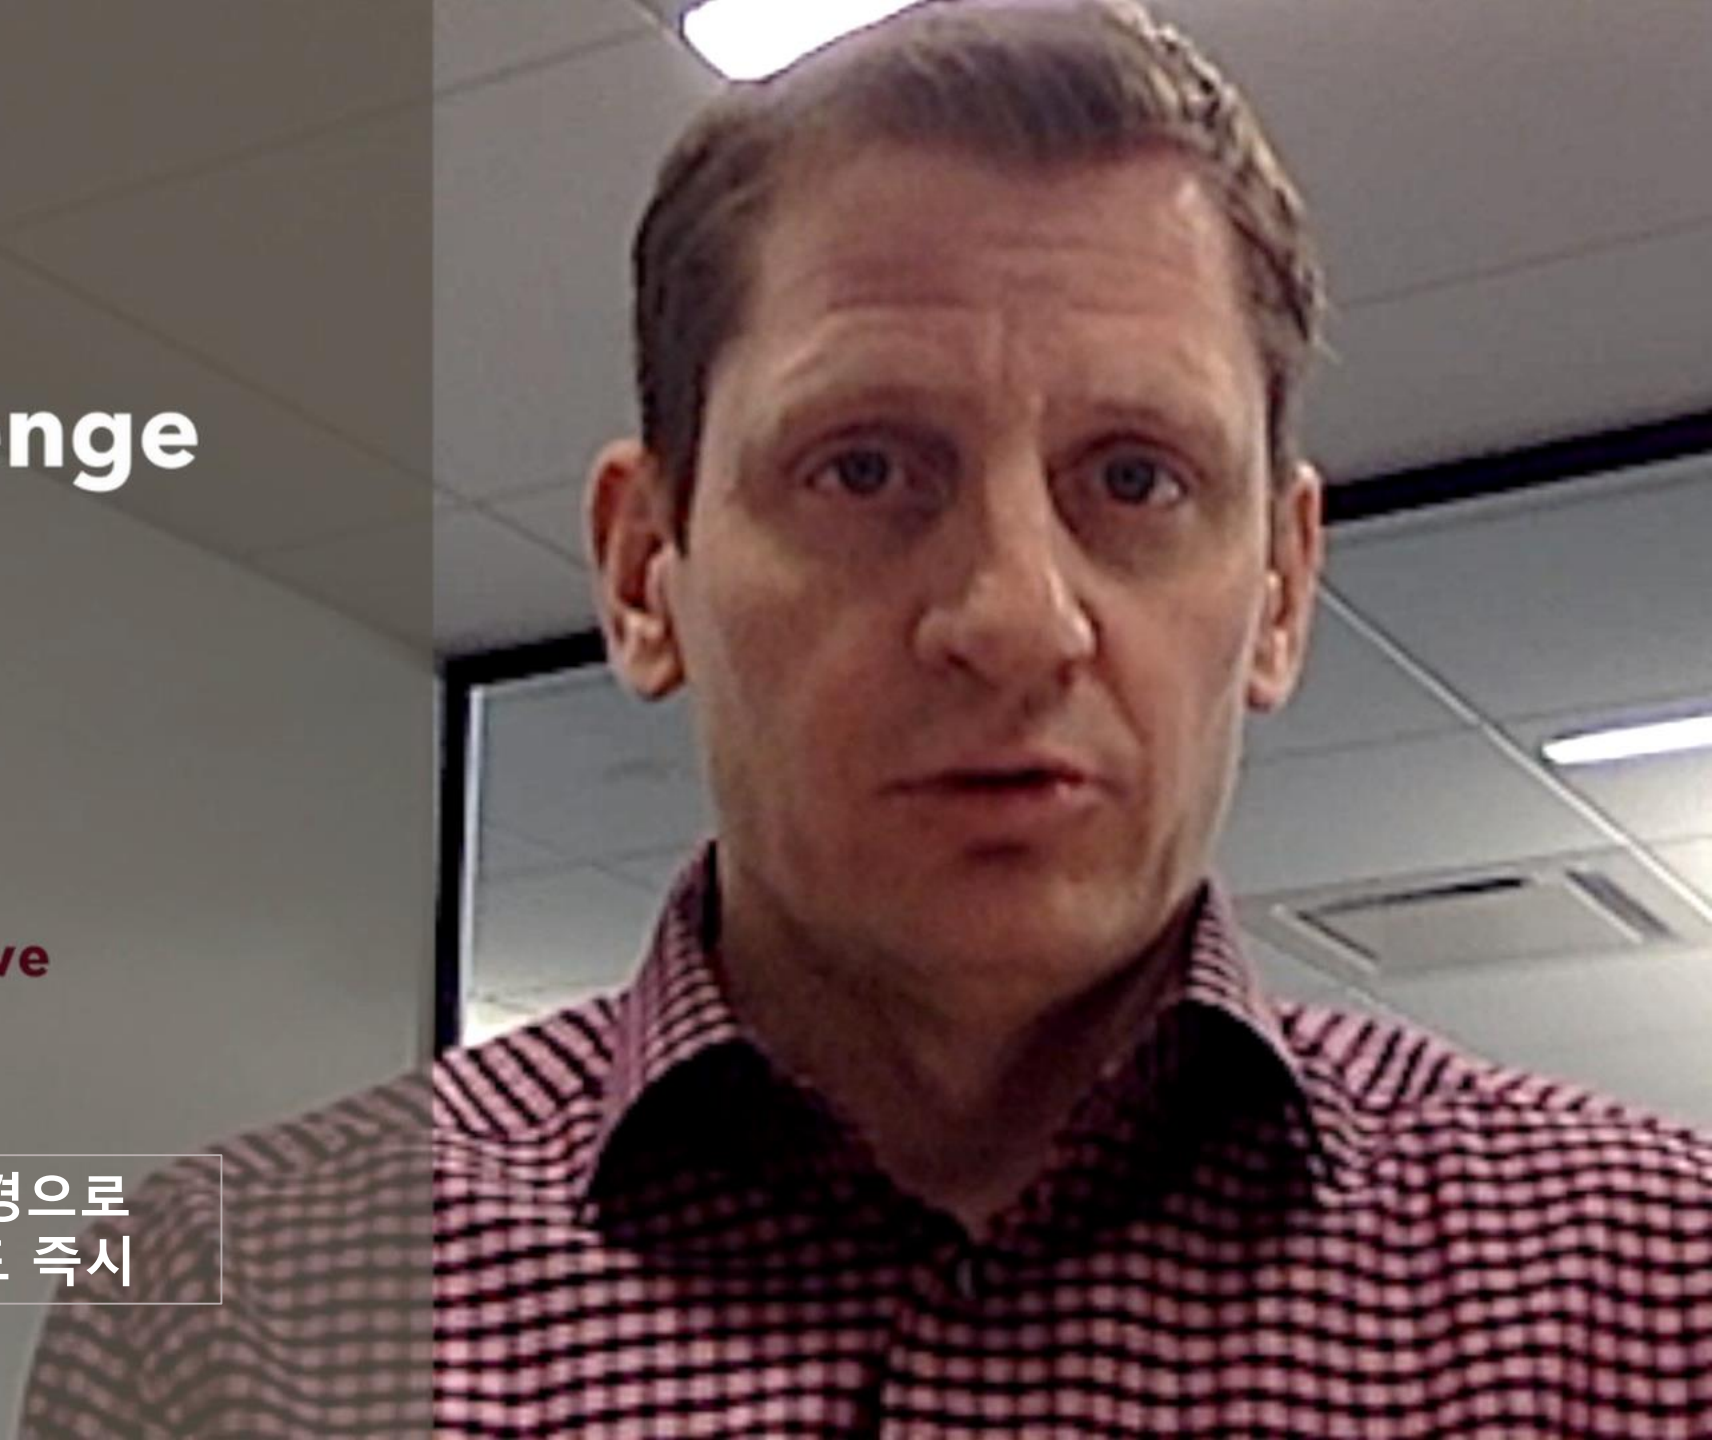


No. 2 Challenge

**Clients are trying
to move
to
Public Cloud.**

**And they want to move
there as soon
as they can.**

기존의 환경을 Cloud환경으로
옮기고 싶어 한다. 그것도 즉시



**Do
data science
faster**

“두가지 큰 변화”

1. Data Science를 위한 똑똑한 기능들이 추가

데이터 사이언티스트도 사용하는 시스템

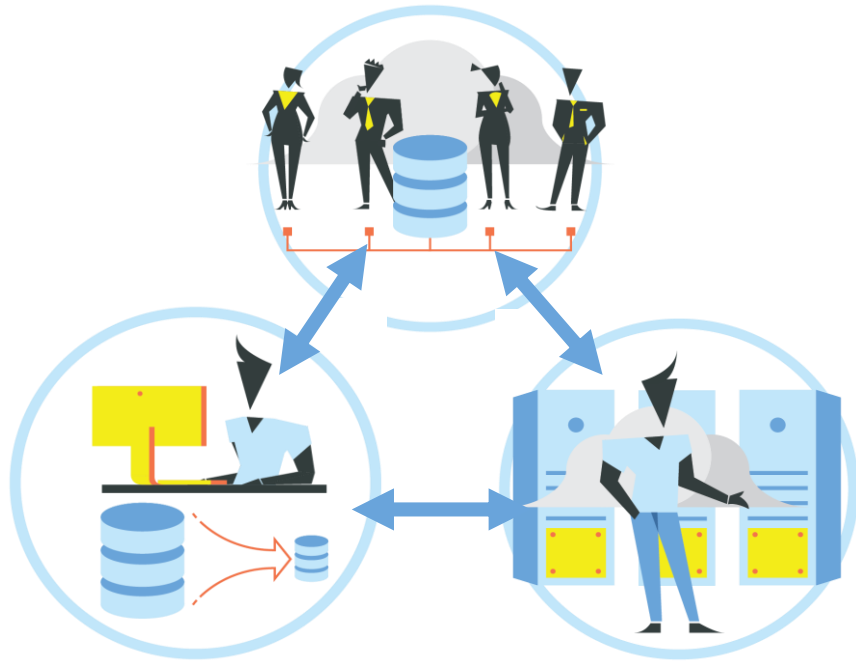


* **DSX (Data Science Experience) Local** : 데이터 과학자가 필요로 하는 데이터, 도구 및 패키지를 관리하고 이에 대한 액세스를 제공하는 on-prem 플랫폼 (Jupyter, Zeppelin, H2O Flow and RStudio / Anaconda for Python 2 and 3* support / Support for Python, Scala, and R languages)

2. Cloud Ready

Write Once – Run Anywhere

**퍼블릭 클라우드
데이터 웨어하우스**
(new data sources, deep analytics)



**온프레미스
데이터 웨어하우스**
(structured, defined set of data,
traditional BI)

**프라이빗 클라우드
데이터 웨어하우스**
(new data sources, deep
analytics)

Common SQL Engine

On-Premises, Private 클라우드 및 Public 클라우드에 있는 모든 IBM데이터 관리 제품군에 손쉽게 배포

운영 환경 통합

모든 데이터 플랫폼을 위한 단일 운영 인터페이스

클라우드로의 빠른 데이터 마이그레이션

클라우드로의 빠르고 안전한 데이터 이동 지원
(IBM Bluemix Lift)

모든 형태의 데이터 지원

데이터 가상화를 사용하여 Hadoop을 비롯하여 더 많은 유형의 데이터를 쿼리에 포함



**LESS
ADMIN**



**MORE
ANALYTICS**

IBM Integrated Analytics System

IBM Integrated Analytics System

Netezza Heritage + New Hardware + New Technology + Cloud Ready

기존 모델의 장점은 더 강화

Simplicity (간편성)

- **인덱스 관리 불필요** : Storage Index / Zonemap / 자동압축 / No Admin
- **물리적 모델링 없이 바로 사용** : Load & Go / 성능개선용 물리모델설계 불필요

Performance (성능)

- **고성능 하드웨어** : 더 빠른 CPU, 더 많은 Cache, 더 큰 Memory
- **Micro latency Flash Storage** : I/O 지연 없는 데이터 전송
- **In-Memory 기술** : 컬럼기반의 인메모리 기술로 N3001대비 수배 성능향상
- **신뢰성 높은 스토리지** : 2-Dimensional flash RAID (99.999% 신뢰 있는 하드웨어 컴포넌트)

In-Place Expansion (확장)

- **모듈 방식 확장** : Server + Storage 증설
- **스토리지 확장 지원** : Storage 증설



M4002

최신 기술 적용

In-Memory (인메모리 기술)

- **집계 및 분석에 용이한 컬럼기반 데이터베이스 기술** : IBM의 BLU기술을 이용하여 성능 및 기술 안정성 확보
- **속도 향상** : Analytic Query PDA N3001대비 수배 성능향상
- **리소스 최적화** : I/O를 최소화 하고 CPU사용률을 절감

Analytic Tool (분석을 위한 툴)

- **데이터 분석가를 위한 환경** : 분석에 필요한 머신러닝 Lib, Jupyter Notebook, Rstudio등을 기본 탑재하여 Analytic Sandbox 분석환경 구축 지원
- **외부 Spark 과의 연계**

Cloud Ready

- **SW Defined Analytic DW** : 동일한 구성으로 Public Cloud, Private Cloud, On-prem DW 구성가능
- **공통 SQL 엔진** : 공통 엔진 / 공통 SQL 사용으로 상호 이관이 빠르고 용이함

Analytic Workload 에 최적화된 IBM 하드웨어

대용량 데이터 분석과 처리를 위한 위한 H/W

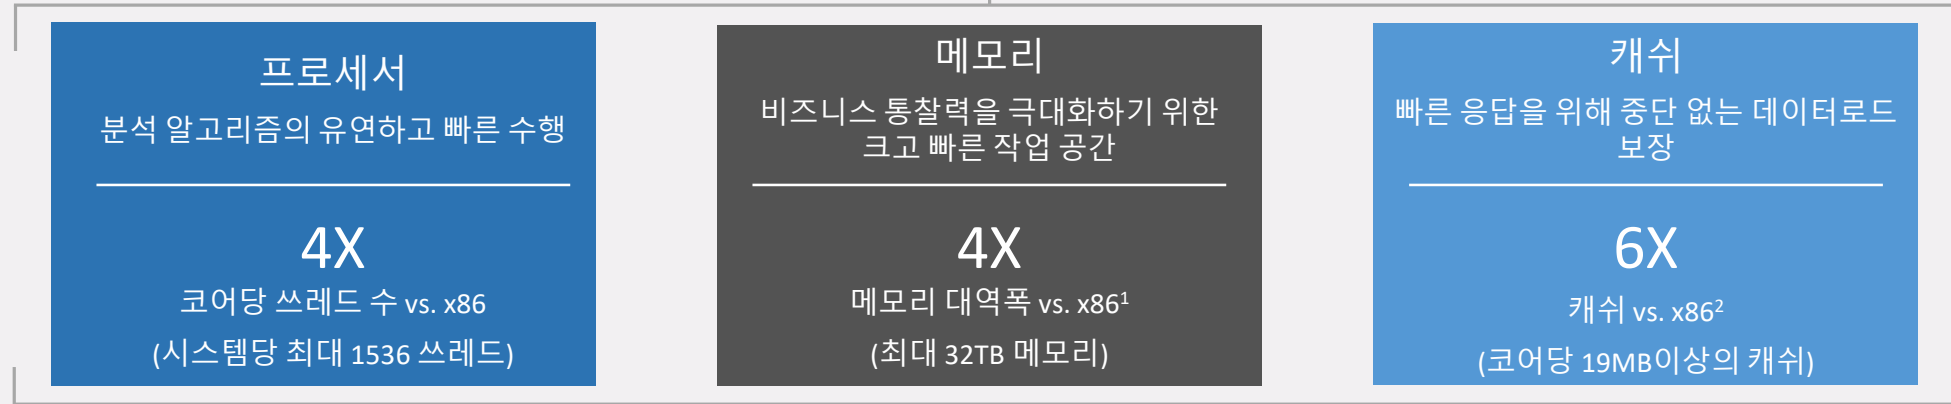
Power Architecture

- 더 적은 노드로 고성능
- 멀티 코어와 SIMD 병렬을 통한 CPU 가속
- 향상된 안정성 및 가용성



Flash Storage, standard

- 고속 데이터 전송 시 준 실시간의 지연을 통해 99.9%의 안정성과 운영 효율성



중단 없는 데이터 로드



방대한 IO 대역폭



병렬 처리



극한의 성능을 위한 Flash 사용



대규모 메모리 프로세싱



1. 특정 x86 및 POWER8 서버에 따라 최대 4 배까지 차이를 보임
2. Intel e7-8890 서버와 12 코어 POWER8 서버를 비교하여 최대 6 배 더 많은 캐시를 제공

Analytic Workload 에 최적화된 System Architecture

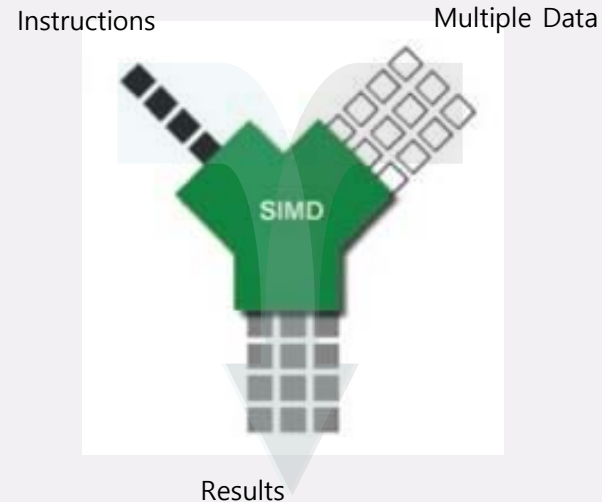
H/W 아키텍처에 최적화된 Columnar 처리

In-Memory Optimized

- 메모리에서 더 유용한 데이터 유지
 - 데이터는 압축 상태 유지
 - LRU Caching보다 진보된 Scan Friendly Caching
- 메모리로 유입되는 데이터량 최소화
 - Columnar access
 - Late materialization
 - Data skipping
- 최적화된 메모리 응답시간 이용
 - 스캔(Scans)
 - 조인(Joins)
 - 집계(Aggregation)

CPU Optimized

- CPU 가속화
 - SIMD 처리 for
 - 스캔
 - 조인
 - 그룹핑
 - 연산
- CPU 사용을 최대한 활용
 - Core 친화적인 병렬처리
- 적은 CPU 처리량
 - 압축된 데이터의 처리
 - Late materialization
 - Data skipping



I/O Optimized

- 읽기 횟수를 줄임
 - Columnar I/O
 - Data skipping
 - Late materialization
 - Scan Friendly Caching
- 효율적인 I/O
 - 특화된 columnar prefetching 알고리즘

Analytic Workload 에 맞는 In-memory 최적화 기술

Dynamic In-Memory 기술을 통해 빠르게 데이터 조회

예시

손쉬운 사용 (Ease of Use)

- 인덱스 불필요, 집계 테이블 최소화, No 튜닝
- Load and Go!

동적 인메모리(In-memory) columnar processing

- 동적 데이터 이동기술을 통해 성능 극대화 (Storage ↔ Memory)

Data skipping (zone maps)

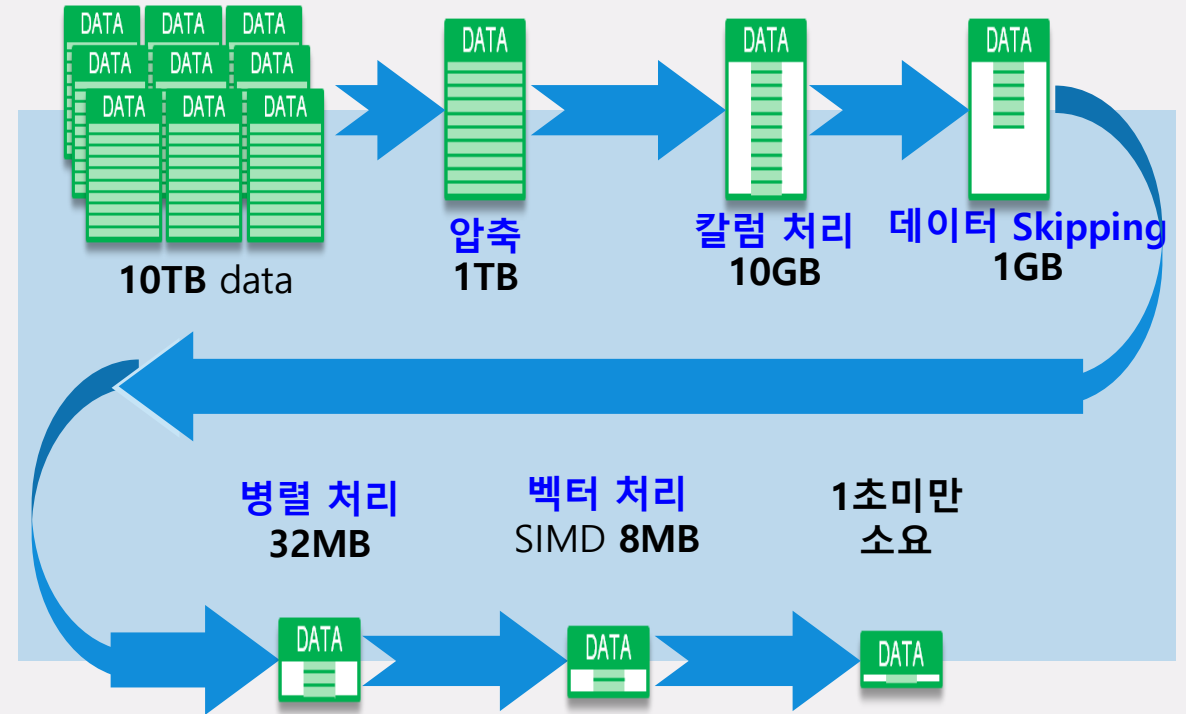
- 물리적인 인덱스 없이 필요한 데이터만 읽을 수 있도록 자동 처리하는 기술

전세계 일등 압축 기술 Actionable compression

- 압축 해제없이 데이터를 조회할 수 있는 특허 받은 압축 처리 기술

10TB 데이터 쿼리 수행 예시

- 데이터 크기 : 100개 컬럼을 가진 10TB 테이블 (10년 동안 축적된 데이터)
- 조회 : 2017년에 얼마나 많은 "SALES"가 일어났는가?
>> SELECT COUNT(*) FROM MYTABLE WHERE YEAR = '2017'



데이터 적재만으로 Query 수행 1초 미만 소요

데이터 가상화 실현

다른 이기종 데이터 소스에 대한 Transparent Access

다양한 DBMS 가 지원되는 **Transparent Access 기능**을 통해 이기종 소스에 대한 가상 뷰를 사용할 수 있음

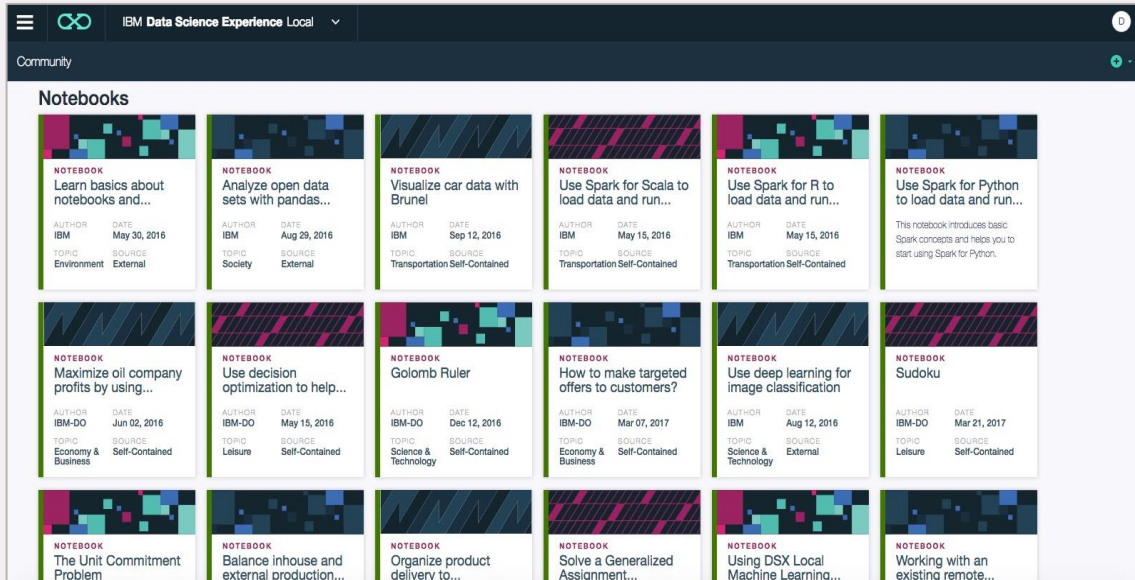


Analytic Model 생성을 위한 도구 제공

DW에서 고급분석과 머신러닝 수행

IBM Data Science Experience (DSX)

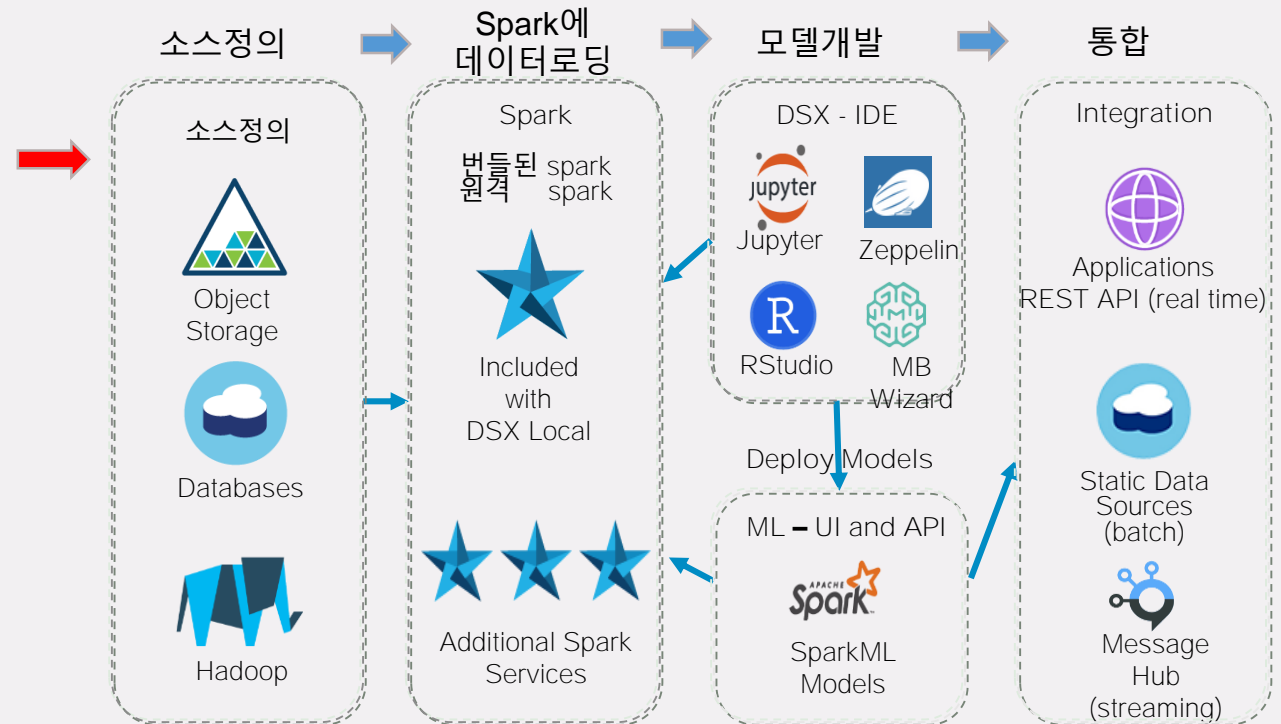
- 분석 애플리케이션을 공동으로 개발 및 배포 하기 위한 플랫폼
- 분석 응용 프로그램을 빌드하고 배포하는 데 필요한 모든 것을 포함
- 모델생성은 Jupyter Notebook에서 프로그래밍 방식으로 작성
- 확장 가능 : remote Spark, 다중 데이터 소스, 라이브러리 추가,LDAP



IIAS와 DSX 의 통합 시너지

- 일반 DB사용자 뿐만 아니라 Data Scientist 까지 사용자 계층 다양화하고 Analytic Sandbox 로 활용

DSX를 이용한 Analytic Flow

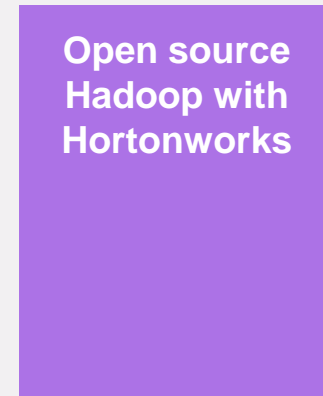
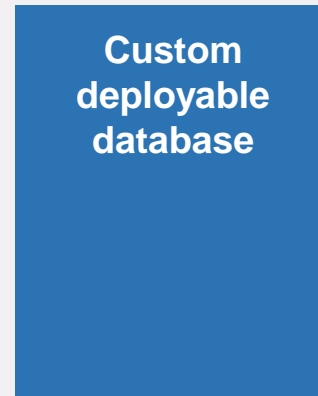
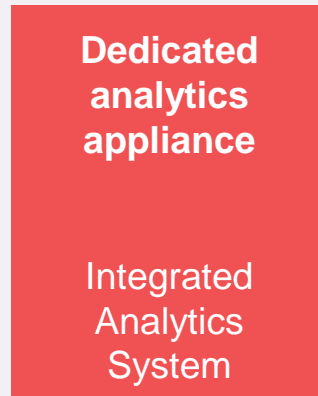


Write Once, Run Anywhere, With a Common SQL Engine

Cloud Ready - IBM Hybrid Data Management Solution

Cloud Ready

- Public Cloud, Private Cloud, 전용 Appliance 이 동일한 S/W Stack으로 구성되어 이식과 연계가 용이



- 신속한 Application 개발**
Write once, run anywhere
One ISV product certification for all platforms
- 운영 호환성**
Reuse operational and housekeeping procedures

- 유연한 라이선스**
Flexible entitlements for business agility and cost-optimization
- 통합**
Common Data Virtualization capabilities for query federation and data movement

Common SQL

- 공통 SQL 사용으로 개발생산성 향상 및 이관 용이

- 표준화된 분석**
Common programming model for in-DB analytics
- 공통 스킬**
One skill set for all deployments
Drive higher efficiencies and portfolio rationalization

Built-in 모니터링 – Home page

IIAS™ 0% 27 12

Type to filter Home

- Home
- Administer >
- Load >
- Run SQL
- Analytics >
- Monitor >
- Settings >
- Connect >
- Help >

Welcome to Integrated Analytics System

Monitor Appliance and Database

- View appliance and database dashboard
- View monitoring data of appliance system
- View monitoring data of database workloads

Administer Databases

- Manage and create tables

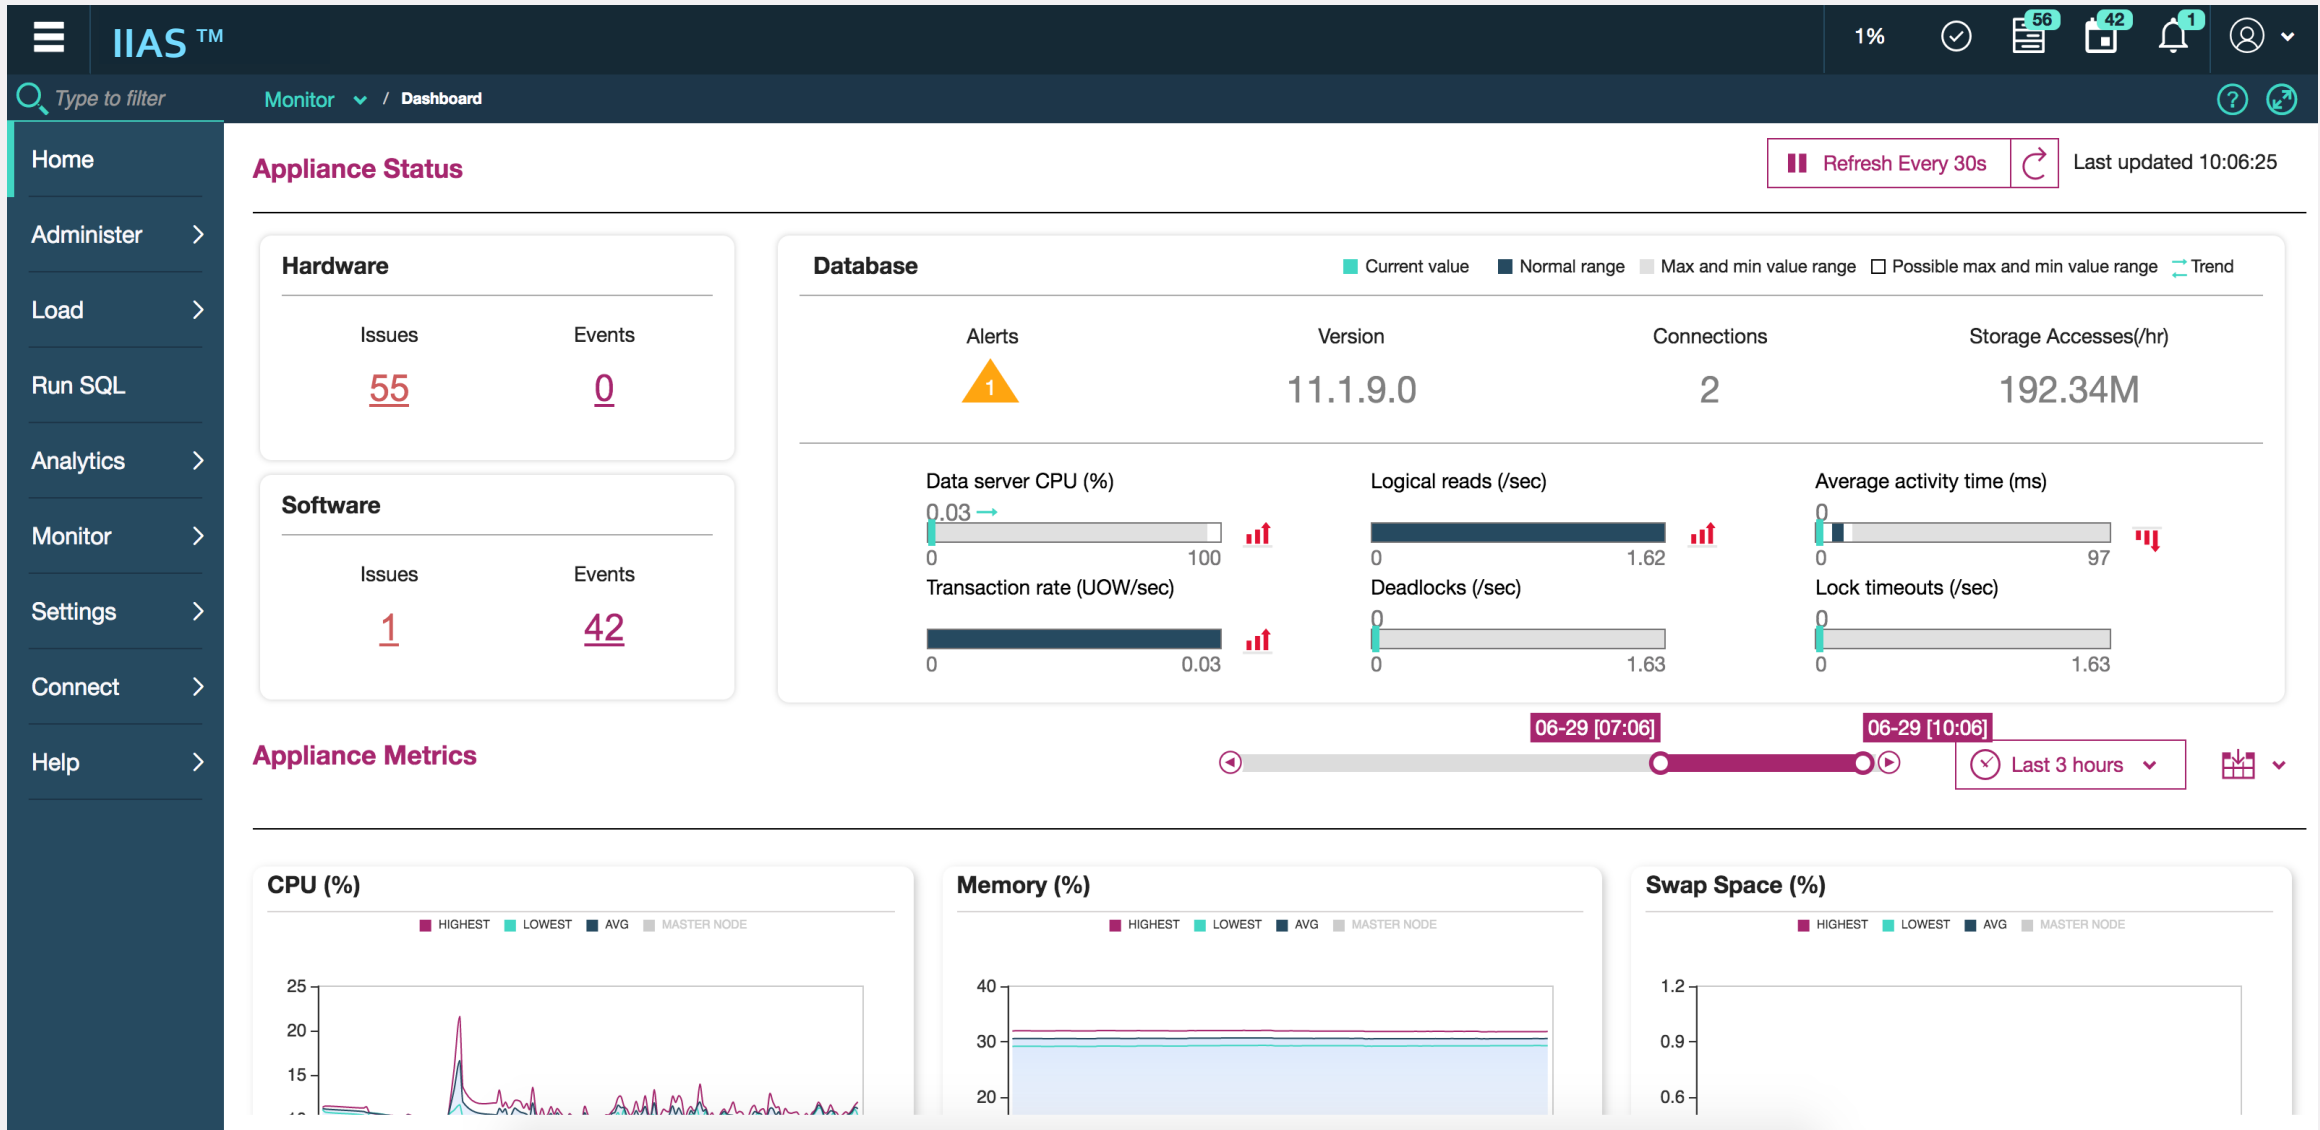
Execute SQL Statements

- Edit and run SQL statements

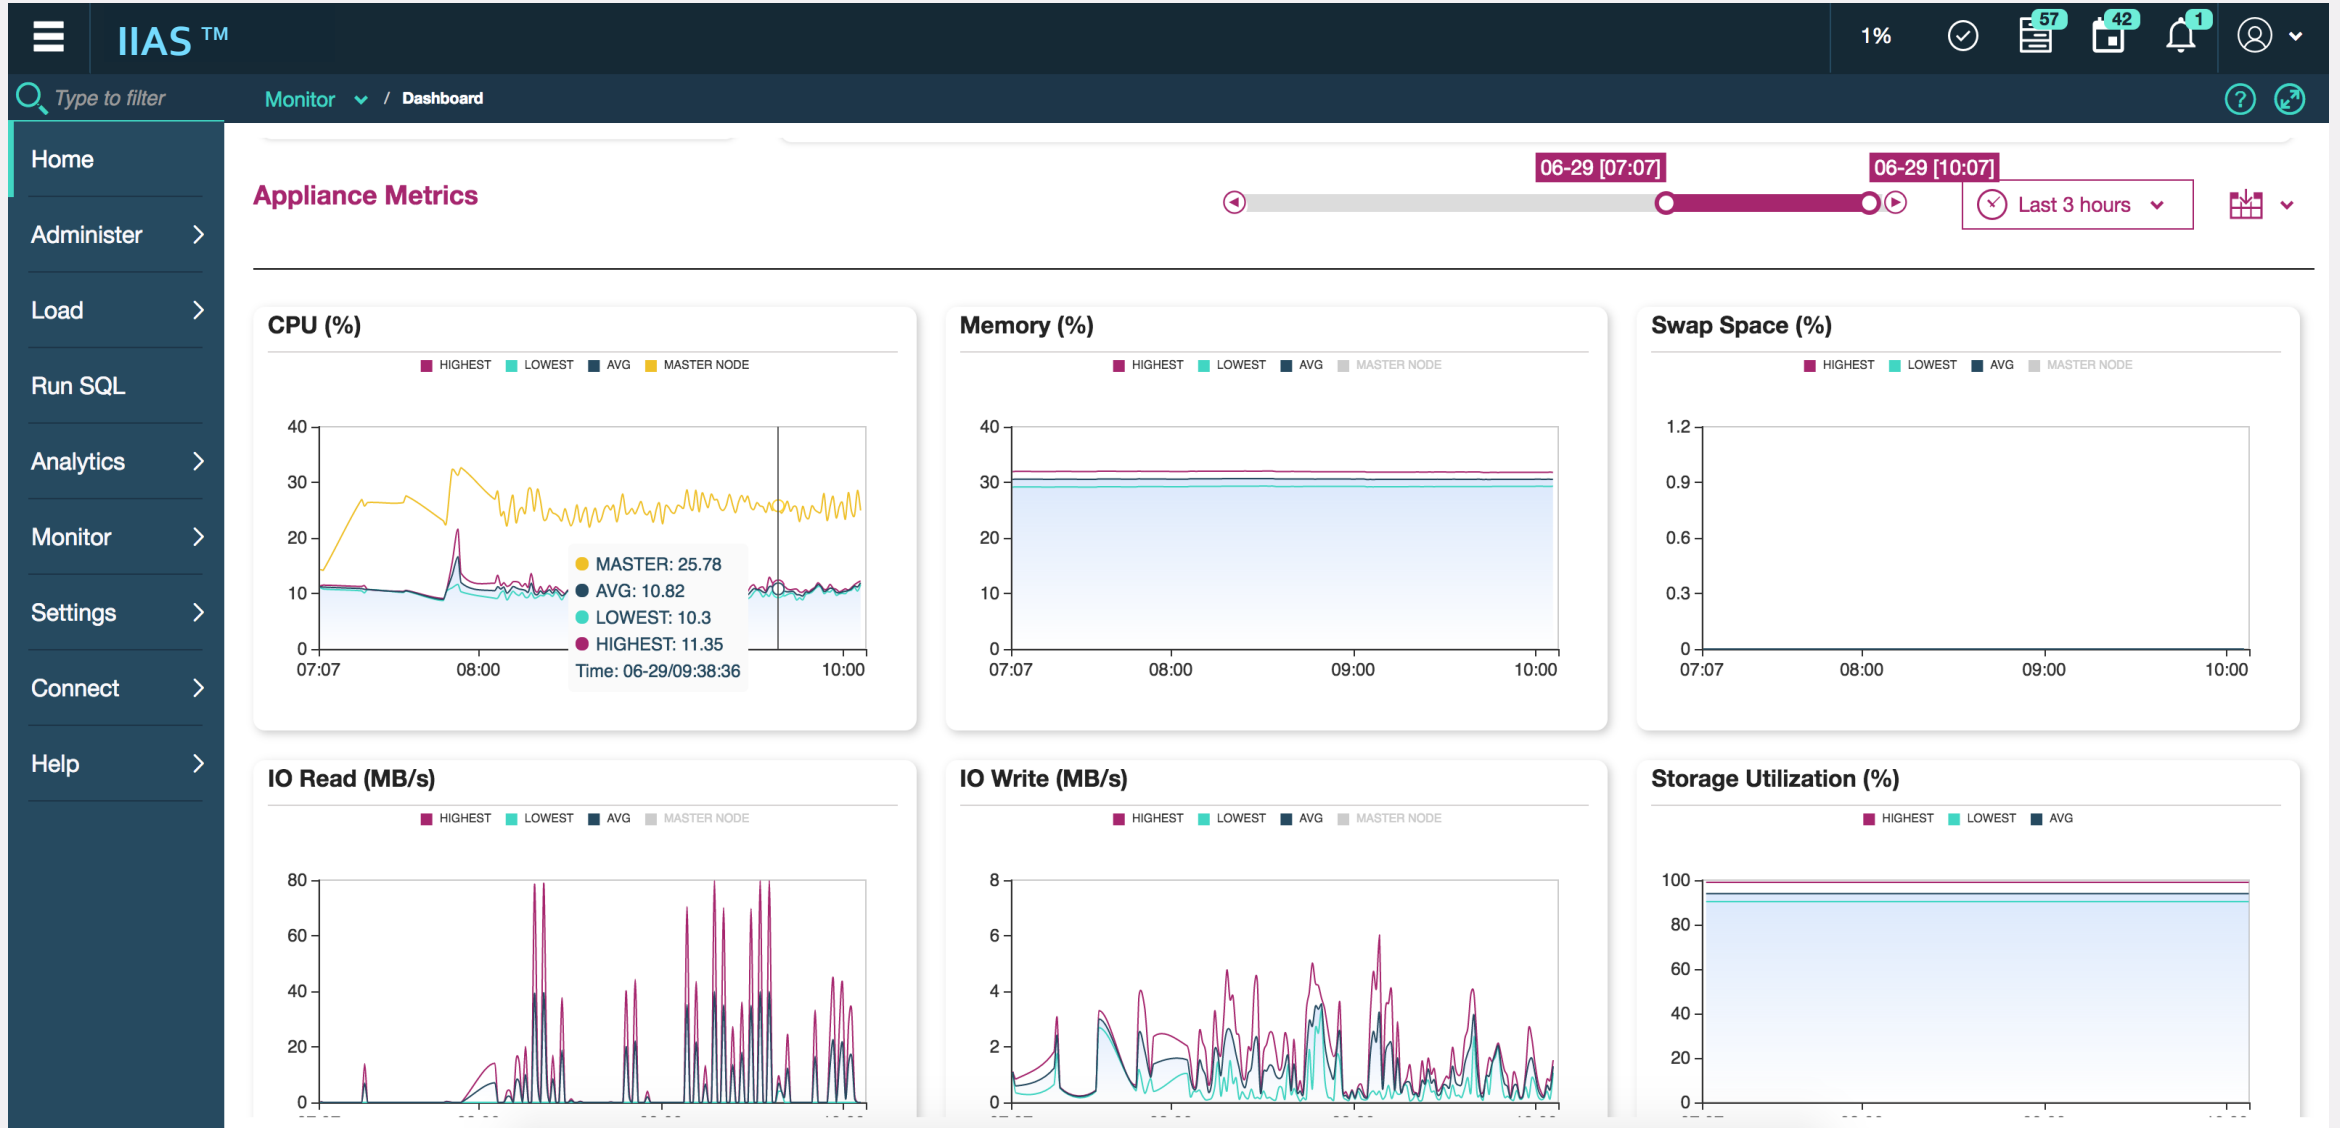
Load data into database

- Load data from Load Hub

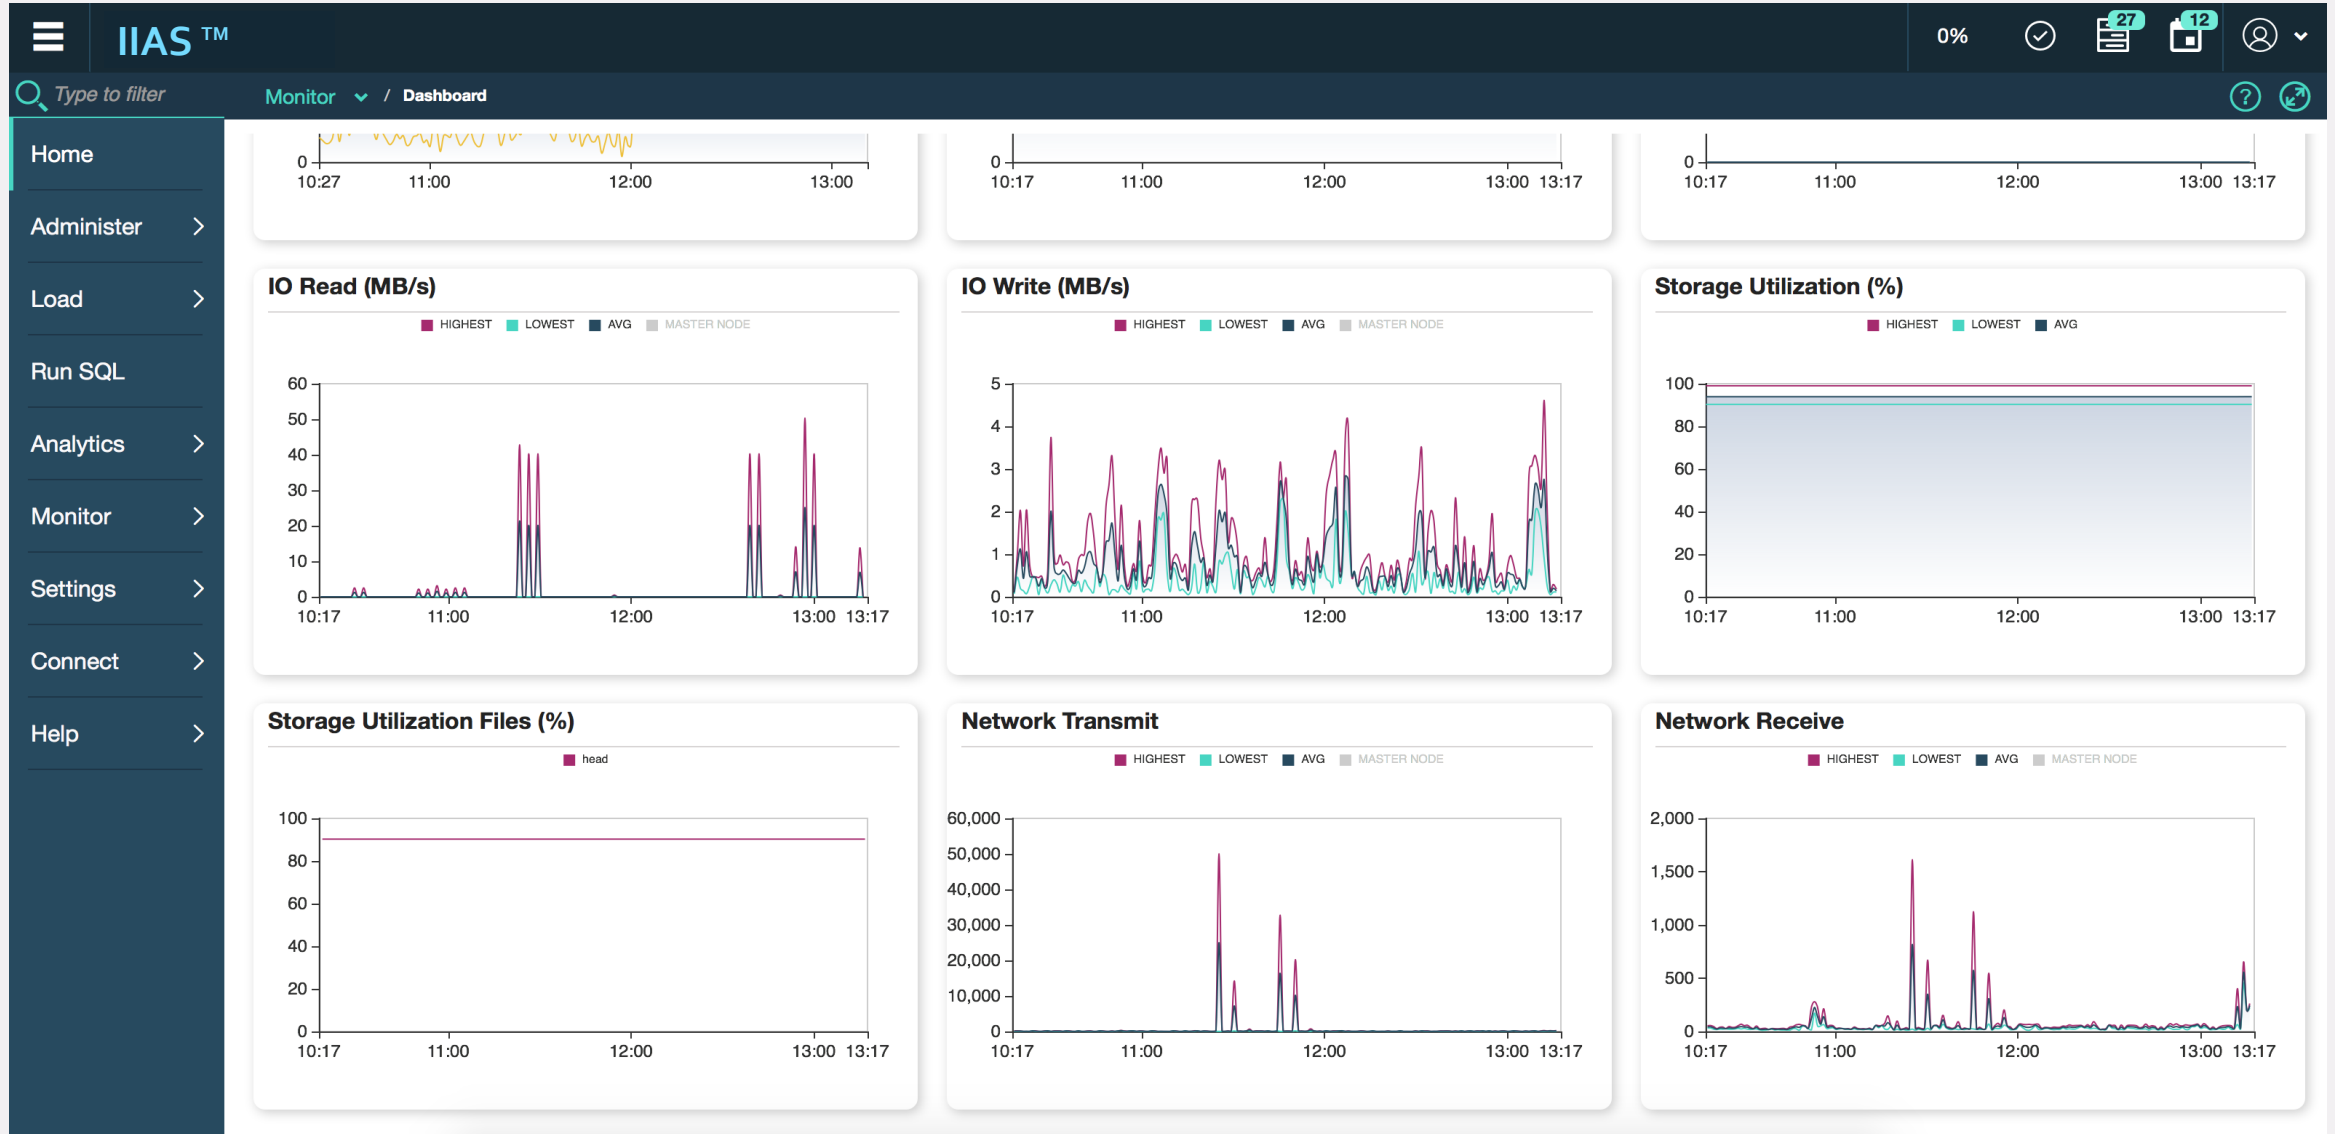
Built-in 모니터링 – Dashboard (Appliance status)



Built-in 모니터링 – Dashboard (Appliance metrics)



Built-in 모니터링 – Dashboard (Appliance metrics) cont.



Built-in 모니터링 – System (Hardware – Rack layout)

The screenshot displays the IIAS™ monitoring interface for a system rack layout. The interface includes a navigation sidebar on the left with options like Home, Administer, Load, Run SQL, Analytics, Monitor, Settings, Connect, and Help. The main content area shows the 'Hardware' tab selected, displaying a rack layout with server slots numbered 2 to 42. A tooltip for 'Node 2' indicates subcomponent failures: 'hw://rack1.node2.cpu4: Subcomponent failed or is not reachable' and 'hw://rack1.node2.cpu2: Subcomponent failed or is not reachable'. To the right, the 'rack1' summary shows 'High Severity Alerts' with three entries, 'Nodes(3)' with a CPU gauge (Lowest Value: 10.15%, AVG Value: 14.27%, Highest Value: 20.19%) and a Memory bar chart (Lowest Value: 35.25%, AVG Value: 38.17%, Highest Value: 40.29%), and 'FlashSystem (1)' and 'PDUs (4)'.

High Severity Alerts

SEVERITY	TITLE	TARGET	START TIME	PMR
Warning	Major component is unrea...	rack1.fcswb	2017-06-27 06:39:09	
Warning	Major component is unrea...	rack1.rpc1	2017-06-27 06:39:09	
Warning	Major component is unrea...	rack1.rpc3	2017-06-27 06:39:09	

Nodes(3)

MAJOR: 10

CPU

- Lowest Value: 10.15%
- AVG Value: 14.27%
- Highest Value: 20.19%

Memory

- Lowest Value: 35.25%
- AVG Value: 38.17%
- Highest Value: 40.29%

FlashSystem (1)

PDUs (4)

Built-in 모니터링 – System (Software)

View and manage the status of your system

Check what components are installed in the system and whether they are working correctly

Refresh Every 30s Last updated 13:31:47

Hardware **Software** Storage Data Slices Notification

LOCATION	STATUS	TYPE	VERSION	DETAILS
spark	OK	application	265	
gpfs	OK	gpfs-cluster	4.2.2.0	
ldap	OK	application	2.4.40	
▶ docker				
▶ os				
▼ dashdb				
dashdb/rack1.node2	OK	application-node		
dashdb/rack1.node3	OK	application-node		
dashdb/rack1.node1	OK	application-node		Leader
▼ docker.dashDB				
docker.dashDB/rack1.node1	OK	docker-container		
docker.dashDB/rack1.node3	OK	docker-container		
docker.dashDB/rack1.node2	OK	docker-container		
▶ ntp				

Built-in 모니터링 – System (Storage)

IIAS™ 0% [27] [12] [User]

Monitor / System

Home Administer > Load > Run SQL Analytics > Monitor > Settings > Connect > Help >

View and manage the status of your system

Check what components are installed in the system and whether they are working correctly

Refresh Every 30s Last updated 13:32:47

Hardware Software **Storage** Data Slices Notification

Utilization

head (%) 1.00 3.37/1024.00 GB Used	data (%) 1.00 109.75/17233.92 GB Used	scratch (%) 1.00 3.60/10240.00 GB Used	local (%) 1.00 2.78/15667.20 GB Used
---	--	---	---

File Systems

File Systems Mount

FILE SYSTEM	DISK	SIZE (GB)	USED (%)	STATUS
platform	nsd.mpath_36005076a81b534...	5120.00	0.05	OK
head	nsd.mpath_36005076a81b534...	1024.00	0.33	OK
scratch	nsd.mpath_36005076a81b534...	10240.00	0.04	OK

Built-in 모니터링 – System (Data Slices)

The screenshot displays the IIAS System Monitoring interface. The top navigation bar includes the IIAS logo, a search bar, and system status indicators (0%, checkmarks, and notification counts). The left sidebar contains navigation options: Home, Administer, Load, Run SQL, Analytics, Monitor, Settings, Connect, and Help. The main content area is titled "View and manage the status of your system" and includes a "Refresh Every 30s" button and a "Last updated 13:35:17" timestamp. Below this, there are tabs for Hardware, Software, Storage, Data Slices (selected), and Notification. A table lists the Data Slices with columns for DATA SLICE, NODE, LOGICAL PORT, and STATUS. All listed slices are in an "UP" status.

DATA SLICE	NODE	LOGICAL PORT	STATUS
14	node0103-fab	0	UP
15	node0103-fab	1	UP
16	node0103-fab	2	UP
17	node0103-fab	3	UP
18	node0103-fab	4	UP
19	node0103-fab	5	UP
20	node0103-fab	6	UP
21	node0103-fab	7	UP
6	node0102-fab	0	UP
7	node0102-fab	1	UP
8	node0102-fab	2	UP
9	node0102-fab	3	UP
10	node0102-fab	4	UP
11	node0102-fab	5	UP

Built-in 모니터링 – System (Notification)

☰
IIAS™
1% 🕒 📄 56 📅 42 🔔 1 👤

🔍 Type to filter
Monitor ▾ / System
? ↻

Home
Administer >
Load >
Run SQL
Analytics >
Monitor >
Settings >
Connect >
Help >

View and manage the status of your system

Check what components are installed in the system and whether they are working correctly

⏸ Refresh Every 30s
 ↻
Last updated 10:16:39

Hardware
Software
Storage
Data Slices
Notification

View Details
Show Open Issues ▾
🔍 📄 ⌵ ⬆

<input type="checkbox"/>	SEVERITY	CATEGORY	STATEFUL	TYPE	TITLE	START TIME	END TIME	TARGET	TARGET TYPE
<input type="checkbox"/>	MAJOR	Hardware	Issues	HW_SERVICE_REQUESTED	Subcomponent failed or is...	2017-06-26 1...	--	rack1.fabswa.port16	port
<input type="checkbox"/>	MAJOR	Software	Issues	SW_NEEDS_ATTENTION	Application container can...	2017-06-28 1...	--	docker.dashDB/rack1.node3	docker-container
<input type="checkbox"/>	MAJOR	Hardware	Issues	HW_SERVICE_REQUESTED	Subcomponent failed or is...	2017-06-26 1...	--	rack1.fabswa.port50	port
<input type="checkbox"/>	MAJOR	Hardware	Issues	HW_SERVICE_REQUESTED	Subcomponent failed or is...	2017-06-26 1...	--	rack1.fabswb.port49	port
<input type="checkbox"/>	MAJOR	Hardware	Issues	HW_SERVICE_REQUESTED	Subcomponent failed or is...	2017-06-26 1...	--	rack1.fabswb.port50	port
<input type="checkbox"/>	MAJOR	Hardware	Issues	HW_SERVICE_REQUESTED	Major component is unrea...	2017-06-26 0...	--	rack1.fsn1	fsn
<input type="checkbox"/>	MAJOR	Hardware	Issues	HW_SERVICE_REQUESTED	Subcomponent failed or is...	2017-06-26 1...	--	rack1.fabswa.port15	port
<input type="checkbox"/>	MAJOR	Hardware	Issues	HW_SERVICE_REQUESTED	Subcomponent failed or is...	2017-06-27 0...	--	rack1.node1.cpu2	cpu
<input type="checkbox"/>	MAJOR	Hardware	Issues	HW_SERVICE_REQUESTED	Subcomponent failed or is...	2017-06-26 1...	--	rack1.fabswa.port17	port
<input type="checkbox"/>	MAJOR	Hardware	Issues	HW_SERVICE_REQUESTED	Subcomponent failed or is...	2017-06-26 1...	--	rack1.fabswa.port18	port
<input type="checkbox"/>	MAJOR	Hardware	Issues	HW_SERVICE_REQUESTED	Subcomponent failed or is...	2017-06-26 1...	--	rack1.fabswb.port18	port
<input type="checkbox"/>	MAJOR	Hardware	Issues	HW_SERVICE_REQUESTED	Subcomponent failed or is...	2017-06-26 1...	--	rack1.fabswb.port15	port
<input type="checkbox"/>	MAJOR	Hardware	Issues	HW_SERVICE_REQUESTED	Subcomponent failed or is...	2017-06-26 1...	--	rack1.fabswb.port16	port

Built-in 모니터링 – Call Home

IIAS™ 1% [Icons: 56, 42, 1]

Type to filter Settings / Call Home

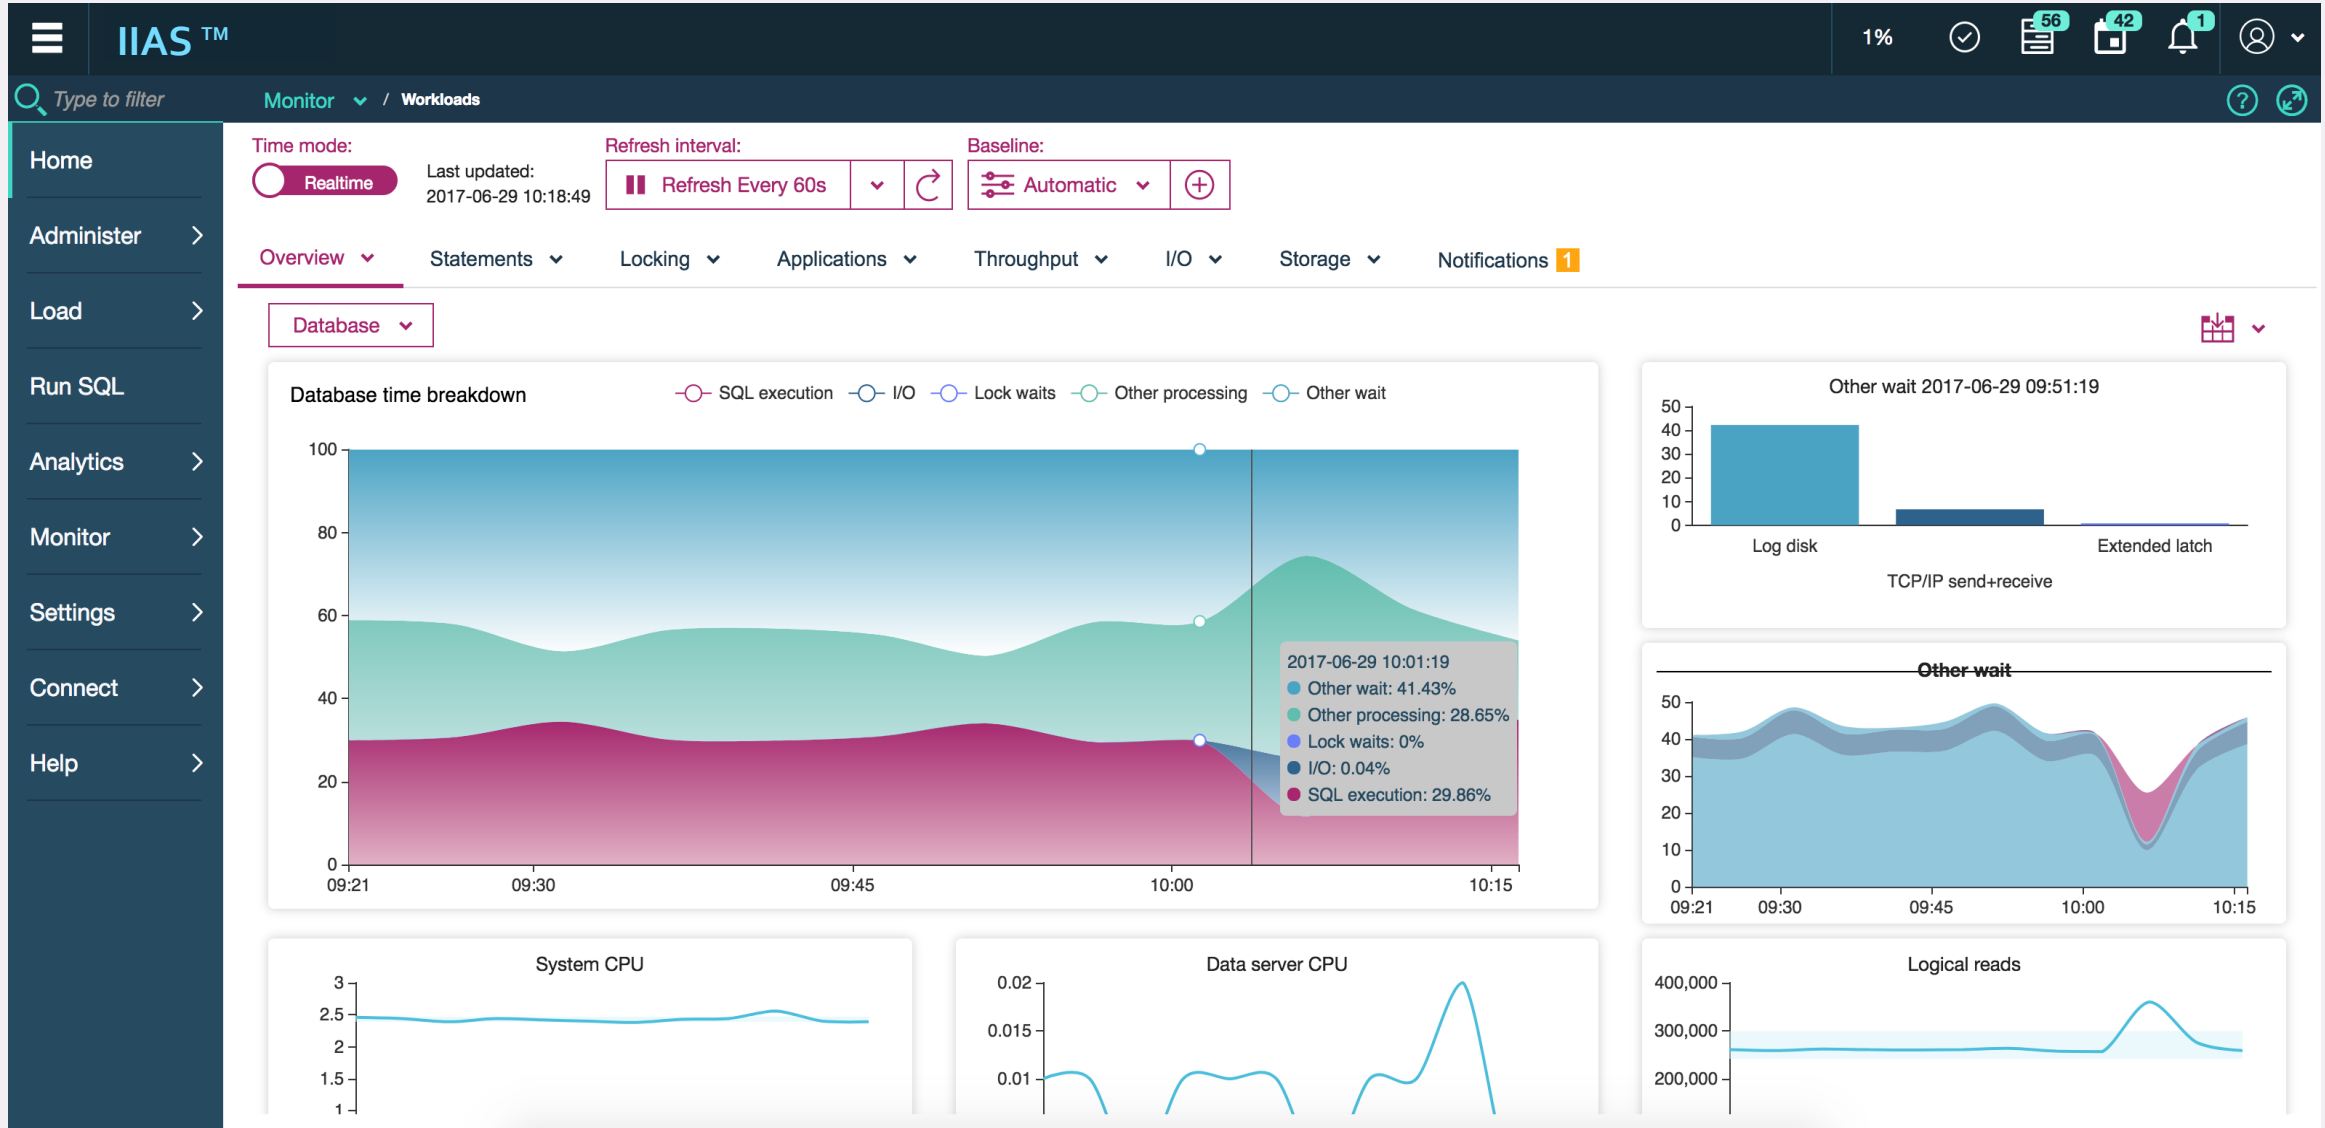
IBM Call Home Settings

Call Home provides communications mechanism to securely send Information from customers systems to IBM.

Customer Information	*Company	TEST CUSTOMER
Contacts	*Address1	ADDRESS 1
Email	*Address2	ADDRESS 2
System	Address3	
Rules		
Advanced	*ICN	575611
Control		

Save Cancel

Built-in 모니터링 – Workloads (Overview)



Built-in 모니터링 – Workloads (Notification)

The screenshot shows the IBM IAS Workloads monitoring interface. The top navigation bar includes the IAS logo, a search filter, and various status indicators (1%, checkmark, 56 notifications, 42 calendar items, 1 alert, user profile). The main content area is divided into a left sidebar with navigation options (Home, Administer, Load, Run SQL, Analytics, Monitor, Settings, Connect, Help) and a main workspace. The workspace shows the 'Monitor' view for 'Workloads'. It includes controls for 'Time mode' (Realtime), 'Refresh interval' (Refresh Every 60s), and 'Baseline' (Automatic). A notification banner indicates '>163384' notifications. The selected notification shows a SQL statement: 'select partition_key,activate_timestamp,activity_collected,activity_id,agent_id,appl_id,application_name,client_acctng,client_ap...'. Below the statement, the 'Analysis' section states: 'The current performance of the statement deviated from its past performance.' The 'Recommended Actions' section lists:

- Get advice and recommendations to help improve performance for a workload (group of statements).
- View access plan - You can use Visual Explain to view the access plan for explained SQL statements as a graph. You can use the information available from the graph to tune your statements for better performance.
- If statement is still running, cancel the statement.
- Investigate the statement.

 The 'History' section notes: 'This alert was not generated previously for this database.' On the right, the 'Details' table compares current performance with historical data:

	CURRENT	HISTORI...	DEVIATION
Average CPU time	42,955	40,479	↑2,476 microseconds
Average activity wait time	1,010	88	↑922 milliseconds
Average statement execution time	1,068	128	↑940 milliseconds
Average lock wait time	0	0	
Number of executions	1	249	
Average rows read	0	0	
Average physical reads	0	0	
Average logical reads	118	115	↑3
Average temporary reads	88	86	↑2
Average sort overflows	0	0	
Average rows returned	0	0	
Average lock escalations	0	0	
Average buffer pool data logical reads	30	29	↑1

Built-in 모니터링 – Workloads (Statements history)

IIAS™ 12%

Monitor / Workloads

Time mode: History 2017-06-26 10:27:00 10:27 10:27 2017-06-29 10:27:00

Data scope: Last 24 hours Baseline: Automatic

Overview Statements Locking Applications Throughput I/O Storage Notifications

View Details Explain Hide System Statements Average Generate report

<input type="checkbox"/>	SQL	NUMBER OF EXECUTIONS	AVERAGE STATEMENT EXECUTION TIME	AVERAGE CPU TIME	AVERAGE ROWS READ	AVERAGE ROWS RETURNED	AVERAGE SORT OVERFLO...
<input type="checkbox"/>	/* IBM_DSSNAP */ VALUE...	655,306	0.000	0.000	0	1	0
<input type="checkbox"/>	select dsmAlertID from ib...	306,105	0.000	0.000	58	0	0
<input type="checkbox"/>	select name, property_val...	305,295	0.000	0.000	16	16	0
<input type="checkbox"/>	select id, userid, lastedit...	305,007	0.000	0.000	0	0	0
<input type="checkbox"/>	select alert_type, severity, ...	304,710	0.000	0.000	0	0	0
<input type="checkbox"/>	INSERT INTO IBM_SAILF...	24,000	0.000	0.000	0	0	0
<input type="checkbox"/>	SELECT COUNT(TRIGG...	15,006	0.000	0.000	0	1	0
<input type="checkbox"/>	select versionnumber,curr...	12,403	0.000	0.000	1	1	0
<input type="checkbox"/>	with unicode_fix(x) as (sel...	10,738	0:00.001	0:00.001	0	0	0
<input type="checkbox"/>	INSERT INTO IBMOTS.S...	10,393	0.000	0.000	0	0	0
<input type="checkbox"/>	insert into ibmots.sql_fact ...	10,357	0.000	0.000	0	0	0
<input type="checkbox"/>	call SYSIBM.SQLGAMES...	10,215	0.000	0.000	0	1	0
<input type="checkbox"/>	SELECT ID, SEVERITY F...	9,820	0.000	0.000	0	0	0
<input type="checkbox"/>	SELECT CREATE_TIME ...	7,920	0.000	0.000	0	0	0

Built-in 모니터링 – Administer database objects

The screenshot displays the IAS (IBM Analytics Server) administration console. The top navigation bar includes the IAS logo, a search bar, and several utility icons. The main content area is titled 'Administer / Tables' and features a 'Create' button. A table lists database objects with columns for Schema and Organization. The 'Application Objects' menu is expanded, showing options like Stored Procedures, User-Defined Types, and User-Defined Functions. A table with 10 rows is visible, all with 'ACSN_EXTTAB' in the Schema column and 'COLUMN' in the Organization column. The bottom right corner shows 'Retrieve rows: 500 / Maximum rows: 500'.

Schema	Organization
ACSN_EXTTAB	COLUMN
ACSN_EXTTAB	COLUMN
ACSN_EXTTAB	COLUMN
ACSN_EXTTAB	COLUMN
ACSN_EXTTAB	COLUMN
ACSN_EXTTAB	COLUMN
ACSN_EXTTAB	COLUMN
ACSN_EXTTAB	COLUMN
ACSN_EXTTAB	COLUMN
ACSN_EXTTAB	COLUMN
ACSN_EXTTAB	COLUMN
ACSN_EXTTAB	COLUMN
ACSN_EXTTAB	COLUMN
ACSN_EXTTAB	COLUMN
ACSN_EXTTAB	COLUMN
ACSN_EXTTAB	COLUMN
ACSN_EXTTAB	COLUMN
ACSN_EXTTAB	COLUMN
ACSN_EXTTAB	COLUMN
ACSN_EXTTAB	COLUMN
ACSN_EXTTAB	COLUMN
ACSN_EXTTAB	COLUMN

Built-in 모니터링 – Administer Table (Privilege management)

The screenshot shows the IAS Administer Table interface. A 'Grant privileges' dialog box is open, allowing the user to grant privileges to a selected user. The dialog has a title bar with a close button (X) and a 'Grant privileges' button. Below the title, there are radio buttons for 'User', 'Group', and 'Role', with 'User' selected. A search bar is present. The main area is divided into two columns: 'All' and 'Select privileges below to grant'. The 'All' column lists users: BLUADMIN_GUARDIUM, BLUADMIN_MON, BLUUSER (checked), and DB2INST1. The 'Select privileges below to grant' column lists various privileges with checkboxes and 'With Grant' options: CONTROL, SELECT (checked), INSERT (checked), UPDATE, DELETE, ALTER, and REFERENCES. At the bottom right of the dialog are 'OK' and 'Cancel' buttons.

Table Explorer: Grant/Revoke privileges

NAME	SCHEMA
ADD1_ERR	ACSN_EXTTAB

Grant privileges

User Group Role

All

BLUADMIN_GUARDIUM

BLUADMIN_MON

BLUUSER

DB2INST1

Select privileges below to grant

CONTROL With Grant

SELECT With Grant

INSERT With Grant

UPDATE With Grant

DELETE With Grant

ALTER With Grant

REFERENCES With Grant

OK Cancel

M4002 모델 용량



서버 & Storage
 IBM S822L 24 core server 3.02GHz (RedHat Linux)
 IBM FlashSystem 900

확장 방식
 In-place Expansion, Tiered storage

네트워크
 Mellanox 10G Ethernet switches
 Brocade SAN switches

- **확장(증설)지원**
 - 선형 성능 확장 : 상위 model 로 증설 가능
 예, 1/3 Rack → 2/3 Rack
 예, 1/3 Rack → Full Rack
 - 스토리지 확장 : Tiered Storage 추가
 예, 1/3 Rack → 1/3 Rack with TS,
 예, 1/3 Rack → Full Rack with TS,

Model	1/3 Rack	1/3 Rack Tiered Storage	2/3 Rack	2/3 Rack Tiered Storage	Full Rack	Full Rack Tiered Storage
Server	3	3	5	5	7	7
CPU cores	72	72	120	120	168	168
Memory	1.5TB	1.5TB	2.5TB	2.5TB	3.5TB	3.5TB
¹⁾ Usable User Disk	Flash 27TB	Flash 59TB HDD 166TB	Flash 54TB	Flash 86TB HDD 299TB	Flash 81TB	Flash 113TB HDD 432TB

1) DBMS Log 저장 등이 제외된 순수 User Data 영역



IBM IAS PoC 오퍼링

- **대상** : PDA 사용 고객사 및 도입을 검토중인 고객
- **검증기간** : 요구되는 기능의 숫자와 데이터 사이즈, 기능의 구현 수준과 연계 대상에 따라서 기간은 협의가 필요합니다. 최소 2주~ 최대 4주 로 예상 됩니다.
- **검증 영역**
 - **성능** : 기존 PDA와 로딩 및 조회 성능 비교
 - **Data Science를 위한 기능 (Data Science Experience)**
 - **솔루션 연계** : 기존 연계 솔루션과의 호환성 검증

Tasks	Week 1	Week 2	Week 3	고객 참여자	비고
Kick-Off 미팅 (진행 일정/ 참석자 확인)				Sponsor, 전체	
검증 시나리오 협의				고객 담당자	
장비 배송 및 환경 구성				고객 담당자	
성능 및 기능 검증				고객 담당자	
최종 보고				Sponsor, 전체	

예시

Legal Disclaimer

- © IBM Corporation 2018. All Rights Reserved.
- The information contained in this publication is provided for informational purposes only. While efforts were made to verify the completeness and accuracy of the information contained in this publication, 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 publication or any other materials. Nothing contained in this publication 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 the applicable license agreement governing the use of IBM software.
- References in this presentation to IBM products, programs, or services do not imply that they will be available in all countries in which IBM operates. Product release dates and/or capabilities referenced in this presentation may change at any time at IBM's sole discretion based on market opportunities or other factors, and are not intended to be a commitment to future product or feature availability in any way. Nothing contained in these materials is intended to, nor shall have the effect of, stating or implying that any activities undertaken by you will result in any specific sales, revenue growth or other results.
- 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.
- All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.
- All references to fictitious companies are used for illustration purposes only.

