Peter Kohlmann <kohlmann@ca.ibm.com>

Tridex – Console Strategy

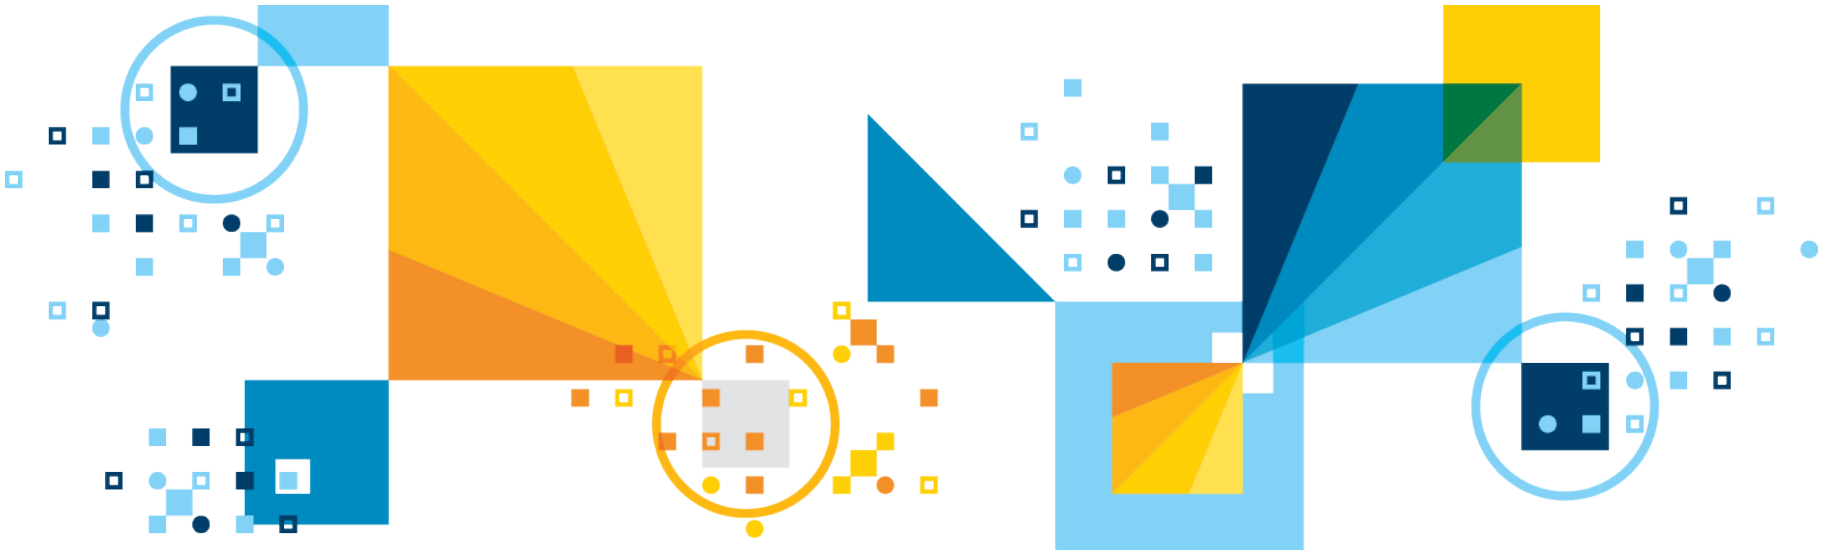

Differentiate the IBM Database Solutions with Core Database Tools for the Data Lifecycle

Backup and Recovery

Design Data Models

Advanced Recovery Feature

InfoSphere Data Architect

Monitor and Optimize Workloads

Access and Connect to Data

DB2 Connect
DS Driver
DS Gateway for ODATA

Data Server Manager

Data Studio
& pureQuery

Develop Data Applications

Administer and Manage Databases

Run and Optimize SQL

Debug and Deploy

```
select * from employee;
```

Today

Data Server Manager Enterprise: Core Capabilities

Monitoring and Alerts

- Proactive alerts
- Overview Dashboards and drilldown
- Real-time, and historical performance data
- Client Monitoring
- HADR and pureScale

Query Tuning

- Advisors to identify tuning opportunities
- Capture SQL and workloads
- Explain query access paths
- Improve statistics quality, improve database and query design

Database Administration and Use

- Build, format, explain and run SQL statements and download as excel
- Perform or schedule routine admin
- Manage privileges, maintenance and availability
- Connect to remote data sources
- Create, and mange database objects

Configuration Management and History

- Keep an inventory of database objects history of change.
- Registry variables, configuration settings, custom key/ value settings to help with diagnosis

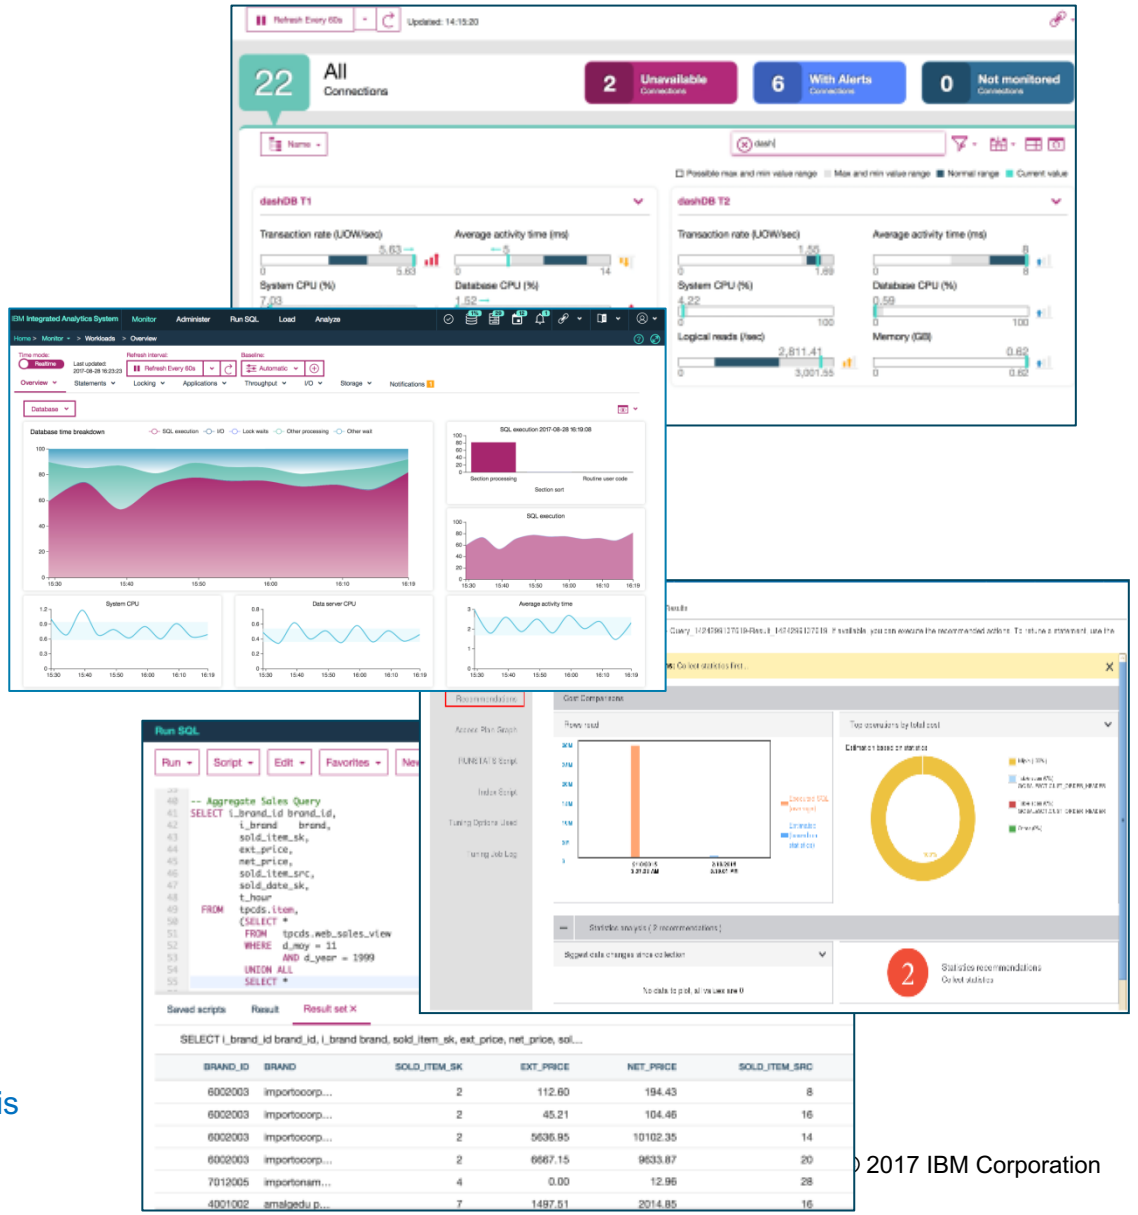

Consoles Powered by Data Server Manager

Cloud Convenience with Appliance Simplicity – Convergence Goal

Walk up and use management

- Alerts
- Key Performance Indicators
- Problem determination
- Hardware Monitoring (Sailfish)
- Software Services (Sailfish)
- Storage
- Call home (Sailfish)

Walk up and use deep dive database monitoring and problem determination

- Database alerts and KPIs
- Active hotspot drill down for in depth monitoring
- SQL History - flight recorder
- Workload monitoring

Walk up and use development and data management

- Load data
- Database administration and security
- Run SQL with Visual Explain
- Download query results as csv or xlsx
- Multiple users with Internal or External LDAP
- Federation and Fluid Query

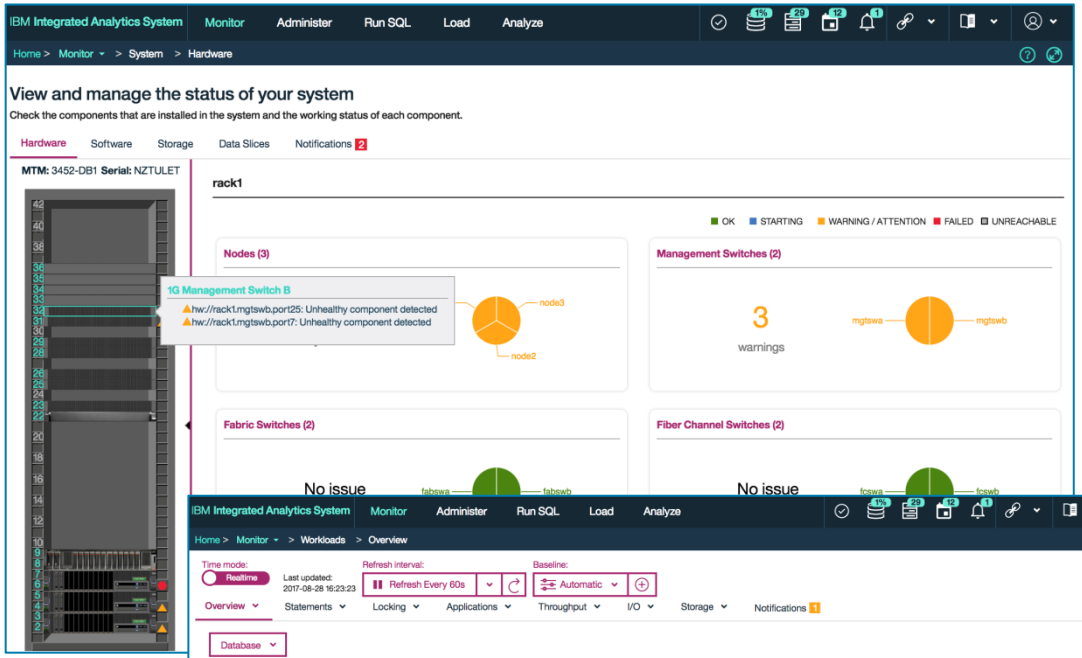

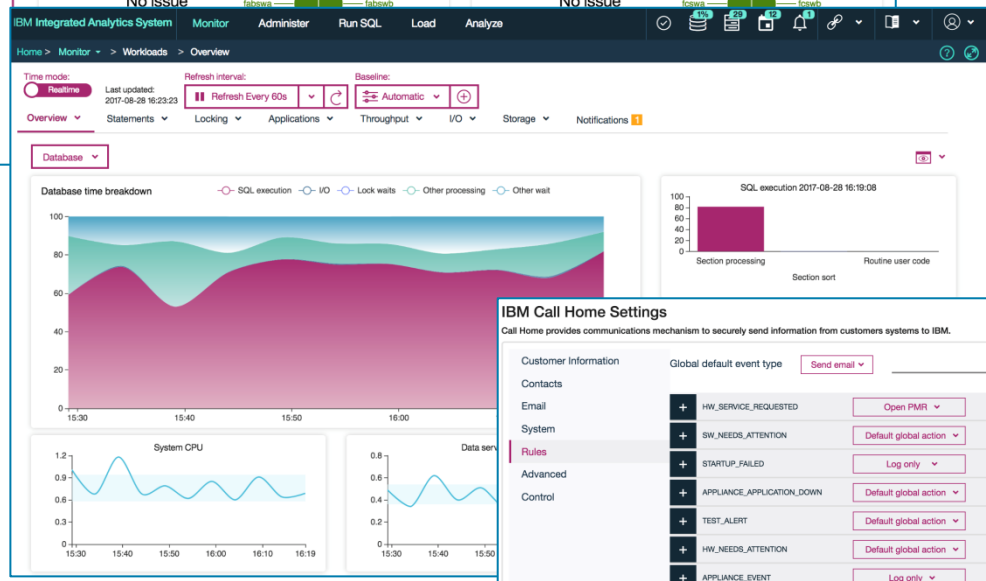

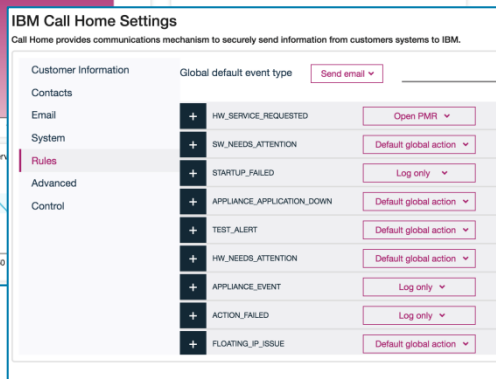

On Cloud Data Management Consoles

Consoles Today

Monitoring and Alerts

- Realtime Data
  - Storage
  - Connections
  - Running SQL

Administration and Use

- User Management
- Load Data from Cloud or Desktop
- Build, format, explain and run SQL statements and download as excel
- Backup and Restore
- Manage privileges
- Explore tables
- Storage by Schema and Tables

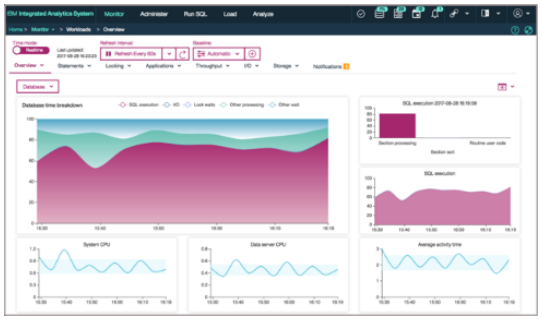

Consoles 3Q

Monitoring and Alerts

- Overview Dashboards
  - Availability
  - Responsiveness
  - Throughout
  - Resource Usage
  - Contention
  - Time Spent
- Realtime Data
  - WLM Queue
  - Storage
  - Connections
  - Running SQL

Administration and Use

- User Management
- Load Data from Cloud or Desktop
- Build, format, explain and run SQL statements and download as excel
- Backup and Restore
- Manage privileges
- Explore tables

DSM for Db2WoC Today

Monitoring and Alerts

- Enterprise view of all your HDM systems
- Proactive alerts
- Real-time, and historical performance data
- Over 20 overview KPIs
- Detailed Monitoring Drilldown Screens
  - In-Flight executions
  - Package cache
  - Stored procedures
  - Individual executions
  - Blocking and waiting connections
  - Connections
  - Locked objects with waiting connections
  - Find locked objects
  - Top consumers
  - Utilities
  - Member summary
  - WLM workload & service class summary
  - Operating system time spent
  - Buffer pools
  - Prefetchers
  - Logging performance
  - Table performance
  - Table space utilization

Administration and Use

- Build, format, explain and run SQL statements and download as excel
- Perform or schedule routine admin
- Manage privileges
- Connect to remote data sources
- Create, and mange all database objects

Data Server Manager for Consoles

- Data Server Manager let you manage all your IBM Hybrid Data Management servers from a single pane of glass
- Alerts and Key Performance Indicators at a glance

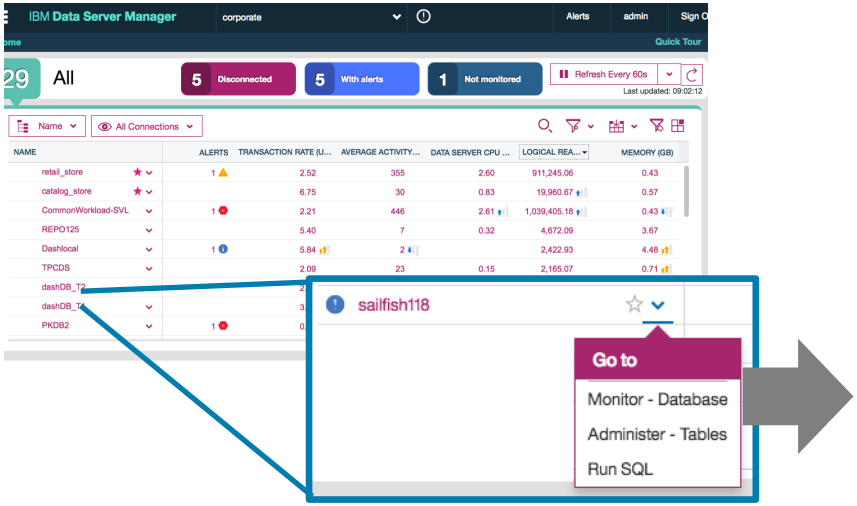

- Dive into detailed monitoring and management for individual databases and servers with a single click
  - Db2 for LUW
  - Db2 on Cloud
  - Db2 Warehouse Private
  - Db2 on Cloud
  - IIAS

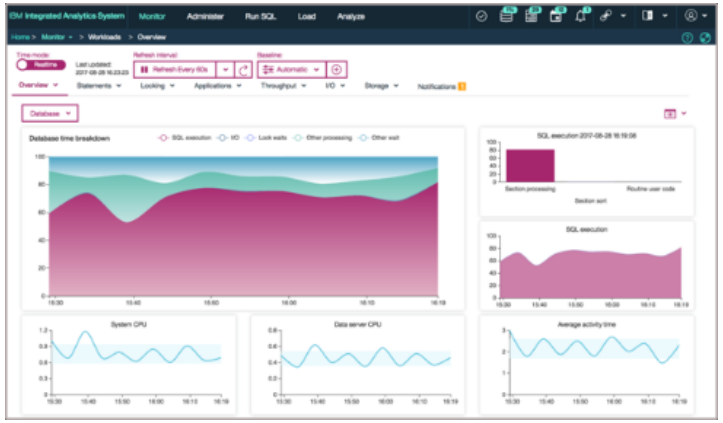

Data Server Manager resources

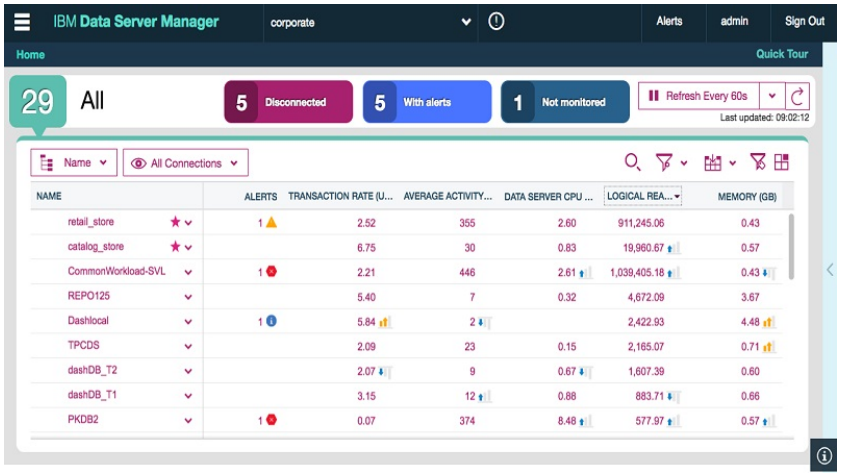

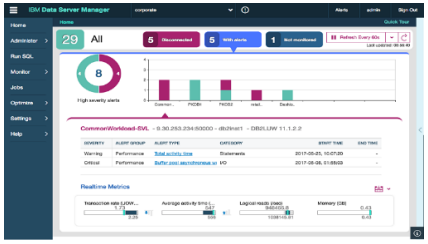

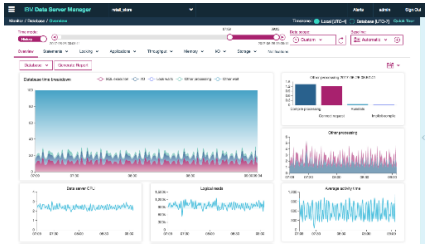

Unlock DSM's full potential by accessing our resources section

Download

Get the latest version of DSM from the marketplace page - http://ibm.biz/IWantMyDSM

Content Catalogue

A video on how to manager the hybrid enterprise with DSM
DSM Community
"How to" videos on Data Server Manager
DSM blogs – "How to" articles
DSM Security Guide and Knowledge Center Guide
DSM Sample code (Jupyter Notebook)
DSM Capacity Planning Guide

Data Server Manager Blog - ibm.biz/DSMBlog

Articles on:

- Mining the DSM History
- Access to Db2 from Python Notebooks
- Historical vs Realtime Monitoring
- What's new in DSM
- Accessing Remote tables to build a virtual warehouse
- How DSM determines what is normal
- Individual Statement History
- Workload Monitoring

New Db2 Extensions make using Jupyter Notebooks a snap

PeterKohlmann | | 347 Views

I just uploaded a new version of the Jupyter notebook that is used in my Blog posting on how to mine history data from the DSM repository. The new version uses Jupyter extensions (magic commands) for Db2 written by George Baklarz. Have a look, you don't have to hard code a userid and password into the notebook anymore. Running SQL is much easier and you can format multiline SQL. The new extensions make using a notebook much easier and faster with Db2. Have a look at https://github.com/DB2-Samples/db2jupyter for instructions and a tutorial on using the extensions. It includes instructions for setting up your own notebooks or you can download the DSX Desktop from IBM.

If you want to see DSM mining and the Jupyter notebook in action, I will be at IDUG EMEA in October and presenting as part of the free Developer Workshop on Thursday. To sign up follow the link: ibm.biz/db2wkshp.

```
In [5]: %%sql -pb
        SELECT WORKDEPT, COUNT(*)
          FROM EMPLOYEE
        GROUP BY WORKDEPT
```

```
<matplotlib.figure.Figure at 0x7f41aee06ef0>
```

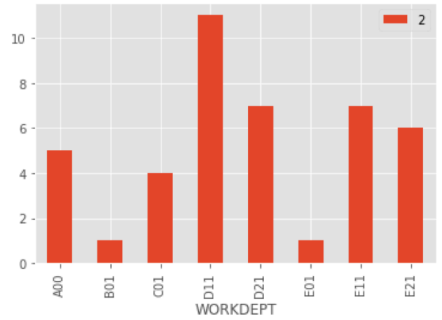

Latest Article: "Can DSM monitor itself?"

Unified Console Vision for Hybrid Data Management

2018-2019

Adaptive

An experience that changes when you need to

- Adapts to different users and their behavior
  DBA, SYS Admin, App Dev, Data Scientist
- Move data and applications between ground and cloud
- Control changing workloads and competing users

Cognitive

Change how you manage your system

- Predicts so you don't have to react
- Understand normal vs abnormal systems
- Manage your systems like a Data Scientist
- Build Apps in Minutes

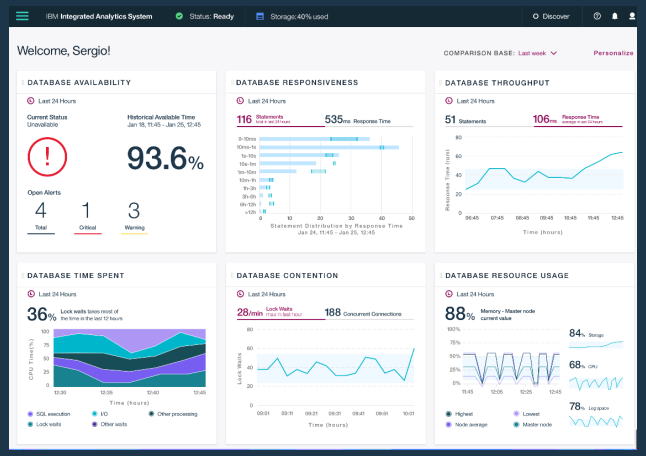

Inclusive

One experience for:

- Db2 on Prem & Cloud
- Db2 Warehouse on Prem & Cloud
- BigSQL
- IIAS
- ICP for Data

Virtual Data Access

- Make remote data act like local tables
- Wide data source support

Open Integration

- Open Source databases: Postgres and Mongo
- Hadoop through BigSQL

Enterprise Ready

The cockpit for the IBM Hybrid Data Enterprise

- Enterprise-wide seamless experience from ground to cloud
- See your whole enterprise on one pane of glass
- Ensure availability, security and recoverability

Unified Console Development Direction

- Design driven unified experience across HDM
- Composible services architecture
- Delivered to fit the offering and the user
- Upgrade from Data Server Manager like a fixpack

Coming Soon

Rapid Triage and Personalization – Today on IIAS

tri·age (trēˈäZH) – Assign a degree of urgency

Each widget tells a story with several metrics

Is my system available? running short or long statements? doing a lot or a little work? in conflict? running out of resources? spending time in the wrong area?

Together the widgets tell the whole story

What Next?

- What is Normal
- Disaster Recovery
- Custom Widgets
- More personas

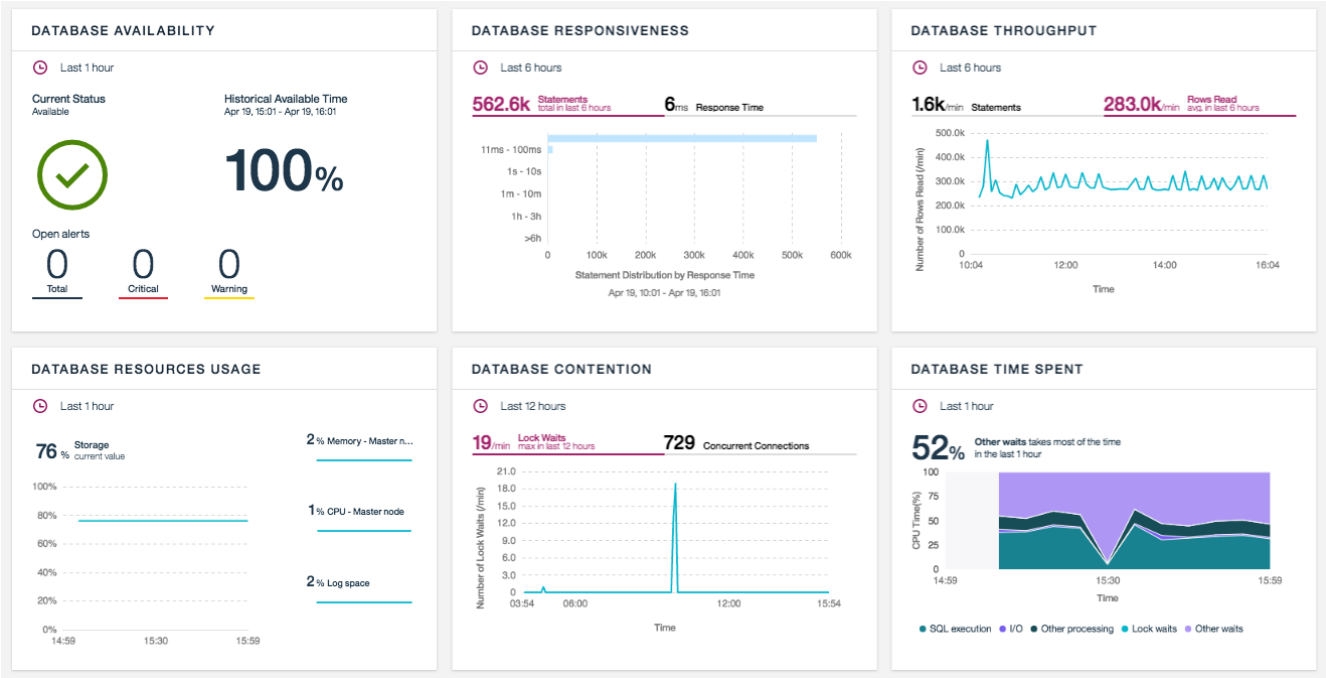

Backup and Restore – today on IIAS

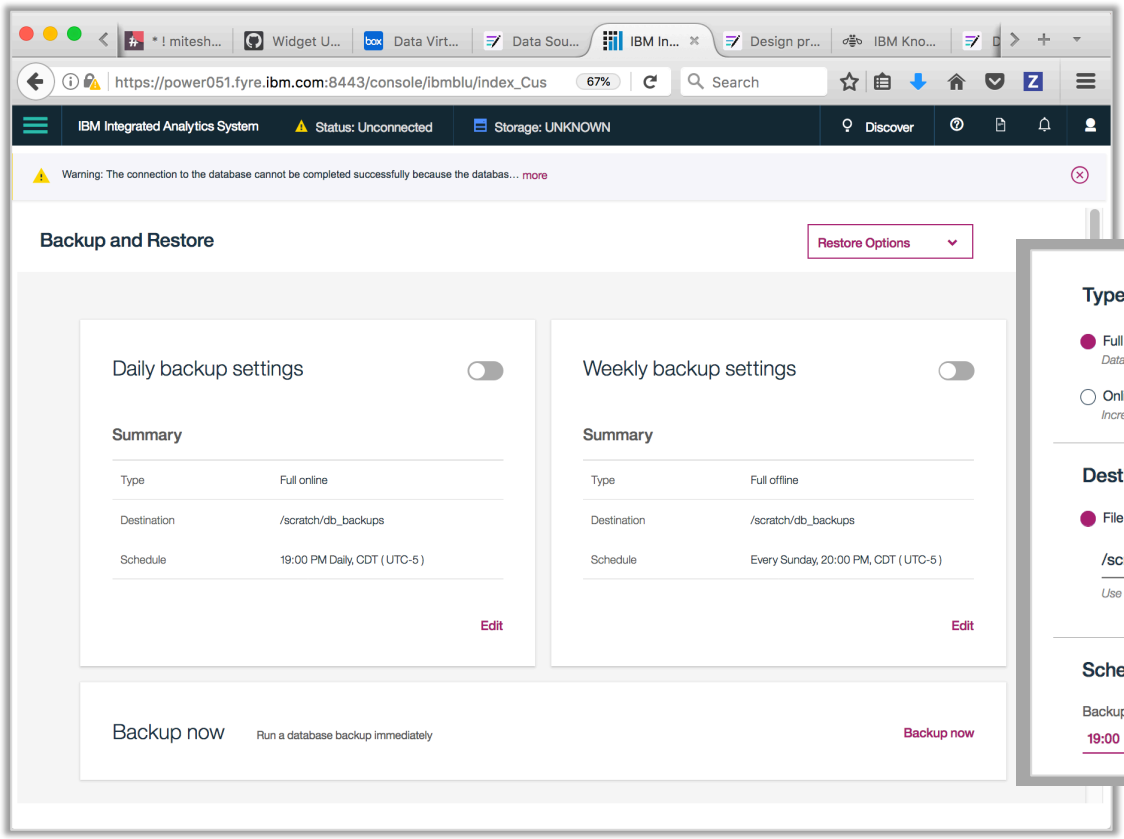

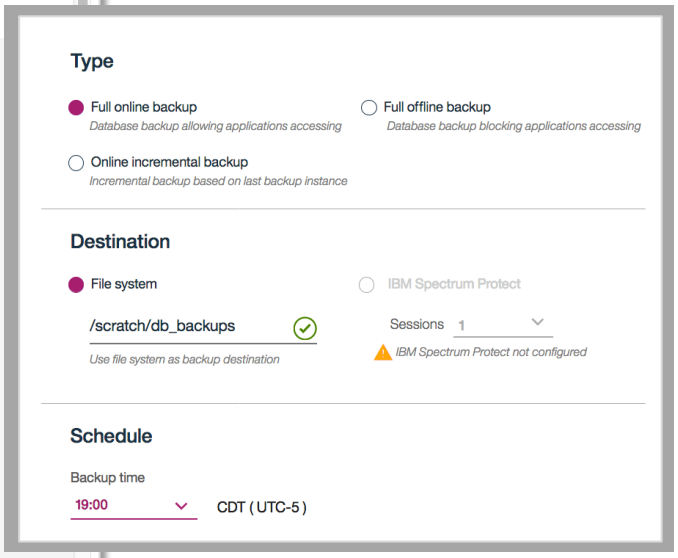

Ability to schedule daily and weekly backups.

Options

- db_backup
- db_restore
- Web console

Backup levels

- full Online backup
- full Offline backup
- Online incremental backup

Location

- file system
- Spectrum storage

IBM **Spectrum** Storage

Disaster Recovery – GPFS AFM – Coming Soon on IIAS

Working with failures

- Write to logs
- Issue Alerts when a failure of any kind occurs
- Mouse-over capability to determine status and next steps

Replication failed – Connection OK

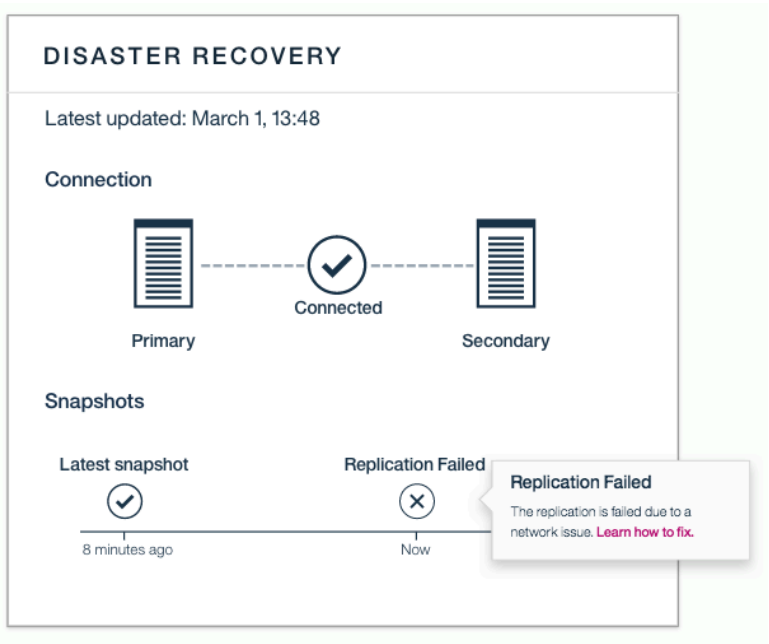

Connection is broken

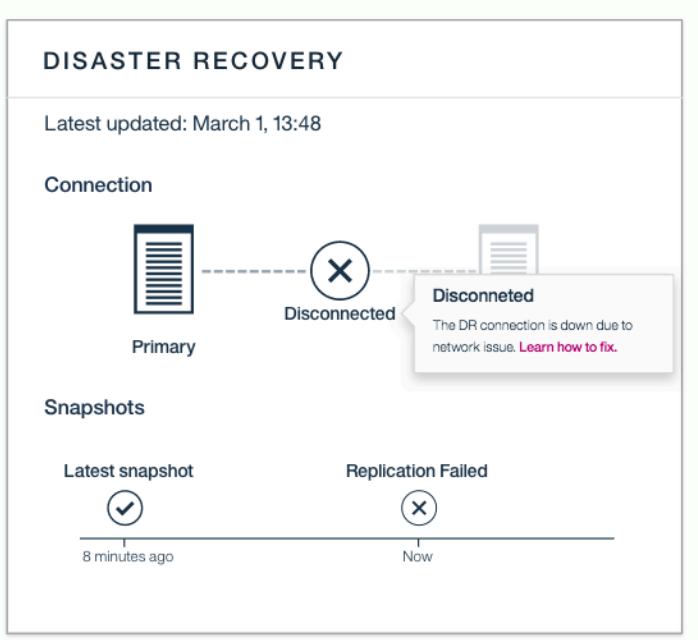

Monitoring your database

Metric Collection

- **What metrics are monitored?**
  - Connections, Statements, Locking, Throughput, Memory, I/O and Storage

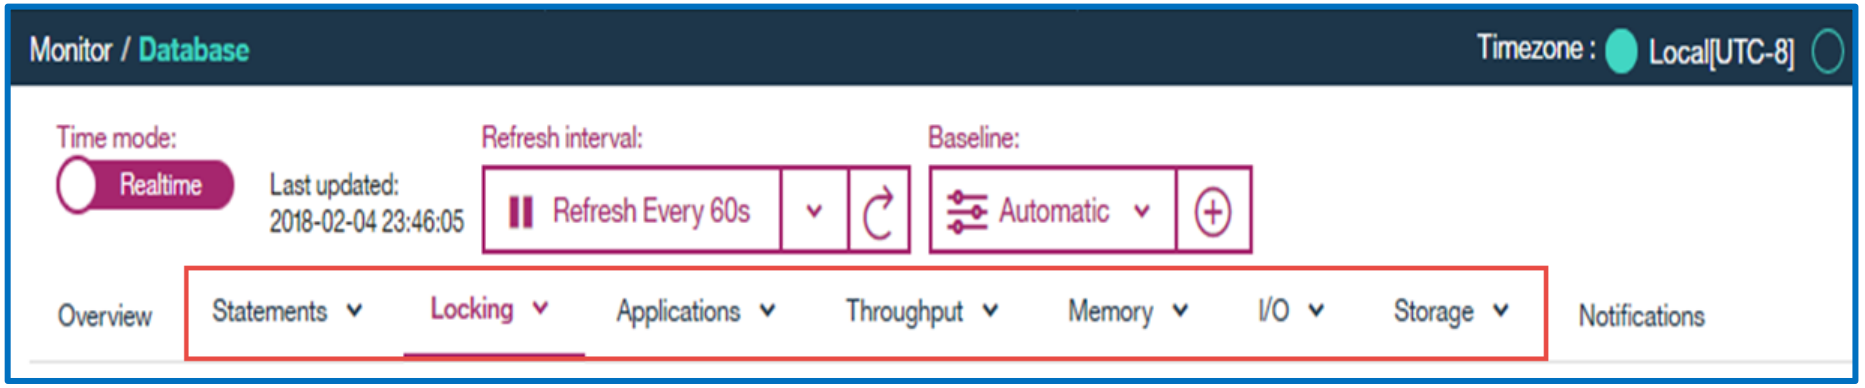

- **How are metrics obtained?**
  - MON_GET_XXXX table functions
  - Event monitors

Normal range and Automatic Baseline

- DSM automatically collects and calculates a normal range for specific metrics or key performance indicators (KPIs)

- For each collected baselined metric, DSM calculates the mean and standard deviation (about 68%) for a 4-hour block of collected data
  - Week broken down to 4-hour blocks (6/day x 7 days = 42 blocks)
    - Examples: Sun 12 am – 4 am, Mon 8 am – 12 pm, etc.
  - Same 4-hour block is aggregated over prior 5 weeks to smooth out anomalies like Black Friday

- What is normal for a given day and time is defined by the baseline for the corresponding 4-hour block
  - For example, if the sampling time is Mon 9:32 am; the Mon 8 am - 12 pm baseline applies
  - This provides a different “normal” for peak hours vs. off hours at 4-hour granularity

Normal range and Automatic Baseline

- "Normal" is when the current value is within this range of values.
- DSM also tracks the maximum and minimum values within the normal range.
- Automatic baselines are dynamic. There is no pruning involved. DSM uses "circular buffering" to keep the latest 5 weeks for comparisons.
- Automatic Baselines used for
  - Buffer Pool Smart Alert – detect when BP hit ratio is abnormally low
  - Statement Performance – detect when statement's current performance metric is outside the baseline
  - Rate of table lock events – detect when

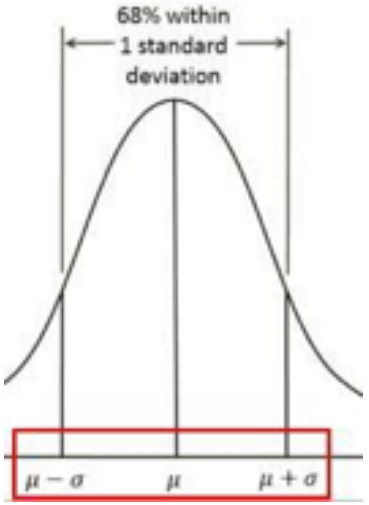

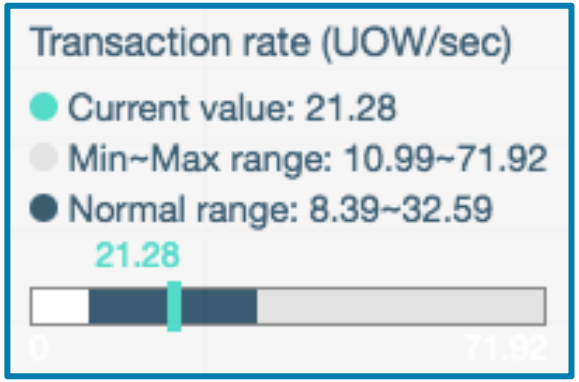

User-defined Baselines

- Enables the user to browse existing Automatic Baselines and save copies of their favorites
- Multiple saved baselines supported
- User names each saved baseline
- Saved baselines are only deleted by user, not by pruning
- User can specify a saved baseline to show as “normal” and used for Smart Alerts

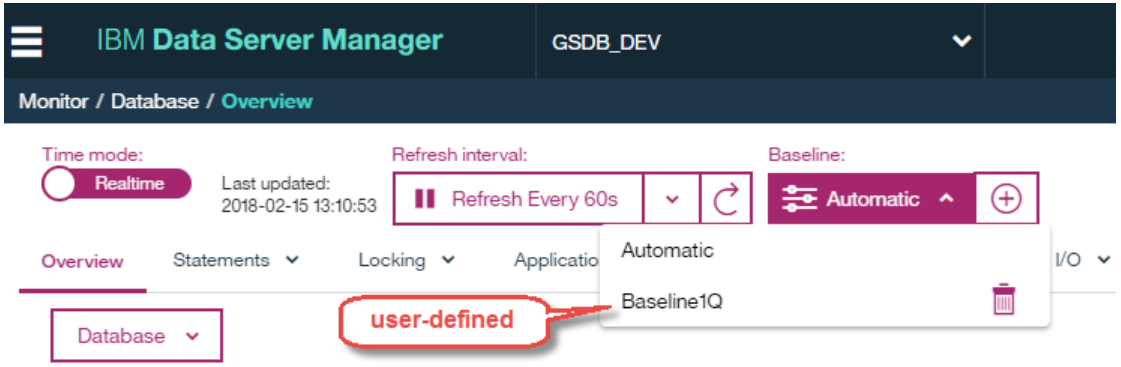

Alert Types

- Threshold alerts – values for warning and critical thresholds
- Availability alerts – resource unavailable or about to reach a threshold of non-availability (e.g. log full)
- Smart Alerts

Rate of table lock events increases	Require enterprise license and repository …	☐
Buffer pool hit ratio degrades	Require enterprise license and repository …	☐
Access plan changes	Require enterprise license and repository …	☐
Poorly performing queries	Require enterprise license and repository …	☐
Configuration changes are detect…	Require enterprise license and repository …	☐
Opportunities are found to reclai…	Require enterprise license and repository …	☐

  - If any of the following metric deviate historical value(10 times for warning alert, 5 times for information alert), a query performance smart alert will be fired.
    - STMT_EXEC_TIME
    - TOTAL_CPU_TIME
    - ROWS_READ
    - ROWS_RETURNED
    - TOTAL_ACT_WAIT_TIME
    - LOCK_ESCALS
    - SORT_OVERFLOWS
    - LOGICAL_READS
    - PHYSICAL_READS
    - TEMP_READS

Statement History

Statement monitoring

- Real-time mode

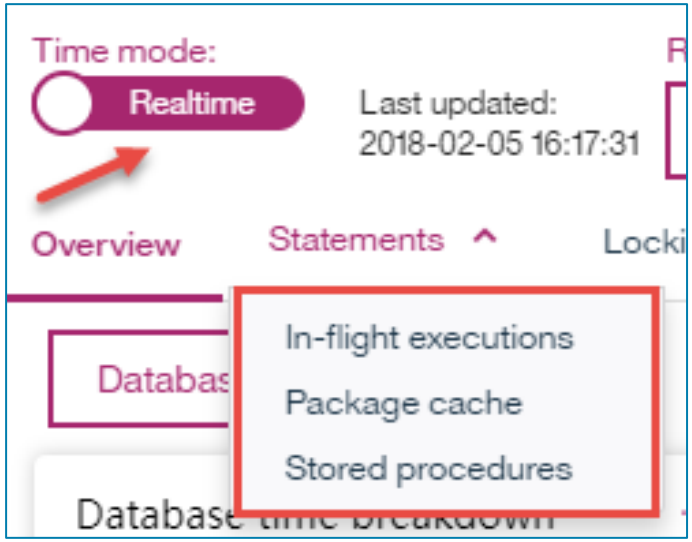

- Historical mode

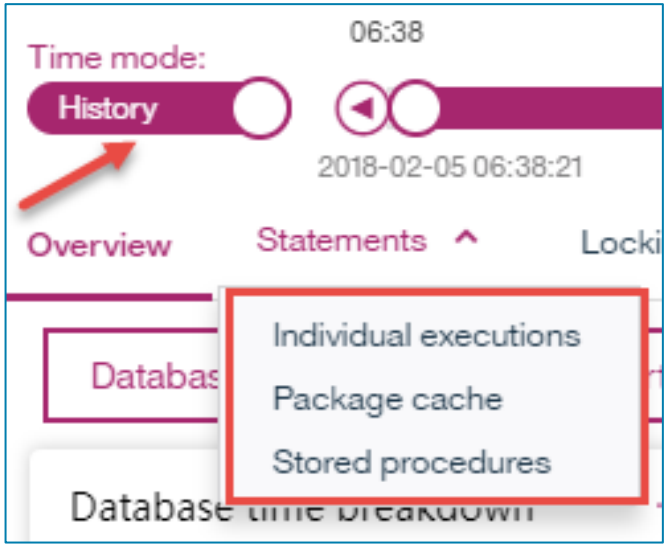

- At each collection, DSM will take a copy of all the information for the most active SQL statements in the package cache. By default DSM collects **the top 20 statements** for each metric.

Statements (real time mode)

- **In-flight executions**
  - Shows “currently” executing statements
- **Package Cache**
  - Shows the contents of the Package Cache as of the last in-memory collection.
  - Monitoring data is not accumulated in the Package Cache until after the statement completes
  - Note that there is no simple rule to determine how long a statement will remain “cached”. A statement can be evicted almost immediately after being executed in which case it will not be captured by DSM; or it can linger for days in the package cache.
- **Stored Procedures**
  - Requires database configuration parameter MON_RTN_DATA to be set to BASE
  - Monitor statements within a stored procedure

Considerations for processing the package cache statements

- **Top N Entries**
  - Only the top N entries from the package cache are displayed. Which 150 entries depends on how you have sorted the data in the grid. For example, if the grid is ordered by CPU time, then the 150 queries that consume the most CPU per execution will be displayed.
- **Averages versus Totals**
  - The statistics maintained for each statement by the package cache are accumulated as totals. Generally, it is more useful to divide the totals by the number of executions to get averages. By default, the grid shows averages. For completeness, you can toggle it to show the totals instead.
- **Column displays**
  - Can customize what metrics are displayed
  - Can sort by a specific metric.
  - Default Sort Column: Average CPU time

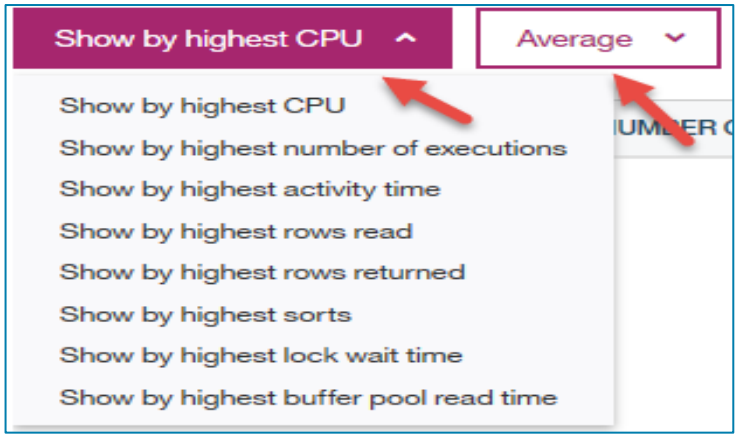

Statements > Package Cache (history mode)

- The big difference between real-time mode and historical mode
  - Real time mode directly retrieve data from the package cache, and does not calculate deltas. If a SQL statement is cached for a long time and execute thousands of times, DSM will show thousands of execution.
  - For historical mode, the delta is calculated between the metric values of the current sampling and the sampling before it. In the above case, if the SQL statement was not executed recently, it will not show in the recent time range.

Statements (History mode)

- **Individual Executions**
  - Show statements obtained from the ACTIVITY and ACTIVITYSTMT event monitors. Need to setup event monitoring
- **Package Cache***
  - Contains the “most interesting” SQL statements obtained from the package cache in the duration or time period specified
- **Stored Procedures**
  - Not available in history mode

* DSM uses the “top N” setting in the monitoring profile to determine how many “interesting” statements to capture. A SQL statement is deemed interesting if the statement is:

- Executing a large number of times
- Running a long time
- Consuming large amounts of resources (CPU, IO, memory)
- Waiting a long time
- Reading a lot of data
- Returning a lot of rows
- Escalating a lot of locks

How to enable statement history event monitoring (DSM 2.1.4+)

1. Create tablespaces and alter WLM service classes in *Monitor >Statement History Prerequisites >New Table Space and Enable WLM*

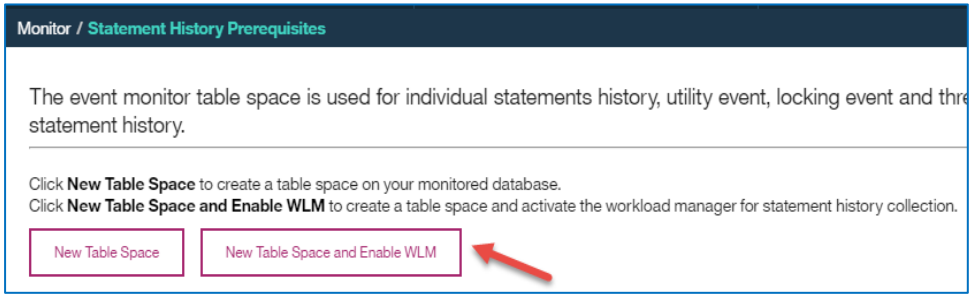

2. Edit and Run the generated SQL Script

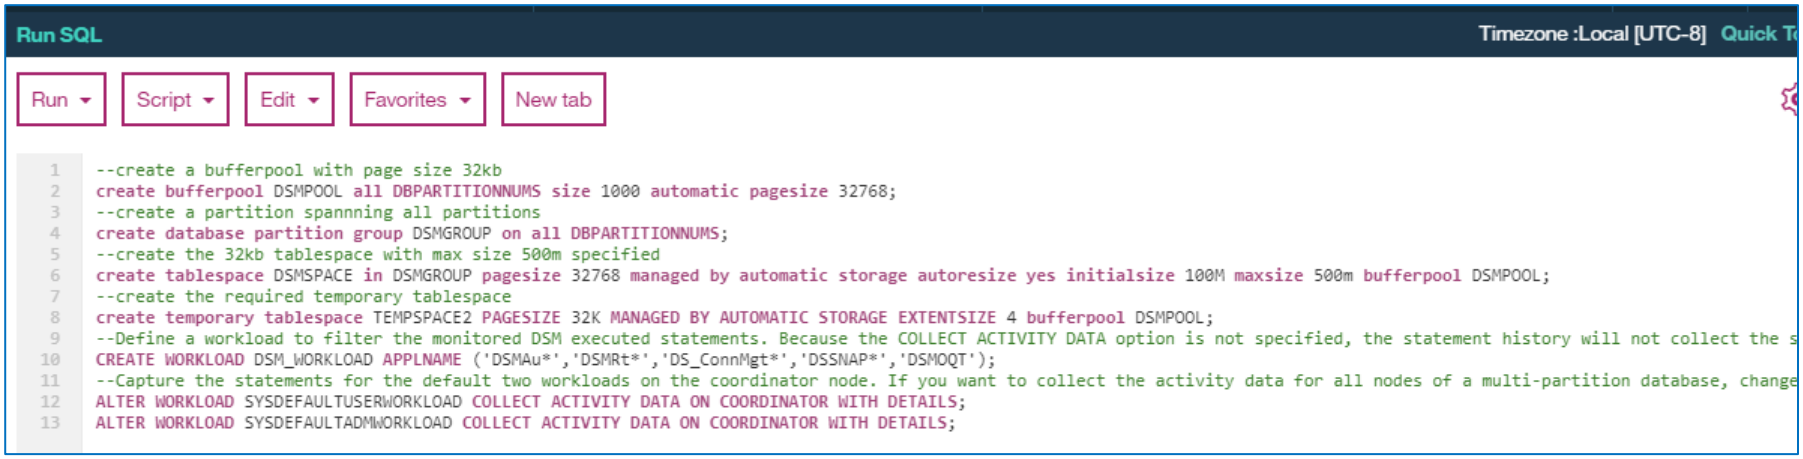

How to enable statement history event monitoring (DSM 2.1.4+)

3. Enable *Individual statements data* in *Monitoring Profiles >Persistence >SQL Statement execution data*

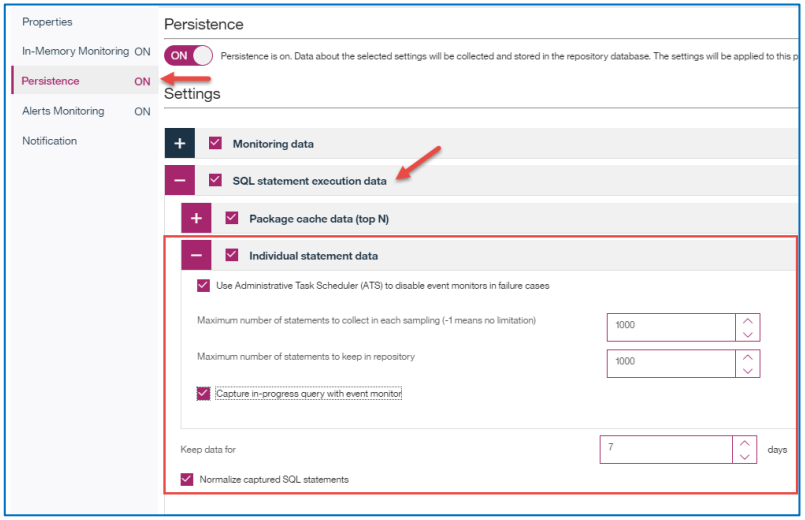

4. Assign tablespace in *Monitor >Statement History Prerequisites.* Select tablespace created in step 2.

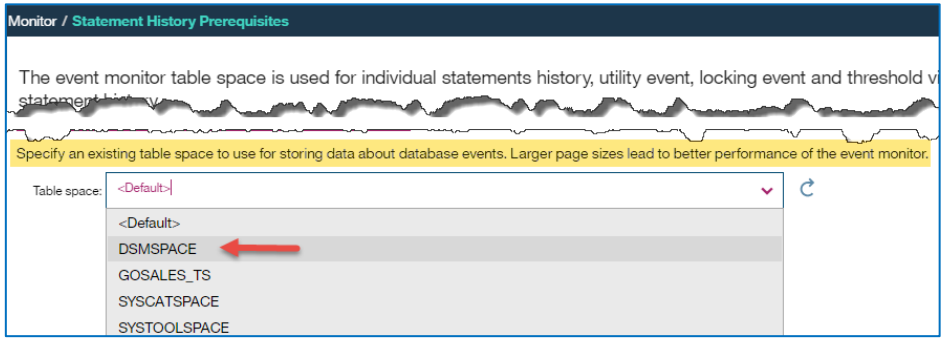

A word about SQL Capture limitations in DSM

- SQL Statement text captured is limited to 16Kb.
- While DSM allows you to set “Maximum number of statements to collect in each sampling” to -1 (no limitation), in actuality if one sampling of SQL statements is > 5K, then the event monitor will be temporarily turned off. Thus SQL statements that would run during this off period not be captured or collected. This is a “melt down” mechanism to prevent DSM from hanging.
- If a table for the event monitor is being created in an automatic storage (non-temporary) or DMS table space, the PCTDEACTIVATE parameter specifies how full the table space must be before the event monitor automatically deactivates. The specified value, which represents a percentage, can range from 0 to 100, where 100 means that the event monitor deactivates when the table space becomes completely full. The default value assumed is 100 if PCTDEACTIVATE is not specified. This option is ignored for SMS table spaces.

Threshold violations and locking event monitor

- **Event monitors are also required for monitoring and alerting**
  - Locking evemts
  - Threshold Violations
- **Set under Monitoring Profiles >Persistence**
- **Enable ATS in Db2 command line**

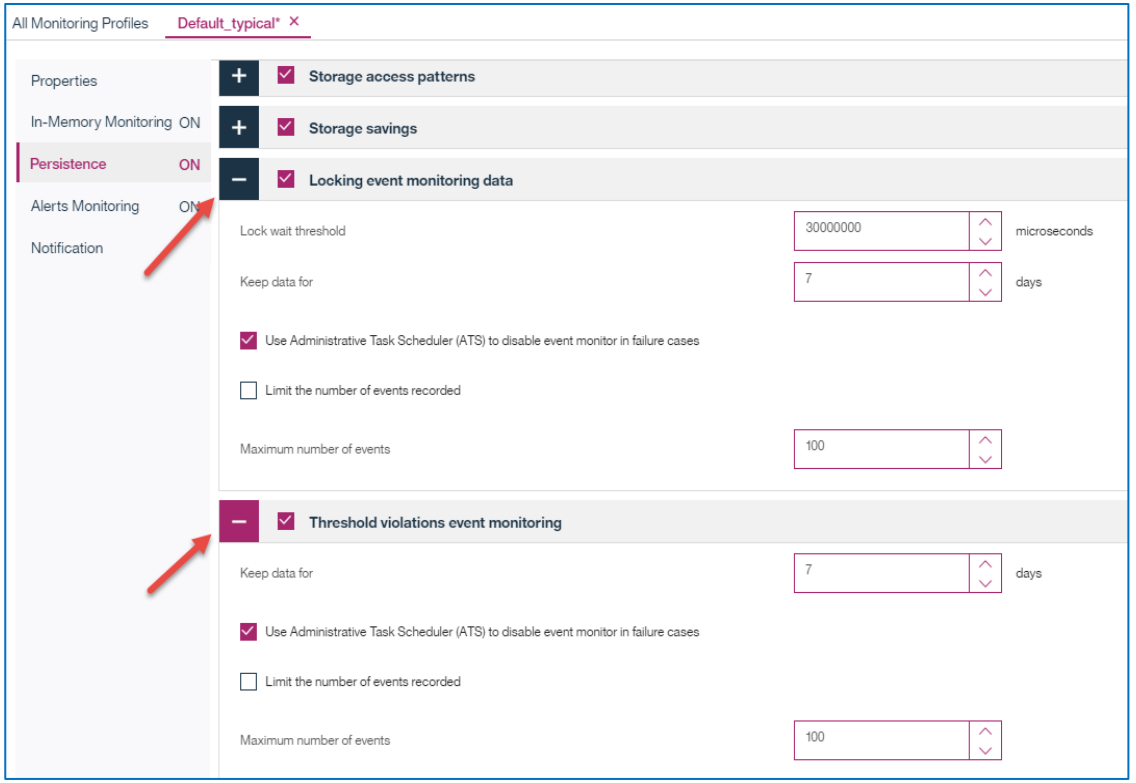

Using Workload Manager for correlating statements to applications

You can find the SQL Statements belonging to an application that is experiencing a problem, e.g. locking, by creating a workload and correlating a workload to the application

- **SETUP your DSM for statement history and WLM workloads**
  - In Monitoring Profile >Persistence >SQL Statements, make sure Individual Statements is checked
  - Setup the event monitor tablespace as shown in the previous slide.
  - Go back to Setup for Statement History. Refresh the tablespace list and select the recently created tablespace (e.g. DSMSPACE). This will assign this tablespace when creating event monitor tables.
- **Correlate workload to statements**
  - In History mode, go to Monitor >Database >Statements tab
  - View Individual Statements. Look for the workload name associated with the application that you suspect is causing performance issues.

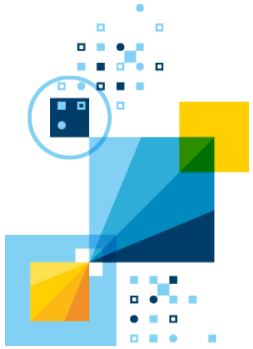

Tools Portfolio

IBM Db2 Advanced Recovery Feature

Save time, reduce errors and meet SLA's

- **What is the problem**? Customers have constant backup and recovery challenges which translates to an average $1.5M in unplanned database outages

- **What is our solution?** Db2 Advanced Recovery Feature gives you additional features on top of Db2 to provide more granular and faster recovery and quickly restore or correct erroneous data

- **Why Db2 Advanced Recovery Feature?**
  Reduce risk during recovery, reduce downtime costs, reduce DBA workload, reduce CPU consumption

- **What's included?**
  - ✓ Db2 Recovery Expert:
    - Recover faster with greater granularity while protecting your critical business data
    - Eliminate data errors before they compound into costly business mistakes
  - ✓ Db2 Merge Backup:
    - Improve speed of your backup & recovery processes
    - Minimize application impact
  - ✓ Optim High Performance Unload for DB2
    - Extract large amounts of data quickly and with minimal impact on system productivity
    - Perform full data and system migrations from one DB2 instance to another

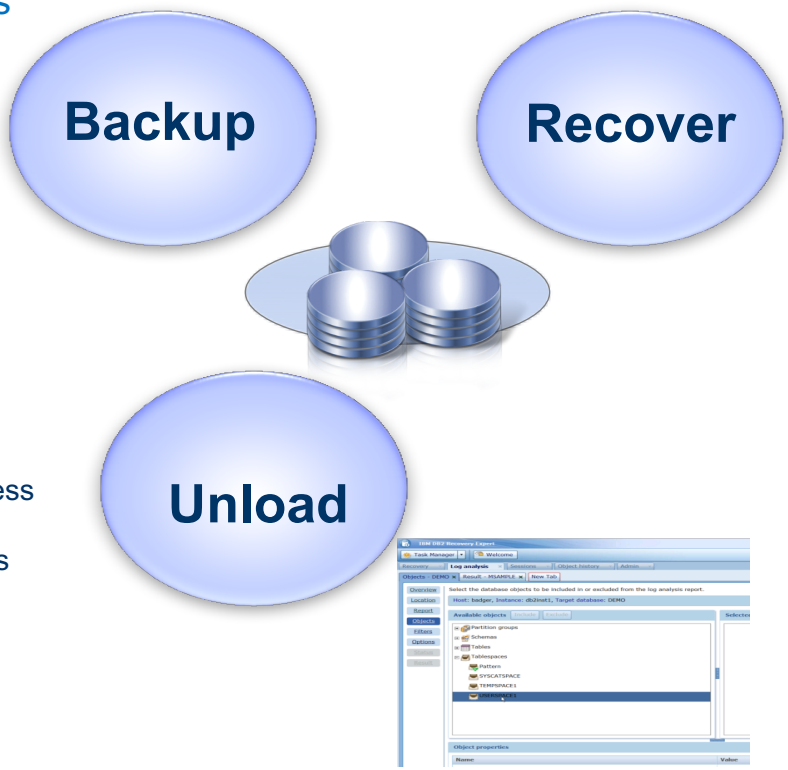

IBM DB2 Advanced Recovery Feature

Save time, reduce errors, and meet SLAs

DB2 Merge Backup

- Improve speed of your backup & recovery processes
- Minimize application impact

DB2 Recovery Expert

- Recover faster with greater granularity while protecting your critical business data
- Eliminate data errors before they compound into costly business mistakes
- Track and report data changes in response to auditing requests

Optim High Performance Unload

- Extract large amounts of data quickly and with minimal impact on system productivity
- Perform full data and system migrations from one DB2 instance to another

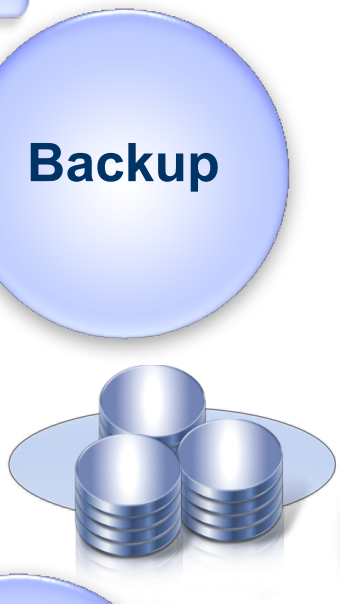

DB2 Connect and Data Server Driver To Access and Connect to Data

- **What is the problem**? Web, mobile, analytics requirements driving workload from distributed and cloud apps to IBM DB managed data
  - ✓ DB2 z/OS and DB2 i usage increases in large % of customers

- **What is our solution?** Data Server Driver provides a common framework supporting connectivity to DB2, dashDB and BigInsights BigSQL data

- **What extends our solution to z/OS and IBM i platforms?** A suite of DB2 Connect Editions that match to user, application and unlimited usage models
  - ✓ Enterprise Edition: User licensing model
  - ✓ App Server Editions: Server licensing model
  - ✓ Unlimited Editions: System z and IBM I licensing model
  - ✓ Advanced Editions: Value add editions improving deployability, performance, and management

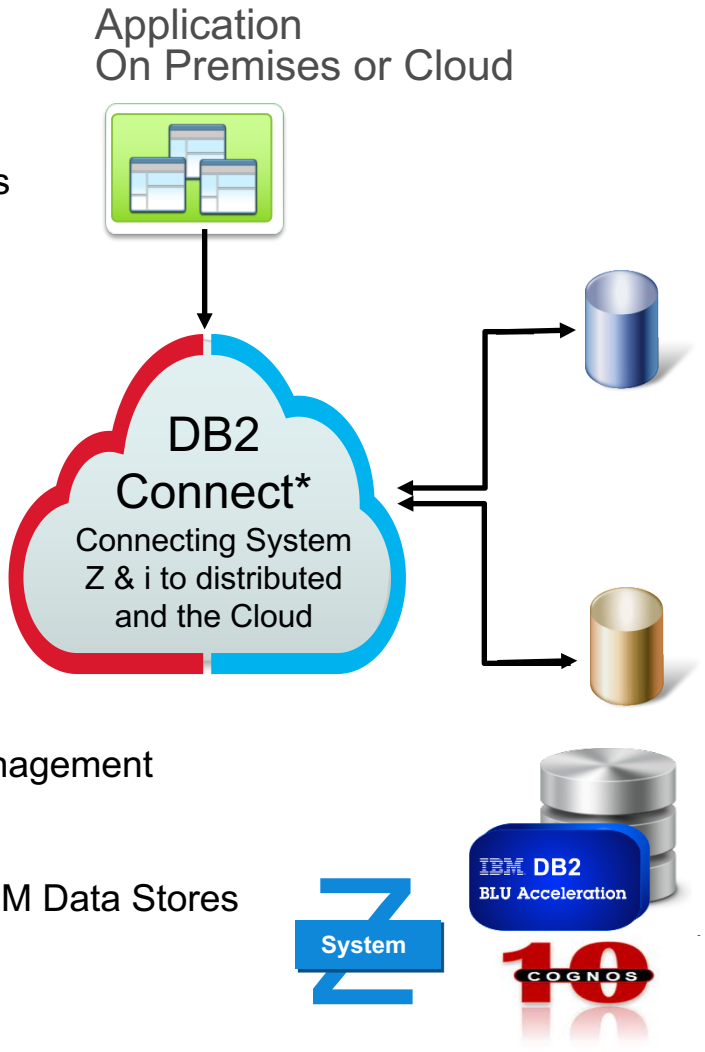

- **What's recent?** ***New with V11***
  - ✓ Common drivers simplify management of access between IBM Data Stores
  - ✓ New platforms such as iMAC driver support
  - ✓ Improved manageability
  - ✓ High availability and performance enhancements
  - ✓ Monthly licenses available

* DS Driver supports IBM distributed database access

IBM Data Studio - For Development

An integrated, modular environment for database development and administration of IBM DB2 for Linux, UNIX and Windows

- Also supports IBM DB2 for z/OS, IBM DB2 for i, IBM Informix and Big SQL
- Enables developers and administrators to create and manage heterogeneous database environments for increased productivity.
- Streamline database development with advanced query validation, object management, procedure development, deploy and debug.
- Improved collaboration through the open source integrated Eclipse environment
- Support for popular source code management and team collaboration tools like Git.

- Available at no charge.
- Support entitlement provided through DB2

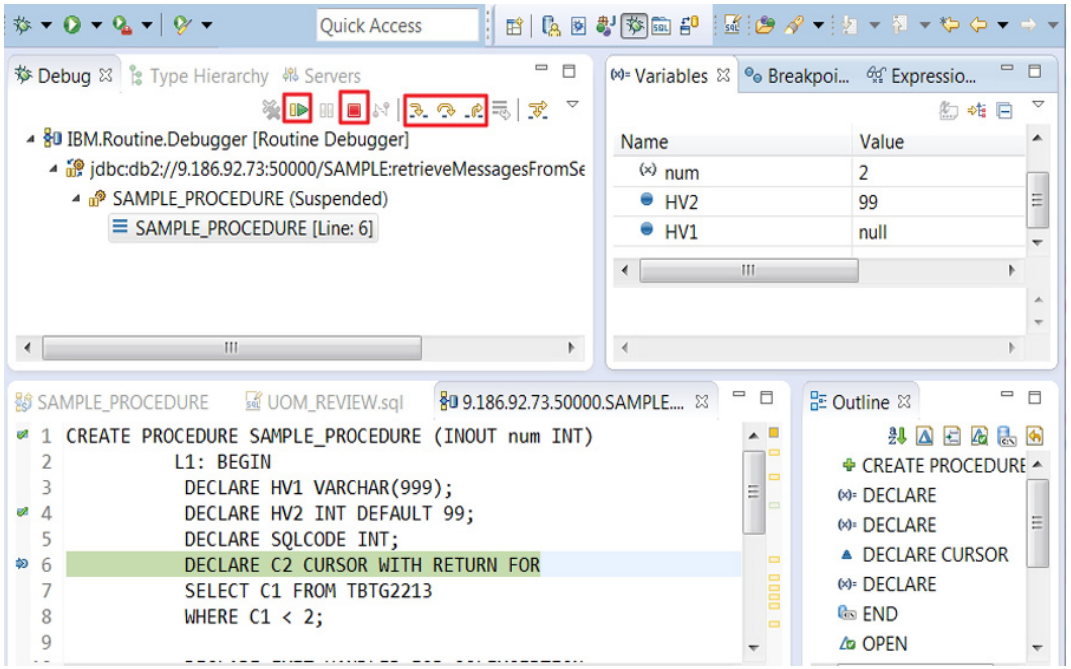

InfoSphere Data Architect – For Data Modeling

A collaborative, data design solution to discover, model, relate, and standardize diverse data assets

- Design and manage enterprise data models
- Enforce model conformance to enterprise standards
- Leverage industry data models for best practices
- Optimize existing investments

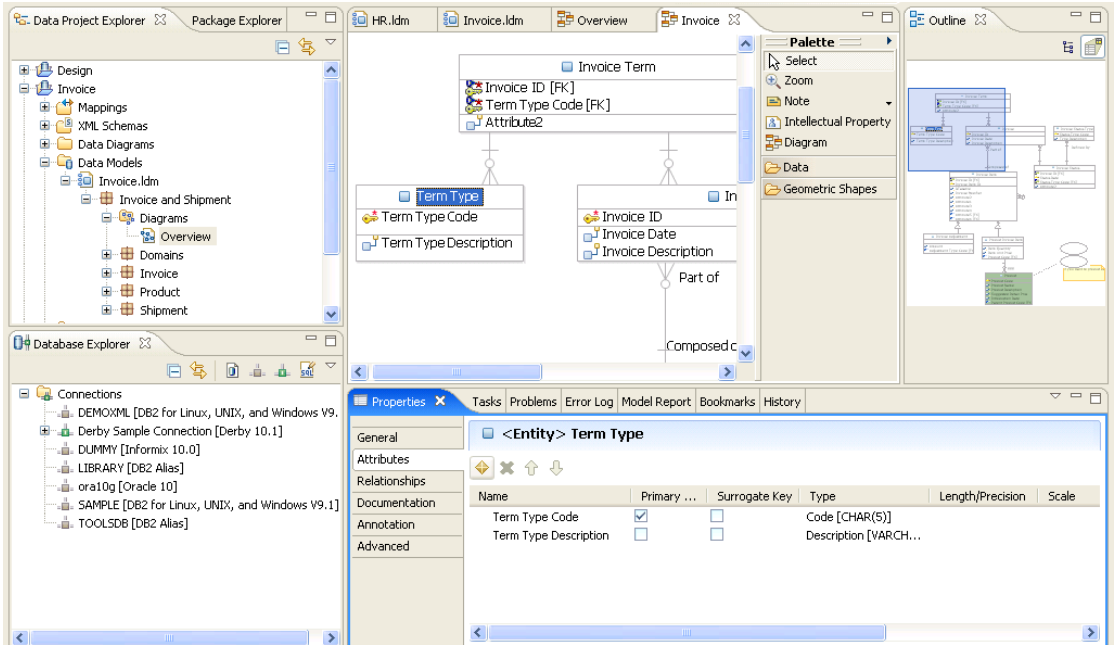

- Design portable Data Models
  - Logical
  - Physical
  - Dimensional
  - Glossary
  - Mapping
- Heterogeneous Database support
- Reverse Engineering
- Visualize Design
- Model Validation
- Naming Standards
- Team collaboration

Data Server Gateway for Odata

Remote Mobile and Open access to IBM Data Stores

- **What is the problem**? Mobile (Android, Windows, iOS) and other open application models need access to IBM DB environments

- **What is our solution?** Data Server Gateway for Odata supports open standards based connectivity to IBM Data Server environments

- **Why Odata?**
  Odata (Open Data Protocol) is an ISO/IEC approved, Oasis standard that defines a set of best practices for building and consuming Restful API's

- **What's included?**
  - ✓ Publish Odata V4 Rest API endpoints/services for selected database tables
  - ✓ Support for most data types
  - ✓ Support for Crud operations
    - ✓Get, Post, Put,(update a full row) Patch (update selected columns) , Delete
  - ✓ JSON and XML payloads
  - ✓ API persistence
  - ✓ Authentication and Authorization
  - ✓ SSL Support

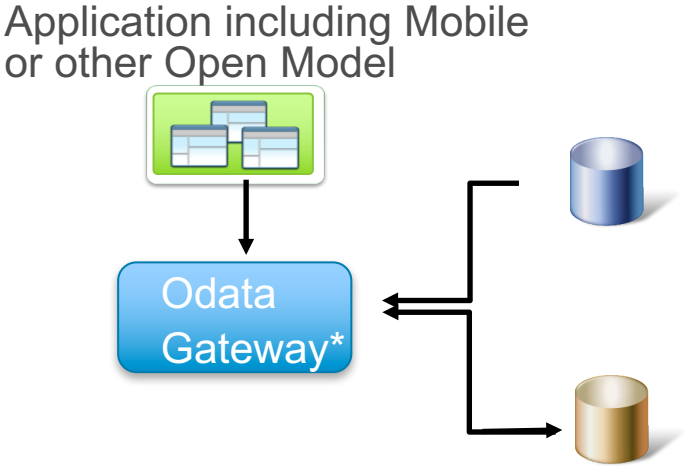

* DS Driver supports Odata Gateway access to IBM Data Stores

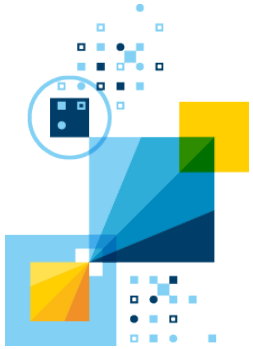

Backup

Graphical Work Load Management Tool

View and manipulate DB2 workload management configurations

- **WLM Expert in a Box**
- **Provided As Is**
- **https://www.idug.org/p/bl/ar/blogaid=430**

How do I buy IBM Data Server Manager

- **IBM Data Server Manager is available for free download from IBM developerWorks**
  - http://www.ibm.com/developerworks/downloads/im/dsm/
- **Base capabilities are available at no extra cost for all DB2 Editions**
- **Enterprise capabilities are available if you:**
  - Purchase IBM DB2 Advanced Enterprise Server Edition or IBM DB2 Advanced Workgroup Server Edition
  - Purchase the Performance Management Offering for DB2 Workgroup or Enterprise

Data Server Manager is DB2 Edition Aware

Settings

- No need to apply license activation kits to turn on DSM Enterprise Features
- Data Server Manager now detects the DB2 for Linux, UNIX or Windows Edition and enables capabilities automatically. Details in Data Connections.

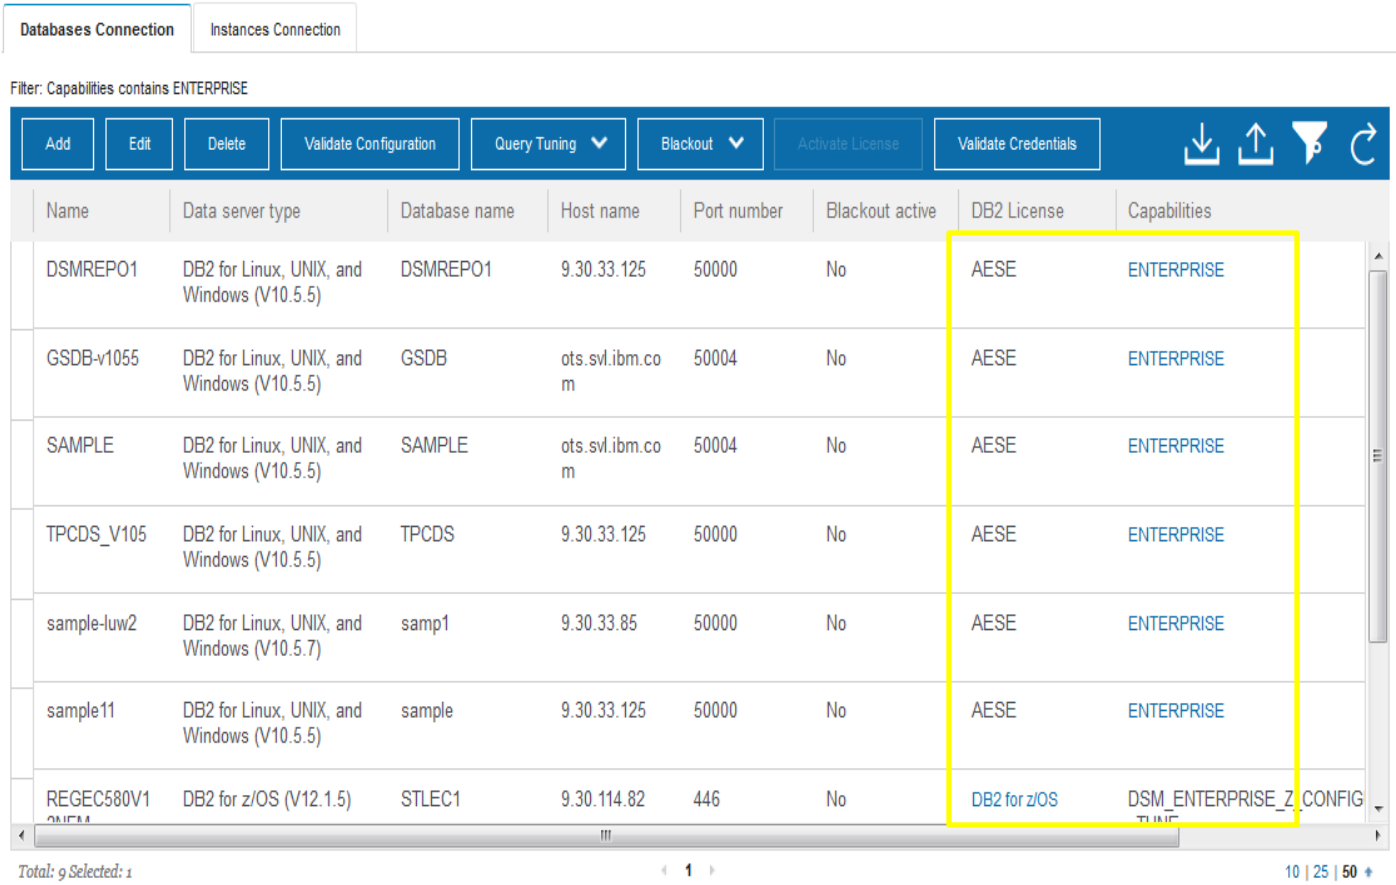

License activation still required when using DB2 for LUW Express, Workgroup and Enterprise servers earlier than 10.5.0.5 that have the Performance Management Offering

IBM Data Server Manager Enterprise, V2.1

Base Edition (BE) and Enterprise Edition (EE) reference and comparison

Feature	Base	Enterprise
Administer databases		
Discover and explore databases, instances and pureScale nodes	Yes	Yes
Edit and review instance and database configuration	Yes	Yes
Manage databases and create, alter, and drop objects	Yes	Yes
Edit and run commands, run external scripts and DB2 utilities	Yes	Yes
Schedule Scripts	No	Yes
Federated support (alter / create federated objects)	Yes	Yes
Manage Roles and Privileges	Yes	Yes
Configure DB2 native encryption	Yes	Yes
Monitor performance and health		
Monitor database health	Yes	Yes
Monitor real-time performance	Yes	Yes
Monitor historical performance	No	Yes
Monitor end-to-end performance	No	Yes
Monitor stored procedures	No	Yes
Configure threshold and user-defined alerts	Yes	Yes
Utilize smart alerts	No	Yes
Configure query plan comparison alerts	No	Yes
Monitor storage access patterns and usage	No	Yes
Monitor DPF environments	Yes	Yes
Monitor pureScale environments	Yes	Yes
Monitor HADR environments	Yes	Yes
Force connections and cancel problematic SQL statements	Yes	Yes
Define blackout periods	Yes	Yes
Generate monitoring reports and share via email	No	Yes

IBM Data Server Manager, V2.1

Base Edition (BE) and Enterprise Edition (EE) reference and comparison

Feature	Base	Enterprise
Manage database manager (DBM) and database (DB) configurations		
Track configuration changes	No	Yes
Compare configurations	No	Yes
Synchronize configurations	No	Yes
Manage performance		
Visualize access plan (EXPLAIN)	Yes	Yes
Tune single SQL using advisors	No	Yes
Tune entire workloads	No	Yes
What-if analysis	No	Yes
Monitor WLM and manage event monitor life-cycle	No	Yes
Identify storage savings	No	Yes
Leverage guided problem determination	No	Yes
Manage database clients		
Track client configuration changes	No	Yes
Compare client configurations	No	Yes
Override client configuration and connection options	No	Yes

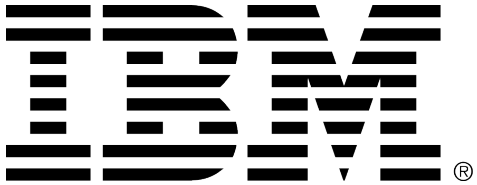

Legal Disclaimer

- © 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.
- If the text contains performance statistics or references to benchmarks, insert the following language; otherwise delete:
  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.
- If the text includes any customer examples, please confirm we have prior written approval from such customer and insert the following language; otherwise delete:
  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.
- Please review text for proper trademark attribution of IBM products. At first use, each product name must be the full name and include appropriate trademark symbols (e.g., IBM Lotus® Sametime® Unyte™). Subsequent references can drop "IBM" but should include the proper branding (e.g., Lotus Sametime Gateway, or WebSphere Application Server). Please refer to http://www.ibm.com/legal/copytrade.shtml for guidance on which trademarks require the ® or ™ symbol. Do not use abbreviations for IBM product names in your presentation. All product names must be used as adjectives rather than nouns. Please list all of the trademarks that you use in your presentation as follows; delete any not included in your presentation. IBM, the IBM logo, Lotus, Lotus Notes, Notes, Domino, Quickr, Sametime, WebSphere, UC2, PartnerWorld and Lotusphere are trademarks of International Business Machines Corporation in the United States, other countries, or both. Unyte is a trademark of WebDialogs, Inc., in the United States, other countries, or both.
- If you reference Adobe® in the text, please mark the first use and include the following; otherwise delete:
  Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.
- If you reference Java™ in the text, please mark the first use and include the following; otherwise delete:
  Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.
- If you reference Microsoft® and/or Windows® in the text, please mark the first use and include the following, as applicable; otherwise delete:
  Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.
- If you reference Intel® and/or any of the following Intel products in the text, please mark the first use and include those that you use as follows; otherwise delete:
  Intel, Intel Centrino, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
- If you reference UNIX® in the text, please mark the first use and include the following; otherwise delete:
  UNIX is a registered trademark of The Open Group in the United States and other countries.
- If you reference Linux® in your presentation, please mark the first use and include the following; otherwise delete:
  Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both. Other company, product, or service names may be trademarks or service marks of others.
- If the text/graphics include screenshots, no actual IBM employee names may be used (even your own), if your screenshots include fictitious company names (e.g., Renovations, Zeta Bank, Acme) please update and insert the following; otherwise delete: All references to [insert fictitious company name] refer to a fictitious company and are used for illustration purposes only.