
The Data Pioneer



Martin Hubel Consulting Inc.

Martin Hubel

+ 1 905-764-7498

martin@mhubel.com

www.mhubel.com

Copyright©2020 Martin Hubel Consulting Inc.

1

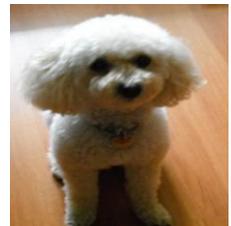
Agenda

- History
- Tasks and priorities
- Environmental setup
- Approaches to support DevOps

2

About Martin Hubel

- Worked with DB2 for z/OS since 1985
 - Started teaching Db2 classes in Sept 1987
- Started with DB2 for LUW in 1993
 - Home town advantage!
- Consulted, taught or presented on Db2 in 32 countries
- Areas of specialty
 - Performance
 - Design
 - Administration
 - Migration



3

Industry Involvement

- IBM Gold Consultant 1994-
- IBM Champion 2008-
- Host, The Db2Night Show
- Guide 1986-1995
 - Project manager, Db2 System Administration
- IDUG 1989 –
 - Given 60+ presentations
- On first board of Central Canada Db2 Users Group, Oct 1985
- On various TABs and CACs for Db2 z/OS and LUW
- DB2 LUW SAP TLE 2010-

Looking Back At Pioneers In North America

- Congratulations! Here is your new farm
 - You have to clear the land first
 - Winter is coming; build yourself a house too



A Lot of Work To Do

- No services existed at all
 - Surveyors laid out the land allotments
 - Roads were planned (maybe) but not built
 - No fire department or any infrastructure
- Pioneers had to be self-sufficient
 - You have your own tools, or what you make
- Pioneers had to solve many problems
 - Both known & unknown issues
 - Experience taught many lessons
- When you start from scratch you have to do everything yourself
 - Data pioneers are in the same position

Fast Forward to Today

- A new business or division has started
 - Whoopee! Success!
- Developers have created a new application
 - Wow, it's so cool!
 - They worked so hard and quickly!
- Users and customers are coming on board
 - How will we manage the growth?
- Demands on the database are growing
 - Let's hire a DBA

7

Startups – Clients of MHC

- Gas loyalty program, Dallas TX
 - Rewards program
- Investment analysis, Ottawa, Ontario
 - Technical analysis of stocks
- ML – AI applications, Toronto, Ontario
 - Health insurance fraud
 - Weekly grocery flyers
 - Brain injury analysis

Copyright©2020 Martin Hubel Consulting Inc.

8

The Data Pioneer

- Now, you are the DBA
 - Many demands on your time
 - Will your company be a victim of its own success?
 - Where to start?
- DevOps has a head start
 - Schedules to meet
 - Need to standardize approach to data management

9

Developer Priorities

- Functionality
 - Per user requirements
- Productivity
 - DevOps
 - Work smart & fast
- Performance
 - May require education on efficient SQL
 - May sacrifice integrity for performance

DBA Priorities

- Data integrity
 - Ensure availability
 - Protect data
- Developer support
 - Data loads
 - SQL support
- Performance
 - System performance
 - Application performance

11

Introducing Data Priorities to Developers

- The delivery clock is running
 - New functions
 - New customers with new function in contracts
- But the environment has exposures within the database
- How can you introduce data integrity with a minimal impact on development schedules?
 - Selling data integrity to developers

12

Selling Integrity to Management & Developers

- Level of availability required?
 - Cost per minute of outage for web application?
- How large is the batch window?
 - How long does the biggest job run?
- For customer data:
 - Can you afford to lose a customer?
 - What is the cost if data is unrecoverable?
 - What steps should you take to ensure customer data is recoverable?

13

Selling Recovery to Management

- Maximum length for an outage before severe financial impact?
 - What is the maximum acceptable recovery time?
- Cost and risk analysis:
 - Company business lost
 - Production capacity lost
 - Personnel downtime
 - Value of lost data

14

Integrity

- Application and data integrity is essential
 - Components on next slide
- Users do not often concern themselves with data integrity
 - They assume it's there – until it's not!
 - Developers need to work with DBA to ensure application and data integrity
 - Many Db2 facilities help enforce data integrity
 - Still require proper use

15

Integrity Components

- Data Validation
 - Checking of input data to ensure it is within range ("correctness")
- Security
 - Prevent unauthorized access to data
 - Possible control by application function
- Recovery
 - The ability to recover from a system or application failure
 - Often includes point-in-time recovery
 - Off site recovery in case data center becomes unavailable
- Operability
 - Procedures for the routine operation and maintenance follow installation standards
 - Complex procedures lead to higher operating costs and errors
- Concurrency
 - Ensure that no updates are lost during concurrent updates

16

DBA Areas of Interest

- This list is in priority order for day 1 if you were the first DBA involved:
 1. Backup and recovery
 2. Naming standards
 3. Security
 4. Data movement
 5. Support of developer practices
 6. Data quality
 7. Design of future databases
 8. Performance

17

One-time Activities (1)*

- Take a backup if there are none
 - Review ability to recover
- Develop naming standards
 - Suggestions are available
- Review or develop security framework
 - Encrypt data at rest
 - Data on the network
 - Work with security department
- * - Adjusted as necessary

18

One-time Activities (2)*

- * Adjusted as necessary
- Develop scripts
 - Db2 can automate some utilities, or write scripts:
 - Back up
 - RUNSTATS – use profiles – set in DB cfg
 - Cleanup – old backups & archive logs – can set in DB cfg
 - Performance reports
 - Data loads / migration
 - Restore / redirected restore
 - Schedule scripts via cron or Db2's ATS (Administrative Task Scheduler)

19

1. Key Activities – Backup and Recovery

- Your backup strategy will determine your recovery options
 - Ensure you have backups that work
- Recovery may be required occasionally
 - Hardware failure
 - Recover to as close as possible to the point of failure
 - Application or software failure
 - Often to a chosen point in time before the failure
 - Disaster recovery
 - Offsite
- Test your recovery scenarios prior needing them
 - Failure to recover in production never acceptable

20

2. Key Activities – Naming Standards

- Standards simplify:
 - Understanding and communication
 - Automation
- DB names are still limited to 8 characters. Possible things to identify:
 - Customer
 - Application
 - Country
 - Server
 - Version, e.g. Production, QA, dev, test, etc.

21

3. Key Activities – Security

- Setup authentication
- Separate environments
 - Production from development and test
 - Customers from each other
- Encrypt data at rest
 - Db2 11.5 makes this an easy default choice
- Encrypt data movement

22

4. Data Movement

- At startup, developers may have used a single database
 - Many customers
 - Development and production
 - Testing
- Migration
 - Between testing levels / production
- Security
 - Should encrypt by default – straightforward in 11.5

23

Splitting the Database

- Licensing
 - DB size up to 100 GB is free in 11.5
- Growth
- Automated builds and deployments for new deployments
 - Base builds / response files
 - Change analysis & version tracking
- Automate via scripts
 - Naming standards / patterns make this easier

24

5. Developer Support

- Understanding units of work
 - Need to identify business units of work
 - Where do commits go?
- Help with concurrency
 - Optimistic or pessimistic locking
- Help with efficient SQL and other Db2 facilities
 - Multi-row inserts and fetch

25

Dealing with Developers – Old School

The DBA Guide

Survival Tips For Developers: How To Talk To Your DBA

Dangerous	Safer	Safest	Ultra Safe
Stop blaming the code.	We may need to review my code.	Can you help me understand this code?	Yes, you are correct.
You need to solve this one, and quick.	We may need to work together.	Can you help me get this to work?	Yes, you are correct.
The database is having issues AGAIN.	We need you to look at the database.	Can you help us find the issue?	Yes, you are correct.
The database doesn't need to be backed up.	We might want to have it backed up.	Can you help us review our recovery plan?	Yes, you are correct.
It worked fine in dev/test, just make it work in production.	We need you to tell us why it worked in dev/test.	Can you help us find what we did wrong?	Yes, you are correct.



<http://thomaslarock.com>

26

Dealing with Developers – The New World

- DevOps have a head start on DBA
 - Deadlines to meet
 - Sprints are ongoing
- DBA must introduce standards with minimum impact on productivity
 - Be part of the solution
 - Developers are likely to avoid people who create problems for them

27

6. Data quality

- Ensure recoverability
 - Take and test backups
 - Develop and test recovery scenarios
 - Include backup and log retention strategies
- Participate in or initiate off-site recovery planning
- Integrate Db2 security into corporate security plan
- Take advantage of Db2's autonomic maintenance features
 - Automatic Runstats, etc.

28

Quick Ways to Wreck Application Data Integrity

- Allow **not logged** operations
 - Talk to developers first before changing DB cfg to prevent
- **Improper** use of Commit
 - Wrongly identify units of work
 - Intermediate commits without restart capability
- **Incorrect validation** of input data
 - Should be caught during testing
- Circumvention of security
 - **SQL injection**

29

8. Performance

- Generally specified as a "given" during requirements phase of project
 - Hardware cost is not considered until it is raised by IT
- Customer satisfaction and personnel productivity are benefits of performance tuning
- Expectations and perceptions about performance must be managed
 - Performance reporting
 - Service Level Agreement (SLA)
- Performance tuning requires a good understanding of the environment

30

Performance Activities

- Data integrity must come first
- Performance tuning leads to quick wins
 - Good for everyone's morale (and your paycheck)
 - Successful tuning demonstrates technical expertise
- DB / instance performance tuning
 - BP analysis
 - Sort parameters
 - DB & DBM config parameters
- Application performance
 - Review SQL performance – match indexes to SQL
 - Add / alter indexes to match SQL
 - Identify and work to resolve concurrency issues

31

Daily DBA Activities

- Production support
 - Ensure availability
 - Check status of overnight scripts
 - Perform database maintenance functions as needed
 - Use Db2 autonomic facilities
 - Resolve and report on outstanding issues
- Development coordination
 - Provide help on SQL use
 - Create and alter object definitions
 - Migrate objects between test environments

32

Db2 DB parameters for routine maintenance

Block non logged operations	(BLOCKNONLOGGED) = NO → YES
...	
Number of database backups to retain	(NUM_DB_BACKUPS) = 12 → 7
Recovery history retention (days)	(REC_HIS_RETENTN) = 366 → 8
Auto deletion of recovery objects	(AUTO_DEL_REC_OBJ) = OFF → ON
...	
Automatic maintenance	(AUTO_MAINT) = ON
Automatic database backup	(AUTO_DB_BACKUP) = OFF
Automatic table maintenance	(AUTO_TBL_MAINT) = ON
Automatic runstats	(AUTO_RUNSTATS) = ON
Real-time statistics	(AUTO_STMT_STATS) = ON
Statistical views	(AUTO_STATS_VIEWS) = OFF
Automatic sampling	(AUTO_SAMPLING) = ON
Automatic reorganization	(AUTO_REORG) = OFF

Copyright©2020 Martin Hubel Consulting Inc.

33

Monthly DBA Activities

- * Or periodically / as needed throughout the month
- Status reports / Documentation
- Restore /recovery as needed
- Data loads on request
- Perform design reviews
- BCP / DRP Testing
- Personal Development
 - Practice – e.g. recovery
 - Training

34

Summary

- In DevOps, responsiveness is the key
 - Be part of the solution
 - Help developers to be more effective
- Integrity must be established and maintained
 - Recoverability
 - Security

The Data Pioneer

Martin Hubel

Martin Hubel Consulting Inc.

martin@mhubel.com

+1 905-764-7498

