

# An Audit a day keeps the lawyers at bay!

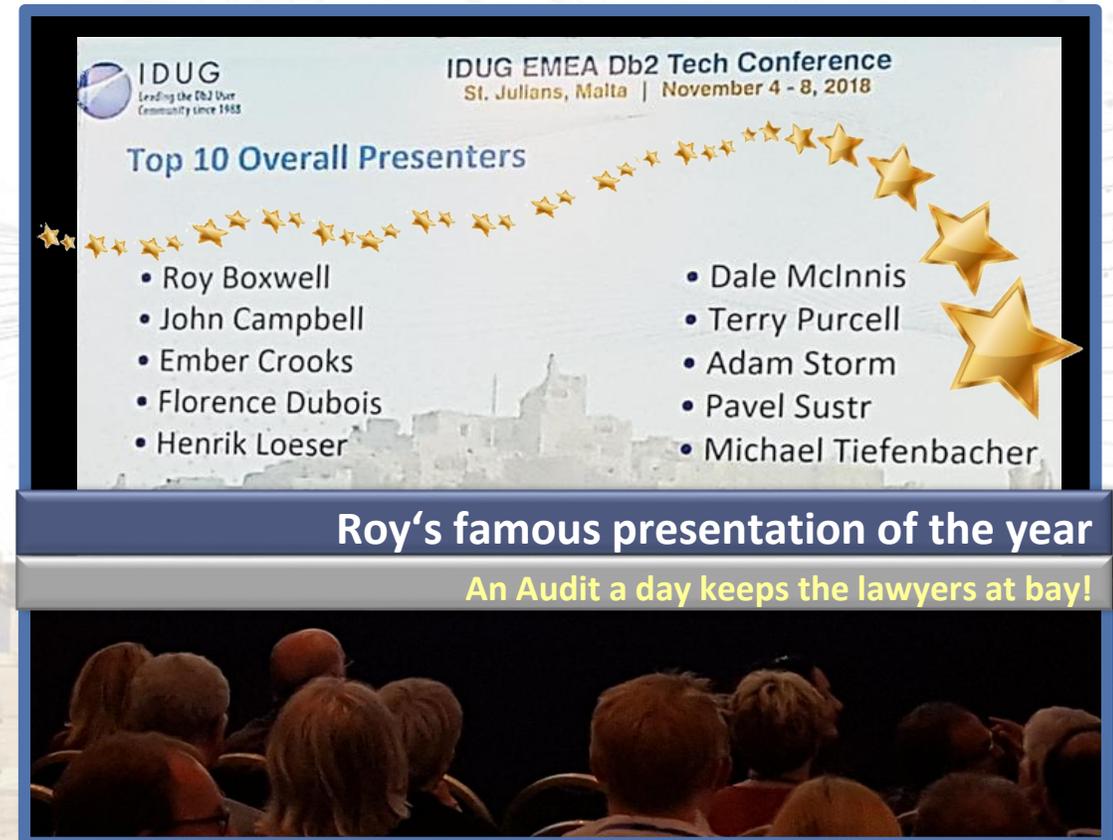
**Roy Boxwell**

**SOFTWARE ENGINEERING GmbH**

Session code: F17

08.11.2018, 09:40

Db2 for z/OS



**IDUG**  
Leading the DB2 User  
Community since 1988

**IDUG EMEA Db2 Tech Conference**  
St. Julians, Malta | November 4 - 8, 2018

**Top 10 Overall Presenters**

- Roy Boxwell
- John Campbell
- Ember Crooks
- Florence Dubois
- Henrik Loeser
- Dale McInnis
- Terry Purcell
- Adam Storm
- Pavel Sustr
- Michael Tiefenbacher

**Roy's famous presentation of the year**  
An Audit a day keeps the lawyers at bay!

# Agenda

1. Audit – do you need it, do you care?!
2. Audit - Voting
3. Audit needs and musts
4. Solution overview and their Pros/Cons
5. The viable way – let Db2 do the magic!
6. The Auditor's way

# Agenda

1. Audit – do you need it, do you care?!
2. Audit - Voting
3. Audit needs and musts
4. Solution overview and their Pros/Cons
5. The viable way – let Db2 do the magic!
6. The Auditor's way

## Audit – do you need it, do you care?!

- The new, European wide, GDPR requires it!

<https://www.eugdpr.org/>

Intro:

The EU General Data Protection Regulation (GDPR) replaced the Data Protection Directive 95/46/EC and is designed to harmonize data privacy laws across Europe, to protect and empower all EU citizens' data privacy and to reshape the way organizations across the region approach data privacy.

Enforcement date: 25 May 2018 – from when those organizations in non-compliance will face heavy fines.

# Audit – do you need it, do you care?!

Definitions:

## Art. 4 GDPR Definitions

For the purposes of this Regulation:

‘personal data’ means any information relating to an identified or identifiable natural person (‘data subject’); an identifiable natural person is one who can be identified, directly or indirectly, in particular by reference to an identifier such as a name, **an identification number**, **location data**, **an online identifier** or to one or more factors specific to the physical, physiological, genetic, mental, economic, cultural or **social identity** of that natural person;

This includes TCP/IP addresses and Email addresses...

# Audit – do you need it, do you care?!

## Definitions:

### Art. 25 GDPR Data protection by design and by default

1. Taking into account the state of the art, the cost of implementation and the nature, scope, context and purposes of processing as well as the risks of varying likelihood and severity for rights and freedoms of natural persons posed by the processing, the controller shall, both at the time of the determination of the means for processing and at the time of the processing itself, **implement appropriate technical and organisational measures**, such as pseudonymisation, which are designed to implement data-protection principles, such as data minimisation, in an effective manner and to integrate the necessary safeguards into the processing in order to meet the requirements of this Regulation and protect the rights of data subjects.

# Audit – do you need it, do you care?!

Definitions:

## Art. 25 GDPR Data protection by design and by default

2. The controller shall implement appropriate technical and organisational measures for ensuring that, by default, only personal data which are necessary for each specific purpose of the processing are processed. That obligation applies to the amount of personal data collected, the extent of their processing, the period of their storage and their accessibility. **In particular, such measures shall ensure that by default personal data are not made accessible without the individual's intervention to an indefinite number of natural persons.**

# Audit – do you need it, do you care?!

Definitions:

## Art. 32 GDPR Security of processing

Taking into account the state of the art, the costs of implementation and the nature, scope, context and purposes of processing as well as the risk of varying likelihood and severity for the rights and freedoms of natural persons, the controller and the processor shall implement appropriate technical and organisational measures to ensure a level of security appropriate to the risk, including inter alia as appropriate:

the pseudonymisation and encryption of personal data;

**the ability to ensure the ongoing confidentiality, integrity, availability and resilience of processing systems and services;**

# Audit – do you need it, do you care?!

Definitions:

## Art. 32 GDPR Security of processing

the ability to restore the availability and access to personal data in a timely manner in the event of a physical or technical incident;

*a process for regularly testing, assessing and evaluating the effectiveness of technical and organisational measures for ensuring the security of the processing.*

In assessing the appropriate level of security account shall be taken in particular of the risks that are presented by processing, in particular from *accidental or unlawful destruction, loss, alteration, unauthorised disclosure of, or access to personal data transmitted, stored or otherwise processed.*

## Audit – do you need it, do you care?!

The fines are also amazingly high.

First, for the “minor” problem of being over 72 hours late when data leaks have occurred (a breach), is 2% of global turnover or €10,000,000 (Nearly \$12,000,000) – whichever is **higher**.

Then, if you are really naughty, like disregarding basic data laws, moving data abroad or ignoring an individual’s rights, then you get hit for 4% of global turnover or €20,000,000 (nearly \$24,000,000) – again whichever is **higher**.

## Audit – do you need it, do you care?!

### Art. 83 GDPR General conditions for imposing administrative fines

Each SA shall ensure that the imposition of administrative fines (...) be **effective, proportionate and dissuasive.**

When deciding (...) due regard shall be given to the following:

the nature, gravity and duration of the infringement taking into account the nature scope or purpose of the processing concerned as well as the number of data subjects affected and the level of damage suffered by them;

the intentional or negligent character of the infringement;

**any action taken by the controller or processor to mitigate the damage suffered by data subjects;**

**the degree of responsibility of the controller or processor taking into account technical and organisational measures implemented by them pursuant to**

**Articles 25 and 32;**

# Agenda

1. Audit – do you need it, do you care?!
2. Audit - Voting
3. Audit needs and musts
4. Solution overview and their Pros/Cons
5. The viable way – let Db2 do the magic!
6. The Auditor's way

---

## Audit – Voting

- Please vote for one of the options below

## Audit – Voting

- Please vote for one of the options below
- Option 1:



## Audit – Voting

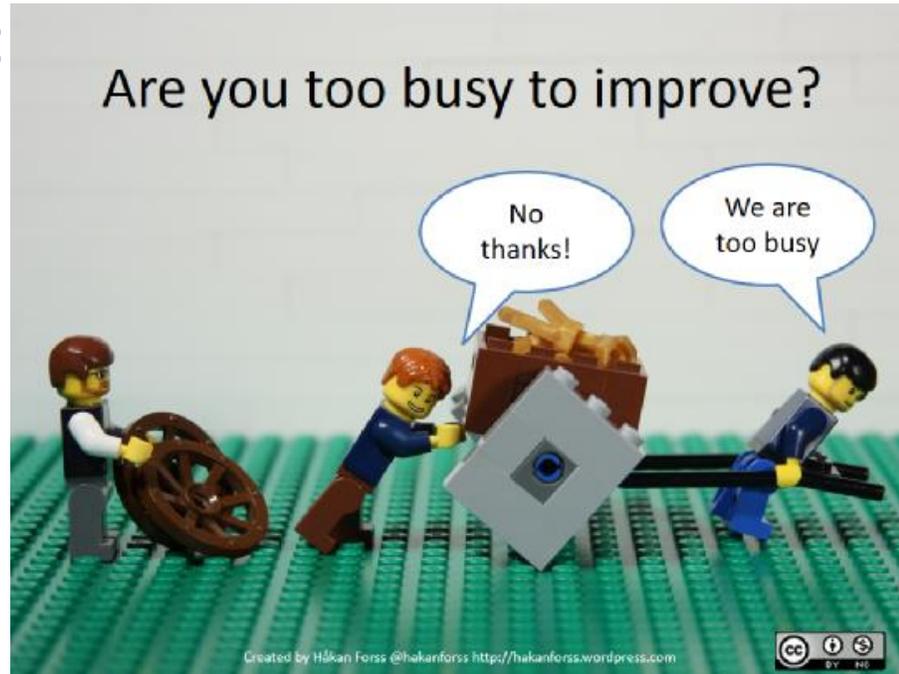
- Please vote for one of the options below
- Option 2:



© Can Stock Photo

## Audit – Voting

- Please vote for one of the options below
- Option 3:



# Agenda

1. Audit – do you need it, do you care?!
2. Audit - Voting
3. Audit needs and musts
4. Solution overview and their Pros/Cons
5. The viable way – let Db2 do the magic!
6. The Auditor's way



## Audit needs and musts

- Critical activities that enterprises should be auditing
  - Privileged Users
    - Access/changes/deletion to critical data
    - Access using inappropriate channels
    - Schema modifications
    - Unauthorized addition of user accounts

## Audit needs and musts

- Critical activities that enterprises should be auditing
  - End Users
    - Unusual access to excessive amounts of data
    - Access to data outside standard working hours
    - Access to data through inappropriate channels
  - Developers, Analysts and System Administrators
    - Access to live production systems
  - IT Operations
    - Inappropriate changes to DB/DB applications

## Audit needs and musts

- ... or in other words:

*Collect as much data as you can, because you probably don't know today what you'll need tomorrow*

**→ breach patterns do change!!!**

- Make sure you include:
  - SELECTs (against sensitive data)
  - DDL
  - DML
  - DCL
  - Utilities (online + offline)
  - Commands
  - Assignment, or change of a user ID/authorization – especially privileged users

## Audit needs and musts

- Be careful what happens outside of a table:
  - Consider clones
  - Consider backups
  - Consider extended statistics in catalog tables, like SYSCOLDIST + SYSKEYTGTDIST
  - Consider utility output (REORG, RUNSTATs)
  - Consider UNLOADs
  - Consider Replication
  - Consider access to the underlying VSAM cluster
- Also consider your INSTALL SYSADM/SYSOPR
  - Sorry DBAs, but Auditing requires a separation of duties

## Audit needs and musts

- Most Home-Grown Solutions are based on the Db2 Audit Trace:
  - Class 1, 2, 7, 8 have very little overhead
    - Access violations (Class 1 IFCID 140)
    - GRANTS/REVOKEs (Class 2 IFCID 141)
    - Assignment, or modification of a user ID/authorization (Class 7 IFCIDs 55, 83, 87, 169, 319)
    - Db2 utility (Class 8 IFCIDs 23, 24, 25, 219, 220)
  - Class 3 (IFCID 142) has very little overhead
    - DDL (only for TB having the AUDIT ALL attribute)

## Audit needs and musts

- Most Home-Grown Solutions are based on the Db2 Audit Trace:
  - Class 4, 5 (IFCIDs 143, 144) has up to 5% overhead
    - 1<sup>st</sup> INSERT/UPDATE/DELETE, SELECT in a UOR
  - Class 10 (IFCIDs 270, 271) has low overhead
    - Trusted context Create/Alter and Column mask/Row permission  
Create/Drop/Alter
  - IFCIDs 90, 91 have very little overhead
    - Db2 Commands

# Agenda

1. Audit – do you need it, do you care?!
2. Audit - Voting
3. Audit needs and musts
4. Solution overview and their Pros/Cons
5. The viable way – let Db2 do the magic!
6. The Auditor's way

## **Solution overview and their Pros/Cons**

There are a variety of existing resources Db2 already provides/comes with:

- Db2 Log
- Db2 Trace
- Db2 Memory (DSC/EDM)
- Db2 Exits

# Solution overview and their Pros/Cons

## Db2 Log:

- Pros:
  - Comes with Db2 and supports all versions
  - No additional overhead
  - No additional costs (except you want to keep logs for a longer period of time than currently and, of course, your analysis)
  - Most companies have log analysis tools they're already familiar with
- Cons:
  - Not all required data is logged
    - SELECTs are especially lacking

# Solution overview and their Pros/Cons

## Db2 Trace:

- Pros:
  - Comes with Db2 and supports all versions
  - No additional costs (except for storing and processing the collected data)
  - Most companies have trace data analysis tools they're already familiar with
- Cons:
  - Depending on the scope (number of IFCIDs/classes), and the type (SMF, OPX, GTF, SRV), the overhead may be significant
  - You need to build your own repository
  - If not using OPX you lose time!

# Solution overview and their Pros/Cons

## Db2 Trace:

- What are the differences:
  - There are different types of traces:
    - Statistics, Accounting, Audit, Monitor, Performance, Global
  - There are different classes
  - There are hundreds of individual IFCIDs
- Depending on your choice, the overhead is unmeasurable to significant
- A key difference in cost is the trace destination!
  - SMF, OPX, GTF, SRV

# Solution overview and their Pros/Cons

## Db2 Trace:

- What are the differences:
  - Processing the data requires simple to more sophisticated knowledge:
    - SMF: System Management Facility:  
Most commonly used, easy to process (use DSN1SMFP) – Once a day “cuts” cost 24 hours
    - OPn/OPX: Buffer Destination Trace  
very efficient, but Assembler needed to process (DSN1SDMP is pretty poor)
    - GTF: Generalized Trace Facility:  
Used for detailed monitoring
    - SRV: Serviceability Routine:  
I have never seen it used

# Solution overview and their Pros/Cons

## Db2 Memory (DSC/EDM):

- Pros:
  - Comes with Db2 and supports all versions
  - No additional overhead
  - No additional costs (except for storing and processing)
- Cons:
  - Not all required data is there
  - Usually you can't access it yourself, unless you hook into it
  - The information is volatile and can get lost quickly

# Solution overview and their Pros/Cons

## Db2 Exits:

- Pros:
  - Partially comes with Db2 and supports all versions
  - No additional costs (except for storing and processing)
- Cons:
  - Not all required data is there
  - Lots of coding necessary to catch and process the data
  - The overhead may be significant

# Solution overview and their Pros/Cons

## Additional Tools:

- Pros:
  - There are various solutions to choose from
  - Usually easy to use and more powerful than native Db2 options
- Cons:
  - Vendors charge for it
  - Implementation and processing overhead may be significant
  - Additional appliances lead to more vulnerability and administration overhead

# Solution overview and their Pros/Cons

## Additional Tools:

- What are the differences?
  - Good solutions have efficient data collectors and share repositories for Audit, Performance Management, Accounting, Analytics ...
  - Some solutions use hooks into the Db2 address space to capture SQL activity – errors can bring down Db2, or the entire LPAR, thus they try to protect Db2 by encapsulating the “foreign” code
  - Some solutions need additional appliances (easily up to 100+ virtual appliances) → all SQL captured is sent (unencrypted!) through the network. If the connection gets lost they try to cache it. Keep in mind that attackers do DDoS attacks!

# Solution overview and their Pros/Cons

## Additional Tools:

- What are the differences?
  - Some solutions exploit zIIP processors
    - Optional (scope)
    - Forced usage
  - Some solutions offer reporting in real-time
  - Some solutions offer alerting
    - This requires a rule, or profile setup  
→ keep in mind that they are based on known patterns
  - and of course solutions differ in
    - Setup (collector per Db2 system/LPAR)
    - Filtering
    - Dedicated support of compliance reports

# Solution overview and their Pros/Cons

## Additional Tools:

- What are the differences?
  - Some solutions have additional capabilities:
    - Covering a variety of databases (Db2 z/OS/LUW, IMS, Oracle, SQL Server, ...)
    - Covering applications (CICS, SAP, ...)
    - Covering dataset activity and Content Managers (VSAM, FTP, SharePoint, ...)
    - Covering Big Data (Hadoop, HANA, ...)
    - Covering vulnerability scanning of up to entire infrastructures (including network, firewall, workstations, ...)
    - Covering logons, connects

→ Depending on your choice it may become complex and expensive and you're locked to a specific vendor!

# Agenda

1. Audit – do you need it, do you care?!
2. Audit - Voting
3. Audit needs and musts
4. Solution overview and their Pros/Cons
5. The viable way – let Db2 do the magic!
6. The Auditor's way

## The viable way – let Db2 do the magic

The most reliable/efficient solution is based on those reliable and robust Db2 key functions we've been using for ages.

Exploiting them results in the most powerful solution:

- You benefit from rock solid features, like:
  - Security
  - Compression
  - Native Db2 functions
  - Extended Client Identification Registers, `sqleseti()`

The only question is: What key Db2 functions are needed?

## The viable way – let Db2 do the magic

Using IFCIDs along with OPX buffers delivers in-depth information without the overhead and delay of SMF processing:

316/318 Dynamic SQL (SELECT, INSERT, etc.)  
(+317 for the full SQL statement)

400/401 Static SQL (SELECT, INSERT, etc.)  
(+SYSPACKSTMT for the full SQL statement)

## The viable way – let Db2 do the magic

Using IFCIDs along with OPX buffers delivers in-depth information without the overhead and delay of SMF processing:

23/24/25	Utility start/phase/stop (219=Listdef and 220=Template)
55/83/87	SQLID setting
90/91	Commands and their completion status
140	Authorization failures
141	Authorization changes
143/144	AUDIT Table access
62/142	DDL and CREATE/ALTER/DROP for tables with AUDIT changes or all
270/271	Trusted Context and Column Masks/Row Permissions

## The viable way – let Db2 do the magic

So now you have all that data for Audit. But also now think about what else you could do with all of it...

Just imagine the performance data contained within...or the usage analysis possible...

The possibilities are endless! This is a fantastic data source created for Audit but available for performance DBAs and even developers!

---

# The viable way – let Db2 do the magic BUT:

## Make sure it's secure!

- Set up and audit access to the repository
- Alert via WTO if someone messes with the IFCIDs you've chosen
- Consider automatically cancelling threads of users violating the rules

## The viable way – let Db2 do the magic

All IFCIDs listed have a much smaller footprint than a blanket  
AUDIT CHANGES/ALL

This is integrated, reliable Db2 technology, OPX is the right target for efficient capturing. Store it in a repository and protect it using proven technology (e.g. RACF, ACF2, Top Secret)

Using Db2 compression reduces storage requirements by exploiting proven, integrated technology

→ No new vulnerabilities like:

- Black Box appliance
- Massive sensitive data transmissions over the network

## The viable way – let Db2 do the magic

Do your (automated) reporting/alerting/analytics as needed:

- SPUFI
- Batch Job
- Enterprise-wide reporting system
- GUI (DRDA based queries are fully zIIP eligible)

If you don't want to improve your Home Grown solution, find a vendor who exploits this technology!

## The viable way – let Db2 do the magic

DSC and EDM provide detailed workload insights, including flushed statements:

- SQL text
- Statement ID
- Date/time
- Current status
- Resource consumption
- Identification/environmental data



## The viable way – let Db2 do the magic

Capture the data e.g. using a STC:

Run a started task 24x7 to catch all the IFCIDs that Db2 will be throwing and store the data.

Process the workload:

Externalize and process the data frequently:

- allow ad hoc data refresh triggered via operator command for the started task (MODIFY)
- capture the SQL Text at trace time



# The viable way – let Db2 do the magic

Use a GUI front end, preferably Eclipse:

Exploit and integrate into Eclipse based GUI front ends

- GUIs can come as a Plug-in for
  - IBM Rational
  - IBM Data Studio
  - Eclipse native
- Existing Db2 connections are used to connect to the mainframe
- Interactive dialogs allow complex and powerful analysis
- Export features can create PDF reports and allow MS Excel handover

## The viable way – let Db2 do the magic

Optionally use a LEEF (Log Event Extended Format) reformatter program for the SIEM system of your choice!



```
LEEF:1.0|Software Engineering GmbH|WorkLoadExpert Audit|6.1|
IFCID 090|cat=success|devTimeFormat=yyyy-MM-dd'T'HH:mm:ss.SSSZ|
devTime=2018-03-09T09:57:33.886+0100|Sev=01|usrName=GABELMA|
name=|usrPriv=|usrGroups=|src=|subsys=DC10|dsn=|plan=MVNXPLAN|
objtyp=|obj=|intent=|SQLid=GABELMA|poe=|submitby=|job=Z100 DC10|
cmd=-DIS GROUP |checkid=|conn=DC10 location Z100DC10 LU DESWEG01.Z100DC10
group DC10 member DC10 connector DB2CALL GABELMA operator GABELMA
workstation DB2CALL tx GABELMA enduser GABELMA|sum=DB2 DC10 GABELMA
Command Issued by id GABELMA:-DIS GROUP
```

## The viable way – let Db2 do the magic

These days most z/OS Audit systems collect data and transfer to a Data Lake of your choice for post processing every two to three hours e.g. WorkLoadExpert, zSecure etc.

This data is typically RACF, SMF and Master Log data on its way to e.g. QRadar, Splunk, AlienVault et al



# Agenda

1. Audit – do you need it, do you care?!
2. Audit - Voting
3. Audit needs and musts
4. Solution overview and their Pros/Cons
5. The viable way – let Db2 do the magic!
6. The Auditor's way

# The Auditor's way



## The Auditor's way

As external auditors become technically smarter and smarter, they like to ask, within the scope of an IT Audit, the following:

Main focus is, in general, the accounting relevant data.

Please provide the authorization concept and how the setup of roles and permissions for the company's databases is realized.

What are the current versions of the database? (Not z/OS really!)

## The Auditor's way

Provide a list of all created users of the database, including the roles e.g. of admins, technical users etc. If there are functional accounts please give a brief note for which tasks.

Is it theoretically possible for an individual person to log on as a technical user?

Who has the access data for technical users?

In case that a manual logon is not possible (e.g. the technical user is protected or revoked) please provide a list of individuals who are capable of changing this status with the result that a manual logon is again possible.

## The Auditor's way

What are the complexity guidelines for the password structures?

Who has access to SYSADM/SYSADMIN?

Is there a protocol in place to check all database activities of DBAs and technical users? e.g. via SMF

Which activities are logged (recorded)?

Who is able to change the parameters of the protocol?

## The Auditor's way

Where, and in which format, are these protocols / logs cached or filed (e.g. VSAM?)

Do the DBAs have the power to modify the database logs / protocols directly or indirectly (e.g. via manipulation of the VSAM files?)

Is there a review in place for the filed protocols? If so, by whom, and how often is it done?

Description of the revision process for the giving/granting of authorizations (monitoring/review)

## The Auditor's way

Proof of the execution of monitoring process over the entire last year.

How long into the past is it possible to restore the databases?

Who has access to the database dumps/back-ups?

Last but not least: Please provide a listing of **all** applications which are using the database and specify **all** relevant tables.

# The Auditor's way

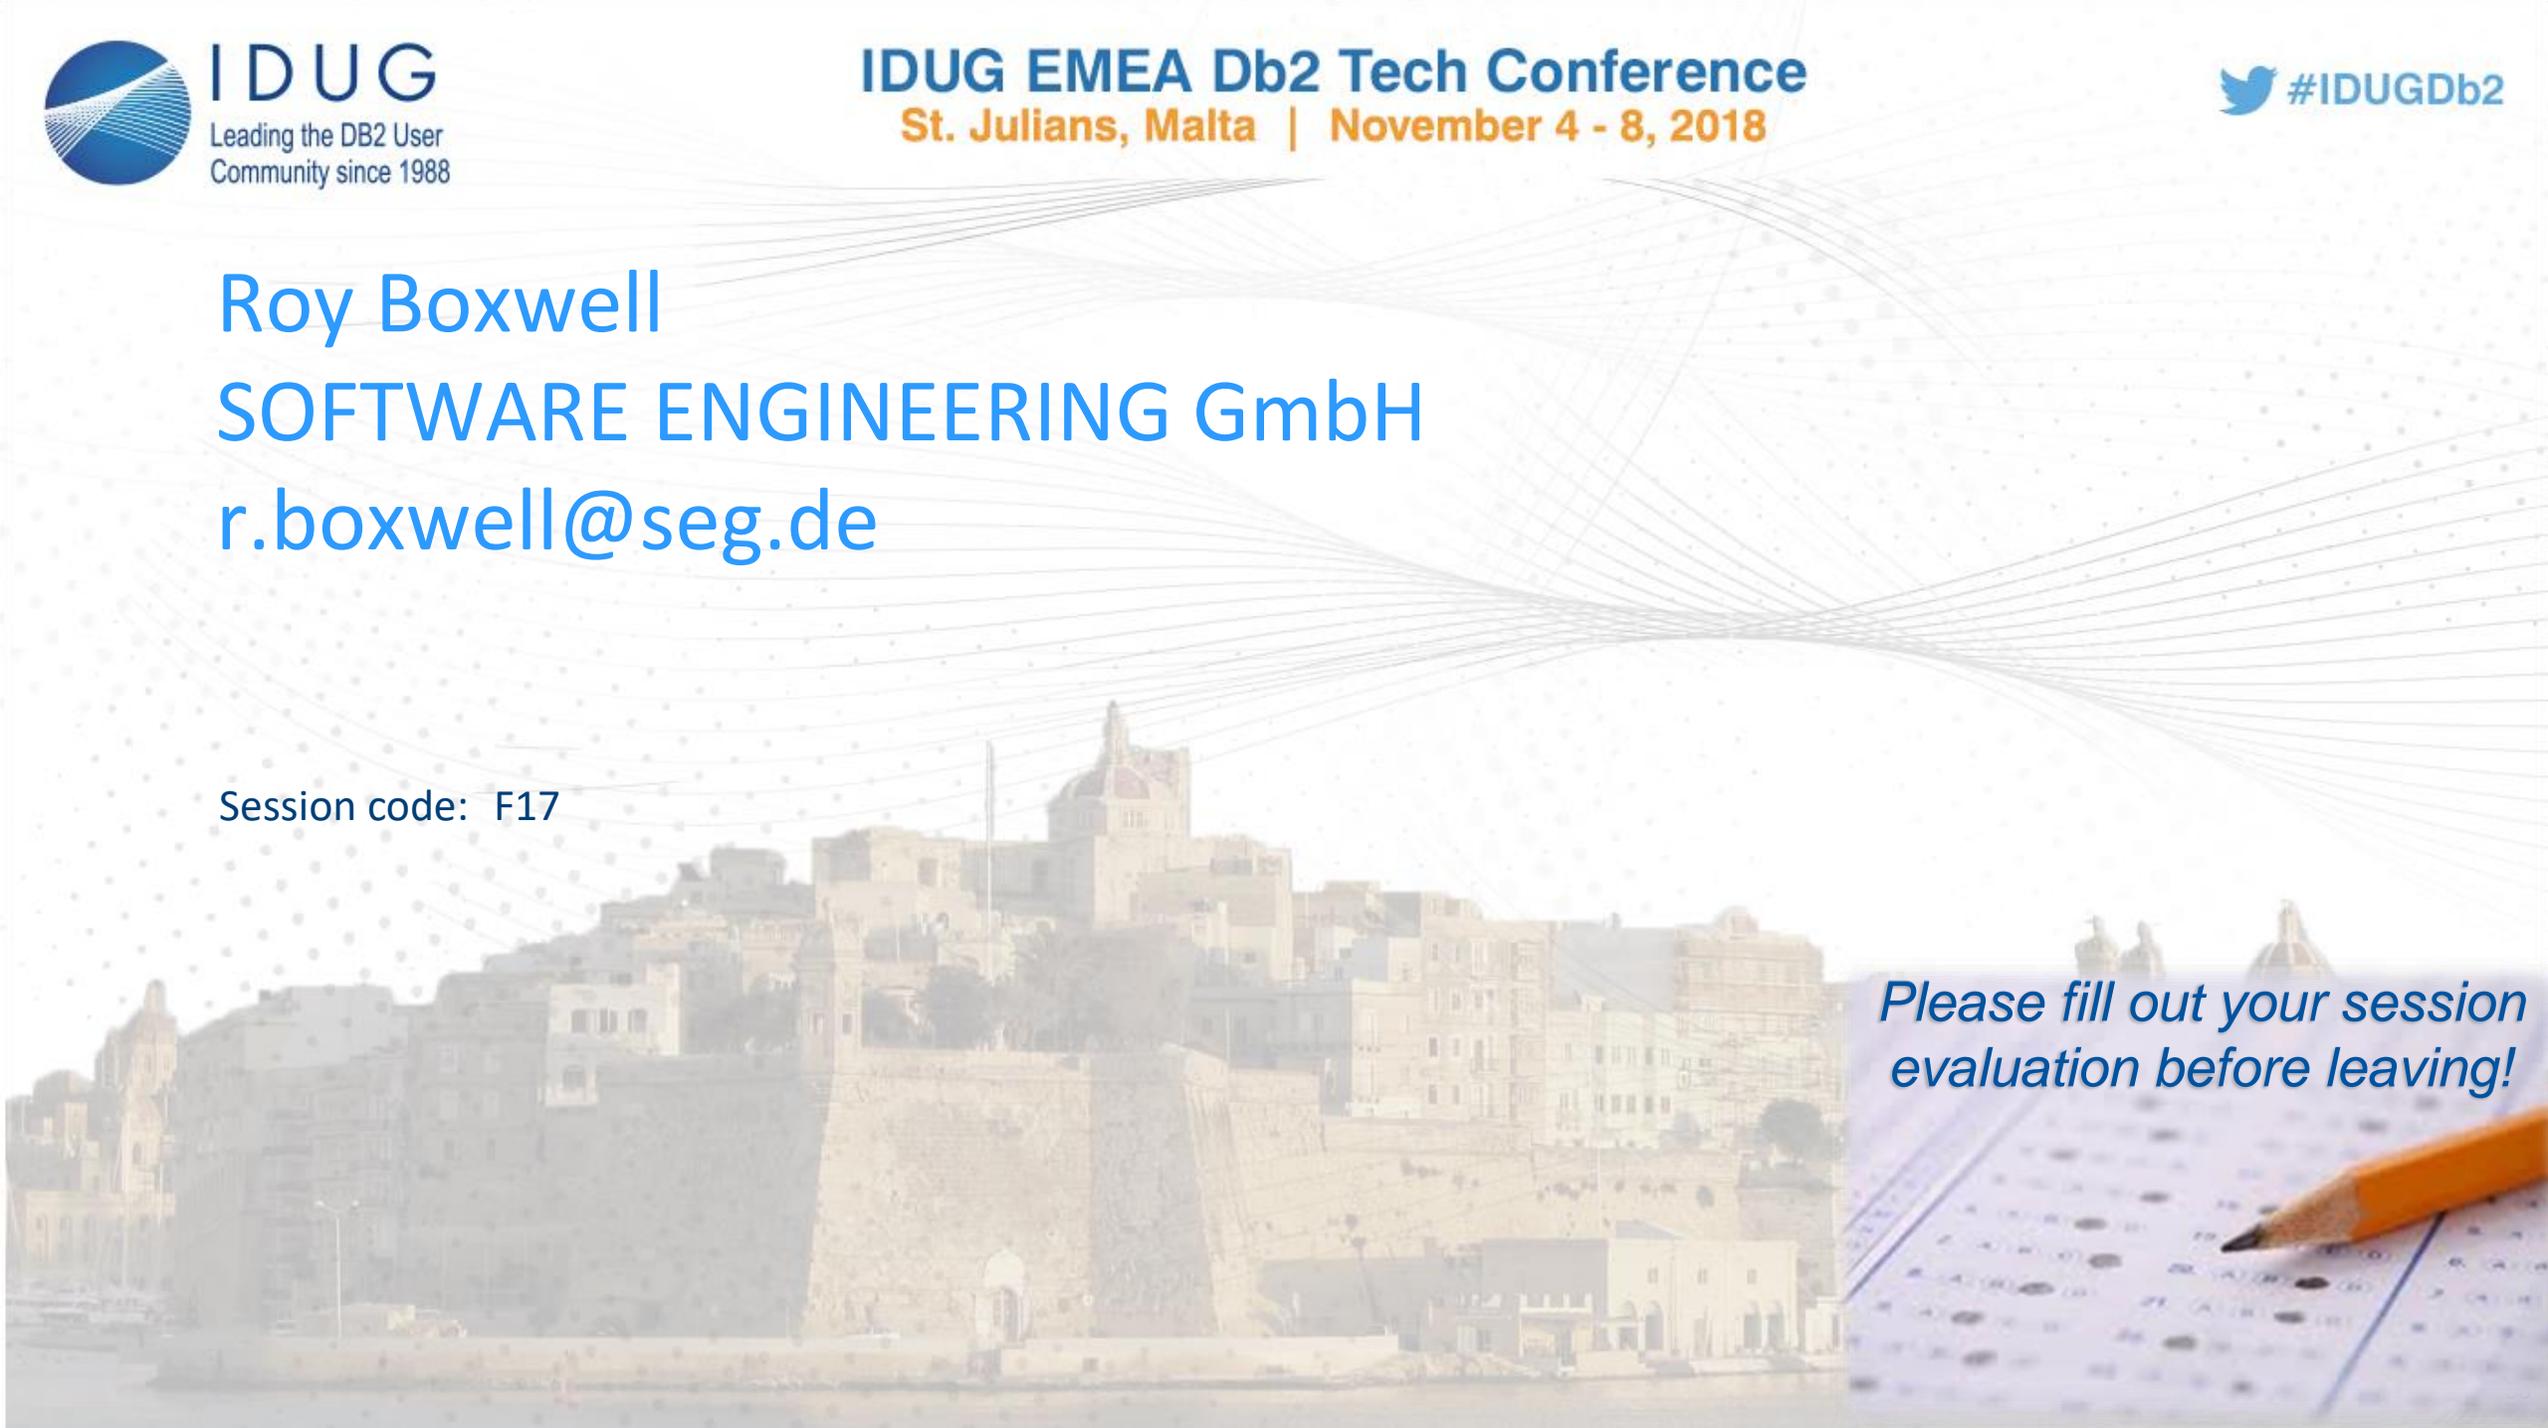
So now you know!

The Auditor's are getting cleverer...

Are you?

Roy Boxwell  
SOFTWARE ENGINEERING GmbH  
r.boxwell@seg.de

Session code: F17



*Please fill out your session  
evaluation before leaving!*