

July 2023, Edition #30

ORAWORLD

e-Magazine for Oracle Users published by the EOUC

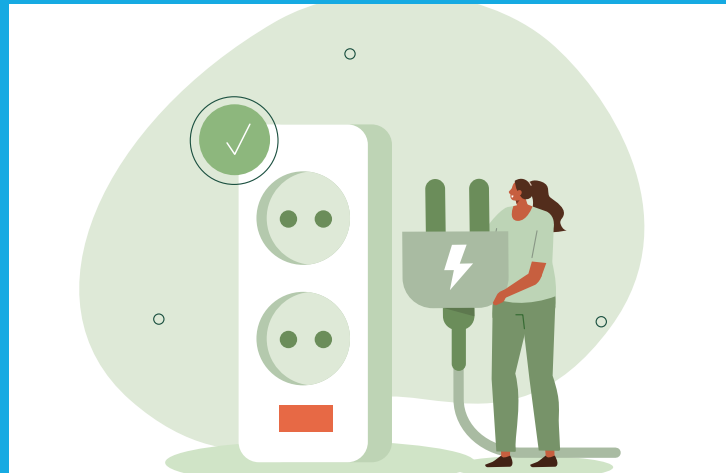
Green Software

Back to OUGN

Push to APEX 23.1!



LLMs, GPTs, and All That Jazz



Green Software – Using Resources Efficiently

page 7



Oracle Application Express (Part 12): Push to APEX 23.1!

page 17



Back to OUGN – Creating a Conference Post Pandemic

page 39

Editorial

Submit Your Article!

I. Work & Life

Comic:
Once Upon a Teams Meeting

Green Software –
Using Resources Efficiently

3 II. Techs & Nerds

4 LLMs, GPTs, and All That Jazz 12

Oracle Application Express (Part 12)
Push to APEX 23.1!
Template Components,
Push Notifications and more... 17

5 Powerful Intercompany
Eliminations of Oracle EPM
Financial Consolidation and
Close Services (FCCS) 24

One Goal, Different Approaches:
Performance Tuning in Practice 31

III. Users & Groups

Back to OUGN –
Creating a Conference
Post Pandemic 39

IV. Past & Future

Call for Papers 41

Events 42

Contact us 44

Legal notice 44

Editorial

Dear Oracle User Group Community Members,

What a pleasure that you have chosen to explore this issue of Oraworld which is packed with Oracle technology and exchange of experience highlights.

This edition concentrates on the most useful information from community leaders and Oracle ACE Pros. You will find something new about the background of the ChatGPT from Heli Helskyaho. I liked to read her article. It explains how you can use this technology within Oracle. In addition, it was really interesting for me because my son is also trying to program with AI models such as ChatGPT on the ProgKids platform. It includes interactive activities with ChatGPT, and games allowing children to experiment and apply their new skills in real life situations.

In addition, this issue provides an Oracle User Group story written by Ann-Sofie Vikström Often, a great example of the first successful in-person conference in Norway after COVID time. This is a nice story about the Norwegian Oracle Community which helps to share experience with their members. Such in-person engagement with the community provides best practices in Oracle technology. Beside that, you will know why Oracle ACE Pro wants to get back to physical conferences.

And of course, another key topic in this issue is APEX 23.1. The article "Push to APEX 23.1" by Carsten Czarski contains several new features and enhancements. Then, I really enjoyed reading about performance tuning and tools in the article "One Goal, Different Approaches: Performance Tuning in Practice" by Dani Schnider and Martin Berger.



Also, for Oracle APPS people, we have a nice story about Powerful Intercompany Eliminations of Oracle EPM Financial Consolidation and Close Services (FCCS) by ACE Pro Mohit Jain and Megha Gupta. Finally, we have a great article about software developers that can bring their contribution to greener software written by Michael Krämer.

This issue is pleasant and easy to browse through. This makes it ideal for Oracle Community users who do not like heavy reading. Enjoy!

Yours,
Andrejs Vorobjovs
LVOUG president (Latvia), EOUC board member



Submit Your Article!

You have an interesting topic to publish in ORAWORLD Magazine? Then submit your article and be part of it!

Please e-mail us your article via the online form at www.oraworld.org.



„Once Upon a Teams Meeting“

www.commitstrip.com




CommitStrip.com

CommitStrip is a daily strip recounting funny anecdotes of life as a coder mixed up with a dash of topical tech news. Find more comics here: www.commitstrip.com

Follow us on  @EOUC @ORAWORLD_Mag #ORAWORLD

 @emeaoracleusergroups @ORAWORLDMag

 www.oraworld.org

NUMBER OF THE MONTH:

1 HOUR

That's all it took Threads - the Twitter's competitor app - to reach 1 million users. If we look at the time elapsed since the launch of other online platforms, this a record hard to beat.

Online Service	Launch Year	Time Taken to Reach 1 Million Users
Threads	2023	1 hour
ChatGPT	2022	5 days
Instagram***	2010	2.5 months
Spotify	2008	5 months
Dropbox	2008	7 months
Facebook	2004	10 months
Foursquare***	2009	13 months
Twitter	2006	2 years
Airbnb**	2008	2.5 years
Kickstarter*	2009	2.5 years
Netflix	1999	3.5 years

*1 million backers **1 million nights booked ***1 million downloads

Source: <https://explodingtopics.com/blog/chatgpt-users>

Threads' strong start is due to a combination of factors: it benefits from the existing scale and infrastructure of Meta - the "mother" app - and many Twitter users have expressed desire for an alternative since Musk took over the platform late last year.

Time will tell if Threads will remain a success because, as Instagram CEO Adam Mosseri pointed out, the challenge for a new social platform is not about getting users to sign up, but rather keeping them engaged long-term.

Green Software – Using Resources Efficiently

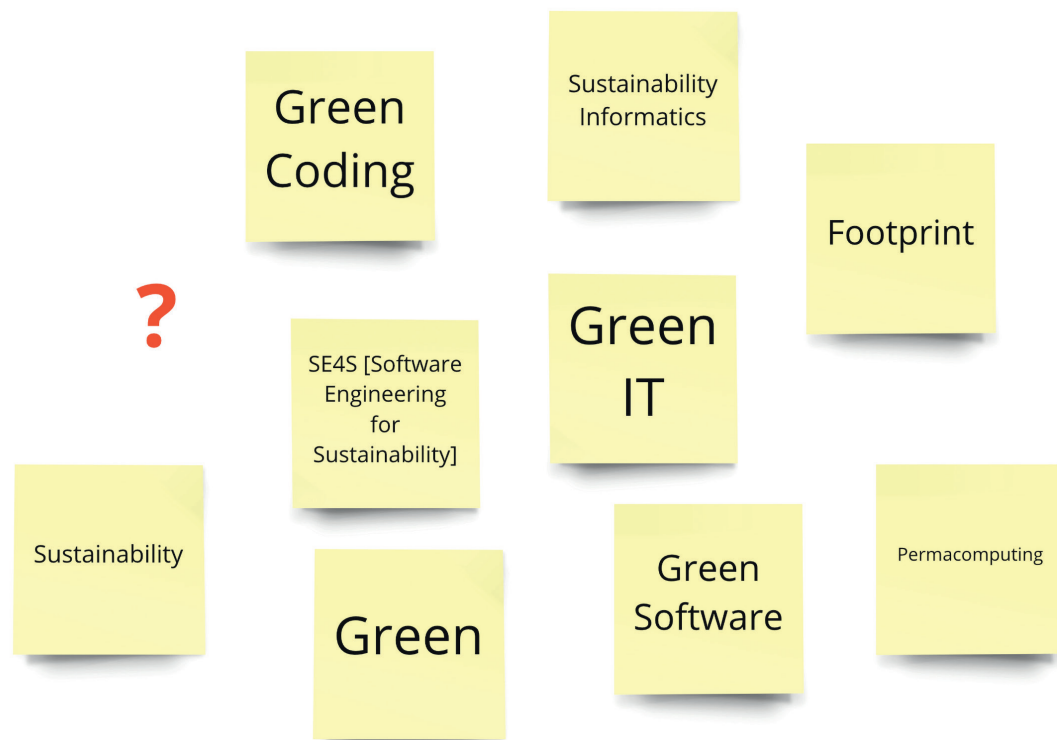
Michael Krämer



In order to reduce the footprint of a software solution or to make it as small as possible directly, there are three levels at which one can start: the infrastructure level, to make the runtime environment as efficient as possible; the concrete implementation of the software, in order to make the program itself efficient; the technical level, for example, to design processes in such a way that they can be implemented efficiently, or to question the costs and benefits of individual functions.

We should not rely exclusively on one of these levels. In this article we will focus on the implementation of the software in order to examine in concrete terms the possibilities that we, developers, have for making the program efficient (see figure below).

Software developers can bring their contribution to greener software by implementing software that executes efficiently at runtime. It is undisputed that there are more and less efficient approaches. But unfortunately we seem to have lost sight of this somewhat. In the area of server-based software, sufficient performance has been available for the vast majority of applications for more than ten years. In addition, virtualization and the use of cloud environments have led to a very strong decoupling from hardware, scaling options do not require hardware delivery and can be implemented in a short time. I see it as a consequence of the environment



we are used to that we have become accustomed to giving less weight to quality criteria efficiency, and instead give preference to maintainability and modularity.

Writing efficient software is not a new job

In other areas, such as embedded computing, efficiency has always been and is high on the priority list. In app development for mobile devices, a more tangible issue is battery life and the resulting energy consumption during use.

Assuming that resource consumption will become more expensive in the future or consumers will increasingly demand efficient products, server-based software will also have to pay more attention to efficiency. It is important at this point to distinguish between fast and efficient software. In some cases, these characteristics complement each other, but software can also become fast through excessive use of resources.

Today, however, it is still the case that investment in optimization is only worthwhile in purely economic terms if a system is highly utilized, because development is still expensive compared to

Hardware. I have set myself the goal of resource efficiency as a quality objective in the architecture of all projects, so that at least an active decision is made about its prioritization. I think this is an important aspect worth to raise awareness on it.

Avoid the blind flight

But what exactly should you optimize now? One developer wisdom is to avoid premature optimization. If we want to trim our software to low resource consumption, we must first find a way to identify the places where it is worth optimizing. In the same way, it should later be possible to check the effect of the optimizations. One problem here, however,



is that even the terms in the field of sustainable software development are not clearly defined. It is even more difficult to find a concrete metric and to calculate it. Global and absolute metrics that make a statement about how many resources a piece of software consumes are currently very difficult because, for example, it is not clearly defined which effects should be included in the measurement and which should not. Examples of such unclear delimitations are the production and supply chain of the hardware or the question of whether data traffic to the end device is included. Comparative measurements are - similar to other metrics for the quantitative assessment of software - are much easier to handle and are sufficient to assess optimizations.

In doing so, one should try to find the simplest meaningful unit. This does not have to be a CO2 equivalent. Within a system and a product group, for example, you can count and compare used nodes [1]. Costs can also be compared in the cloud within such a category. It becomes more complicated when different product categories are used and managed services such as databases are used. The major cloud providers have dashboards that show CO2 consumption, but these have their pitfalls. They require permissions to access the billing account, which users from technical departments often don't have available. Some products don't seem to show up in the dashboards at all; for example, on a larger project, we don't see the use of Elastic Container Service in the AWS Carbon Footprint Tool. In addition, the statistics are only available after up to three months, which is clearly too long for a feedback loop, especially since it is not to be expected that the used software will remain free of changes over such a period of time.

A possible alternative is provided by the cloud-jewels model of Etsy [2], which also makes it possible to compare different

types of cloud services and different providers - apples with oranges, so to speak - by means of heuristics. This model also forms the basis for the Cloud Carbon Footprint software [3]. In general, it can be said that such global measurements, which estimate the consumption of an entire component or platform, can be implemented in cloud environments. However, due to the observation periods, only the long-term trend may be meaningful. In private environments away from the public cloud providers, you can actually perform such measurements with an energy meter and select the observation period yourself. Effects such as fluctuating loads or increasing customer numbers over time must be calculated from this data in any case. The article on Etsy Cloud Jewels also gives examples of this.

Off to the white box

If you want to measure more specifically, there are several options for this as well. Powertop [4] provides the estimated energy consumption per process and can also be used in cloud environments. Since it was actually designed to measure battery consumption in laptops, the accuracy of the values in mains operation is apparently not entirely assured. If you work with Java, Powertop only provides statements per JVM instance. In the Java environment there is the tool JoularJX [5], a Java agent that outputs which methods of a software consume how much energy. This is a unique feature and at first gives great pleasure in exploring the behavior of one's software at runtime. JoularJx works via

RAPL [6] on Linux, an interface to read energy consumption over hardware as accurately as possible. There is also a version for Windows, but none for Mac. Unfortunately, this interface is only usable in bare metal environments. Virtualized or container environments are not compatible with it. Another practical difficulty is that JoularJX writes



files every second, resulting in a lot of files quickly. Assuming that very few production environments run directly on bare-metal, one must set up a separate system to examine one's software. It should be noted that it must be comparable to the production system in terms of behavior.

In the search for other, more practical approaches, the use of a JVM profiler seemed plausible. In fact, it can be used to find out very well which components and methods of a software require the most CPU time. These are the candidates for optimizations, as it was already away from Green Software. An alternative to profilers are metrics tools like Micrometer [7] or Dropwizard Metrics [8], which can do the same thing and are great for running persistently in systems. In any case, one should make sure to find the relevant parts of a system where optimizations can have a visible effect in the end.

As about optimization of components or methods - there are a lot of papers and studies about the energy consumption of code sections or classes, for example for the the Java Collections [9]. In practice, however, this will probably be of lower importance.

Recipes

Concrete potentials that can be found in many current systems are, for example, to use the autoscale function of a cluster manager also to scale down. From an energy perspective, it would be optimal to completely shut down systems that are not being used and only start them up again when needed.

Cloud functions or lambdas offer these possibilities, but have not been of great importance in many projects so far. This could be due to the limitations they initially had.

The possibilities have expanded considerably in recent years. Especially KNative offers the possibility to use such a technology with more control and freedom in an own cluster. You should definitely investigate this option if your solution includes components whose load is not uniform.

Another aspect that can be highlighted is Continuous Integration. Long running tests are not only a potential problem for Green Software, they are often a barrier to developer productivity and satisfaction. Revise the test strategy, make the software easy to test, and then run simple but meaningful tests, taking into account the test pyramid. This is an example of how the drive for Green Software can create synergies that have further positive effects elsewhere. Another option is to stop running all tests on every commit. One indicator to shed light on Continuous Integration is when CI servers and development tools require a lot of resources.

Critical questioning of data flows is another such recipe: generic or generated data structures in particular tend to contain large amounts of data, only a fraction of which is then used in several components. Sorting or filtering of data should also take place only once and then be passed on accordingly. Of course, the functions of the database should be exploited where possible. This is especially important when ORM frameworks are used and tables are mapped to object structures 1:1. Two remarks are still permitted. In order to be able to concretely assess changes, it is important to fix some of the possible variables. Therefore, for simplicity, it was assumed that using fewer computer resources also results in less load. Thus, although no emissions can be derived, comparisons can be made. Second, the level of expertise provides very significant leverage in some situations. Particularly in the context of an agile development process, it is often possible to incorporate the findings from




the implementation in order to be able to be able to design further features in an optimal way.

Good luck with the implementation, every contribution counts.

Resources

- [1] AWS re:Invent 2017: Deep Dive on Amazon EC2 Instances, Featuring Performance Optimiz
<https://www.youtube.com/watch?v=mZy6E2I5Rek&t=815s>
- [2] Etsy Cloud Jewels
<https://www.etsy.com/codeascraft/cloud-jewels-estimating-kwh-in-the-cloud>
- [3] Cloud Carbon Footprint
<https://www.cloudcarbonfootprint.org/docs/getting-started>
- [4] Using Powertop to Lower System Power Usage
<https://cloud-cow.com/content/using-powertop-to-lower-system-power-usage/>
- [5] JoularJx
<https://www.noureddine.org/research/joular/joularjx>
- [6] Energy measurements in Linux
<https://blog.chih.me/read-cpu-power-with-RAPL.html>
- [7] Micrometer Application Monitoring
<https://micrometer.io/>
- [8] Metrics - Dropwizard
<https://metrics.dropwizard.io/4.2.0/>
- [9] Energy Profiles of Java Collections
<https://dl.acm.org/doi/10.1145/2884781.2884869>

 This article was first published in Java aktuell magazine issue 02/2023



About Michael Krämer

Michael has been developing software for over 15 years and works as a software architect at INNOQ. In his projects he is very committed to developing components with clear responsibilities and to develop technically solutions for business requirements. He is also involved with machine learning and the integration of ML models into production-ready software environments and gives trainings for software architecture.



LLMs, GPTs, ^{Heli Helskyaho} and All That Jazz



MAY I HELP YOU ?

Everybody is talking about ChatGPT and other similar tools. What are they and how can they be used? ChatGPT, as well as Bard, Bing, DALL-E, Midjourney, Codex and many more, belong to a machine learning category called Generative AI (GenAI). The idea of a GenAI is to generate something, for example text, images, videos, audio, and 3D models. GenAI learns patterns from existing data to generate new and unique outputs. It does not really “know” things, it just uses those patterns and combines them. A technology called transformer neural network was first introduced in 2017. Large Language Models (LLMs), that for example ChatGPT uses, are based on this transformer architecture and have made significant advancements in natural language processing. The acronym GPT comes from words Generative Pre-trained Transformer. We will discuss the technology in later issues of ORAWORLD. In this article we will talk about how a GPT tool can be used and what are the risks and limitation you should be aware of. We will use ChatGPT as an example.

Because of the technology used it really does not matter what language is used or if two languages are used together. For example, in **Fig. 1** we ask the ChatGPT to write a platitude for LinkedIn. We ask it in Finnish and in the request, we tell we want to have it in English.

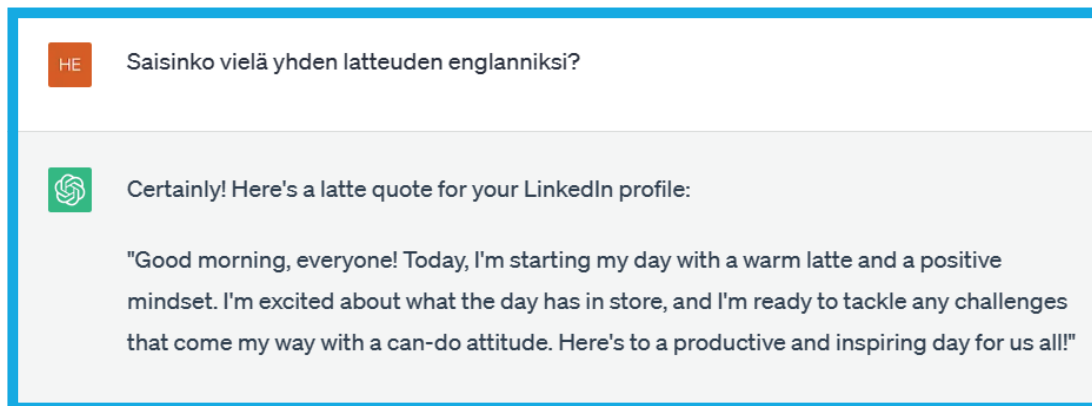


Fig. 1: Asking for a platitude. The request is in Finnish asking for a reply in English.

In **Fig. 2** you can see part of a poem written by ChatGPT about GPT.

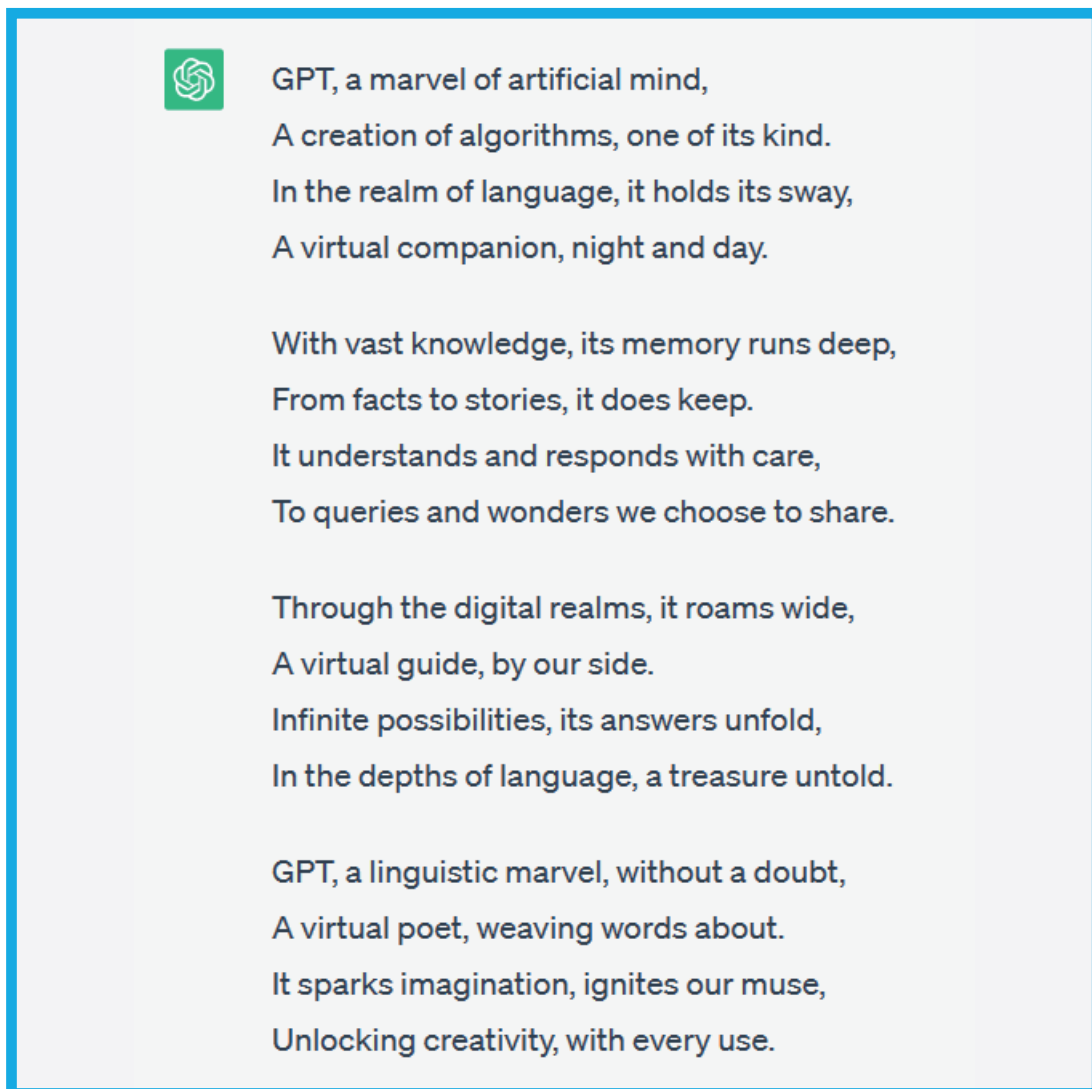


Fig. 2: Part of a poem about GPT written by ChatGPT.

The model improves very fast. Some months ago, I asked it to tell me how to create a homography matrix for images using PL/SQL. That is a task easily performed using Python libraries, but I wanted to do it using PL/SQL. In reply I got an advice to use a PL/SQL package that does not exist. Now that I asked the same question, I got this:

```
CREATE OR REPLACE FUNCTION calculate_homography_matrix(
  x1 IN NUMBER,
  y1 IN NUMBER,
  x2 IN NUMBER,
  y2 IN NUMBER,
  -- Add more parameters as needed
  homography_matrix OUT SYS.DBMS_TYPES.BLOB
) RETURN NUMBER IS
  -- Define local variables if needed
BEGIN
  -- Perform the necessary calculations to generate the homography matrix
  -- Assign the resulting matrix to the homography_matrix output parameter
  -- Example: Creating a dummy 3x3 homography matrix
  homography_matrix := SDO_GEOMETRY(
    2002, -- 2D Transformation
    NULL, -- Coordinate system identifier
    NULL, -- Coordinate system SRID
    SDO_ELEM_INFO_ARRAY(1, 3, 1), -- Element information
    SDO_ORDINATE_ARRAY(x1, y1, x2, y2, 1, 1) -- Coordinate values
  ).TO_BLOB;
  RETURN 1; -- Return a success status code or an appropriate value
END;
```

Listing 1

ChatGPT can produce code but the code **must** always be tested. As mentioned earlier what GPT does is **generate** text. The text can be correct or it might not. The “invented” results that are not correct are called hallucination. Hallucination is one of the research and improvements topics in the field. It has been tried to solve for example by letting the GPT connect to internet or to use a vector database. Even though you use a GPT tool, you must know your area to verify that the answers are correct and whenever possible, test them.

GPT only “knows” areas that the data used to train it covers. If it has never been trained with data, let’s say about Oracle technology, it would not be very helpful in generating PL/SQL code for example. Also, the data used for training might have been bias and cause bias in generated results. There might also be ethical issues with the results. Once again: the user is responsible on publishing any kind of results and therefore responsible if the result is not ethically appropriate. A very important issue is security. When you use the tool with your data, make sure the data is not sensitive. You do not know where the data is stored and how it is used. The copyright

issues are still unsolved globally. My recommendation is to always mention that the result was generated by the tool you used. When you understand these basics, risks and limitations, you are ready to use the tools.

All the GenAI tools are used by prompting. The better your prompt is the better results you get. If the result is not good enough, you can prompt more. Let's see an example. I'll prompt ChatGPT with the following:

"I am new to database designing. I should write a report to my boss about the process. What should that report include?"



ChatGPT:

"When writing a report about the process of database designing for your boss, it is important to provide a comprehensive overview of the key aspects and considerations involved. Here's a suggested outline for your report:

Introduction:

Provide a brief introduction to the purpose and importance of database designing. State the objectives of the report and what your boss can expect to gain from it.

Database Design Overview:

Explain the concept of database design and its significance in organizing and managing data effectively. Describe the benefits of a well-designed database, such as improved data integrity, efficiency, and scalability.

Requirements Gathering:

Discuss the initial phase of database design, which involves identifying and documenting requirements. Explain the process of gathering requirements, including interviewing stakeholders, analyzing existing systems, and defining data entities and relationships.

Conceptual Design:

Describe the creation of a conceptual data model based on the gathered requirements. Explain the use of techniques like entity-relationship diagrams (ERDs) to represent entities, attributes, and relationships.

Logical Design:

Detail the translation of the conceptual data model into a logical data model. Discuss the normalization process to ensure data integrity and eliminate redundancy. Explain the

use of data definition language (DDL) to create tables, define constraints, and establish relationships.

Physical Design:

Discuss the conversion of the logical data model into a physical database design. Explain considerations such as storage structures, indexing, and partitioning for optimal performance. Mention any platform-specific optimizations or configurations applied.

Security and Access Control:

Highlight the importance of data security and access control in database design. Discuss techniques like authentication, authorization, and encryption to protect sensitive data. Mention any security standards or regulations that must be considered.

Data Migration and Integration:

Discuss the process of migrating data from existing systems to the new database. Explain the challenges and strategies involved in data integration, especially if multiple data sources are involved.

Testing and Quality Assurance:

Explain the importance of testing and quality assurance in ensuring the reliability of the database. Discuss different types of testing, such as unit testing, integration testing, and performance testing. Mention any tools or methodologies used for testing.

Documentation and Maintenance:

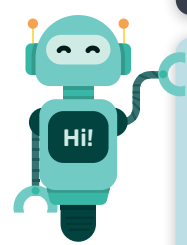
Emphasize the need for comprehensive documentation of the database design and ongoing maintenance processes. Discuss the importance of documenting schemas, data dictionaries, and procedures for future reference. Explain the role of regular maintenance activities, such as backups, data archiving, and performance tuning.

Conclusion:

Summarize the key points covered in the report. Highlight the benefits and significance of a well-designed database. Offer recommendations for further improvement or future projects.

References:

Include a list of references or sources consulted during the preparation of the report. Remember to tailor the content and level of detail to your boss's understanding and specific requirements."



Now, I need to read it carefully and decide if it makes sense. Since I am new to the topic, I do not know what Conceptual Design means. I can continue this conversation by asking:

"What is "Conceptual Design"?"



ChatGPT:

"Conceptual design, also known as conceptual data modeling, is an early phase in the database design process. It involves creating a high-level representation of the database system, focusing on the overall structure, entities, attributes, and relationships between them. The main objective of conceptual design is to capture the business requirements and provide a conceptual understanding of the database before diving into the technical details.

Here are key aspects of conceptual design:

Business Understanding: Conceptual design starts with a thorough understanding of the organization's business requirements, processes, and data needs. This involves gathering information from stakeholders, subject matter experts, and existing documentation to identify the main entities and their relationships.

Entity-Relationship Modeling: The primary technique used in conceptual design is entity-relationship modeling. This modeling approach represents the main entities in the system as well as their attributes and relationships. The entities are represented by rectangles, attributes are depicted as ovals, and relationships between entities are represented by lines connecting them.

Entities and Attributes: Conceptual design focuses on identifying and defining the main entities and their attributes. Entities represent real-world objects, concepts, or events that are relevant to the system being developed. Attributes describe the properties or characteristics of these entities. For example, in a student management system, entities could include "student," "course," and "instructor," while attributes of the "student" entity might include "name," "student ID," and "date of birth."

Relationships: Conceptual design determines the relationships between entities. Relationships depict the associations, connections, or dependencies between entities. For example, a student can be related to a course through a "takes" relationship, indicating that the student is enrolled in that course. Relationships are classified as one-to-one, one-to-many, or many-to-many, depending on how entities are

related to each other.

Cardinality and Participation: Cardinality defines the number of occurrences of one entity that can be associated with another entity in a relationship. Participation describes whether an entity's involvement in a relationship is optional or mandatory. These concepts help define the nature and constraints of the relationships between entities.

Abstraction and Simplification: Conceptual design aims to abstract and simplify the complex real-world domain into a manageable representation. It focuses on the essential entities, attributes, and relationships without getting into implementation-specific details. This level of abstraction helps ensure that the design captures the core requirements and is not influenced by specific technical considerations.

The output of the conceptual design phase is typically an Entity-Relationship Diagram (ERD) that visually represents the entities, attributes, and relationships identified during the process. This diagram serves as a foundation for further stages of the database design process, such as logical design and physical implementation."

I can continue on this topic with ChatGPT until I have my topic covered. I can later go and see the chat chain if needed.

ChatGPT is easy to use if you just remember you are responsible for the results when using it. To start with it maybe ask the following:

"Create a beginner's guide to use ChatGPT".

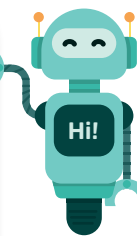


The guide I got from ChatGPT is:

"Welcome to the beginner's guide to using ChatGPT! ChatGPT is an AI language model designed to provide helpful responses and engage in conversations on a wide range of topics. Here are some steps to get started:

Understand the Basics:

ChatGPT is an AI model and can generate text based on the prompts you provide. It is trained on a vast amount of internet text and can provide information, suggestions, or creative responses. ChatGPT doesn't have real-time awareness, so it doesn't have up-to-date information beyond its September 2021 knowledge cutoff.



Frame Your Prompt:

Begin by framing your conversation with a clear and specific prompt. For example, instead of asking "Tell me about dogs," try "What are some interesting facts about Golden Retrievers?" Providing context helps ChatGPT understand your query better and provide more relevant responses.

Experiment and Iterate:

Feel free to experiment with different phrasings or prompt structures to get the desired response. If you receive an answer that's not quite what you were looking for, try rephrasing the question or providing more specific details.

Use System Messages:

You can use system-level instructions to guide ChatGPT's behavior. For example, you can start with "You are a helpful assistant who knows a lot about science" to receive scientific explanations or facts. System messages set the behavior for a few turns of the conversation and can be a useful tool.

Be Specific:

When asking questions or seeking information, try to be specific and provide relevant details. Instead of asking "What should I do in London?", try "What are some recommended vegetarian restaurants in London near Hyde Park?"

Be Patient and Iterate:

ChatGPT is an AI model and may not always provide perfect or complete responses. If you encounter an unhelpful or incorrect answer, try rephrasing or providing additional context to guide it better.

Verify Information:

While ChatGPT strives to provide accurate information, it's essential to verify critical or time-sensitive details from reliable sources. Cross-check important facts, statistics, or advice with trusted websites or experts.

Explore Different Topics:

ChatGPT is designed to discuss a wide range of subjects, so feel free to explore various topics and have fun with your conversations. From discussing movies and books to getting assistance with technical questions, ChatGPT is here to help. Remember, ChatGPT is a tool, and while it can provide helpful information and engage in conversations, it's always good practice to verify critical details and use your judgment when making decisions. Enjoy your interactions with ChatGPT!"

I hope this article helps you to start with GPTs. In the next issue of ORAWORLD, we will talk more about the topic. Happy GPT!

**About Heli Helskyaho**

Heli is the CEO for Miracle Finland Oy and an Oracle ACE Director. Heli has been interested in data over 30 years. Data and using the data are her favorite things in IT.

Oracle Application Express Carsten Czarski (Part 12)

Push to APEX 23.1!

Template Components, Push Notifications and more...

APEX 23.1 is available since May 2023. Marquee features are the brand-new *Template Components*, support for *Push Notifications*, the new *Object Browser*, and the ability to *run APEX page processes in the background*. In addition, this new release contains several smaller new features like the *Invoke API process type* now supporting REST, support for *Property Graph* queries on Oracle Database 23c, a new *Dashboard for APEX Instance Administrators* and much more.

Template Components

Template Components are a brand-new approach to define re-usable UI components with support for actions, menus, and custom attributes. Template Components are technically Plug-Ins, but unlike “traditional” Plug-Ins, Template Components do not include procedural PL/SQL Code. A template component Plug-In basically consists of a *HTML template* (the **Partial**), and parameters to control the dynamic parts of the template.

Universal Theme in APEX 23.1 includes a few Template Components out of the box: **Avatar, Badge, Comments, Media List** and **Timeline**. **Fig. 1** shows how the new **Badge** Template Component is configured in Page Designer.

Classic Report Templates allowed to achieve pretty similar results in the past – however, Template Components are way more flexible and powerful: For instance, Template Components do not require the region data source to return specific column names – the actual column names are rather *mapped* to the Template Component Parameters. Template Directives, Actions and Menus allow to build much richer UI than a classic report templates ever supported. **Fig. 2** shows how the result of the **Badge** Template Component looks like.

As stated above, Template Components are technically Plug-Ins. Thus, Template Components are developed and maintained within **Shared Components, Plug-Ins**. **Fig. 3** shows, how the definition of a Template Component Plug-In looks like: the most important part is the **Partial**, which contains the actual HTML template, including template directives. All placeholders within the partial template are mapped to Custom Plug-In Attributes, which define labels, help texts, attribute types and other properties. These will show up in Page Designer as region attributes.

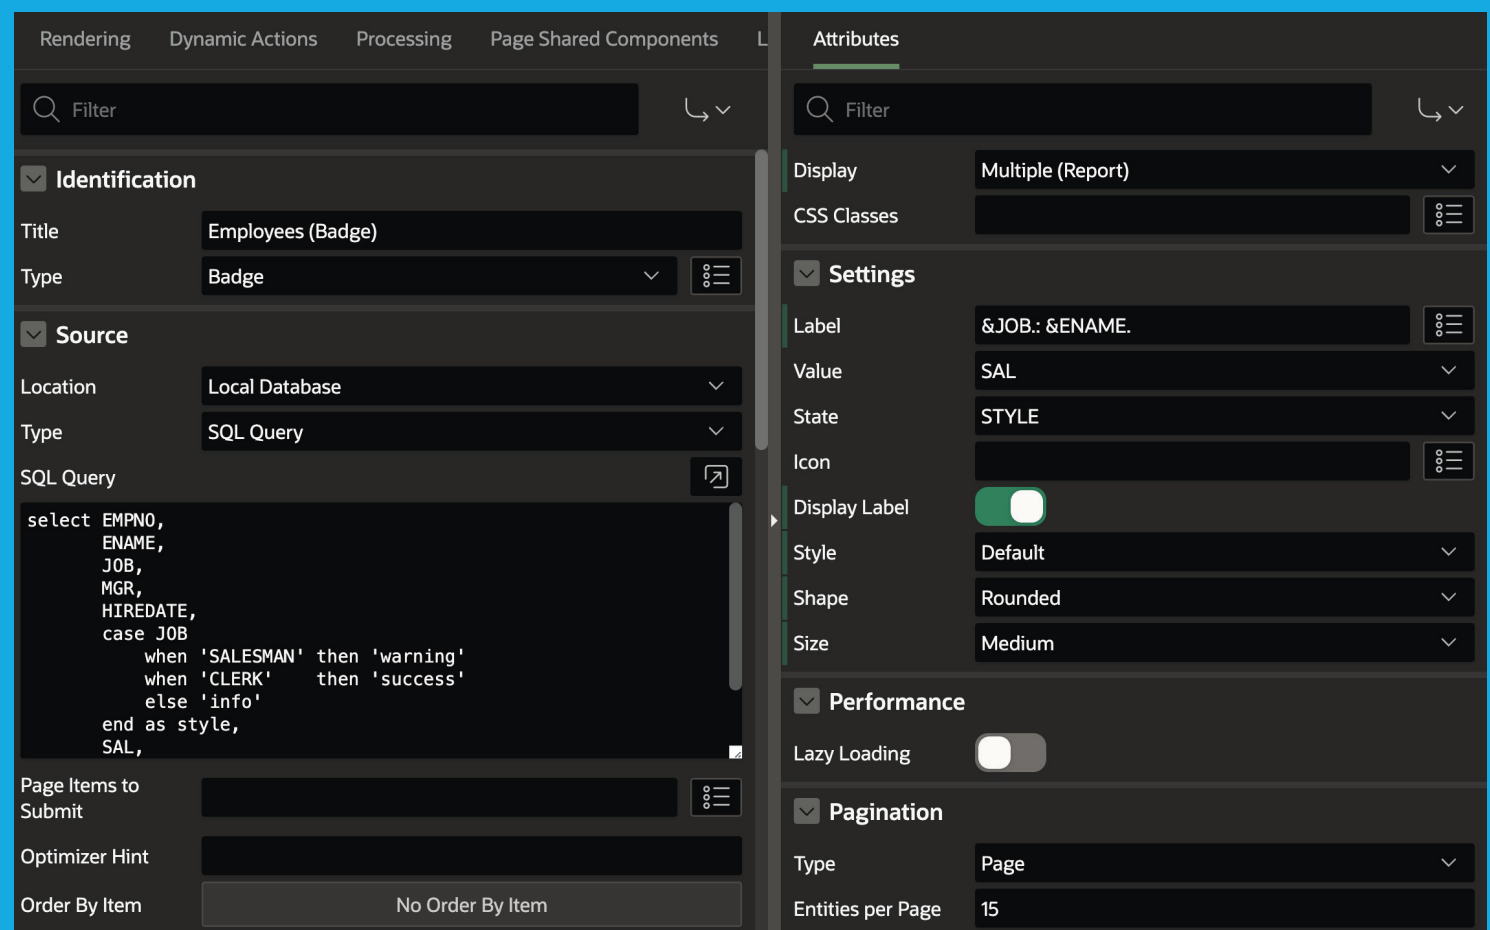


Fig. 1: Configure the “Badge” Template Component in Page Designer.

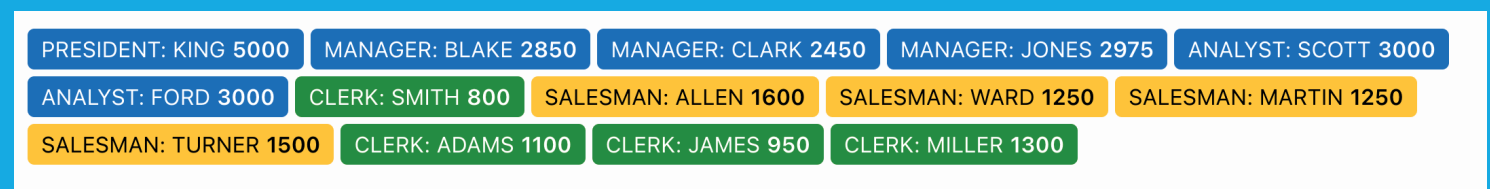


Fig. 2: The Badge Template Component in Action.

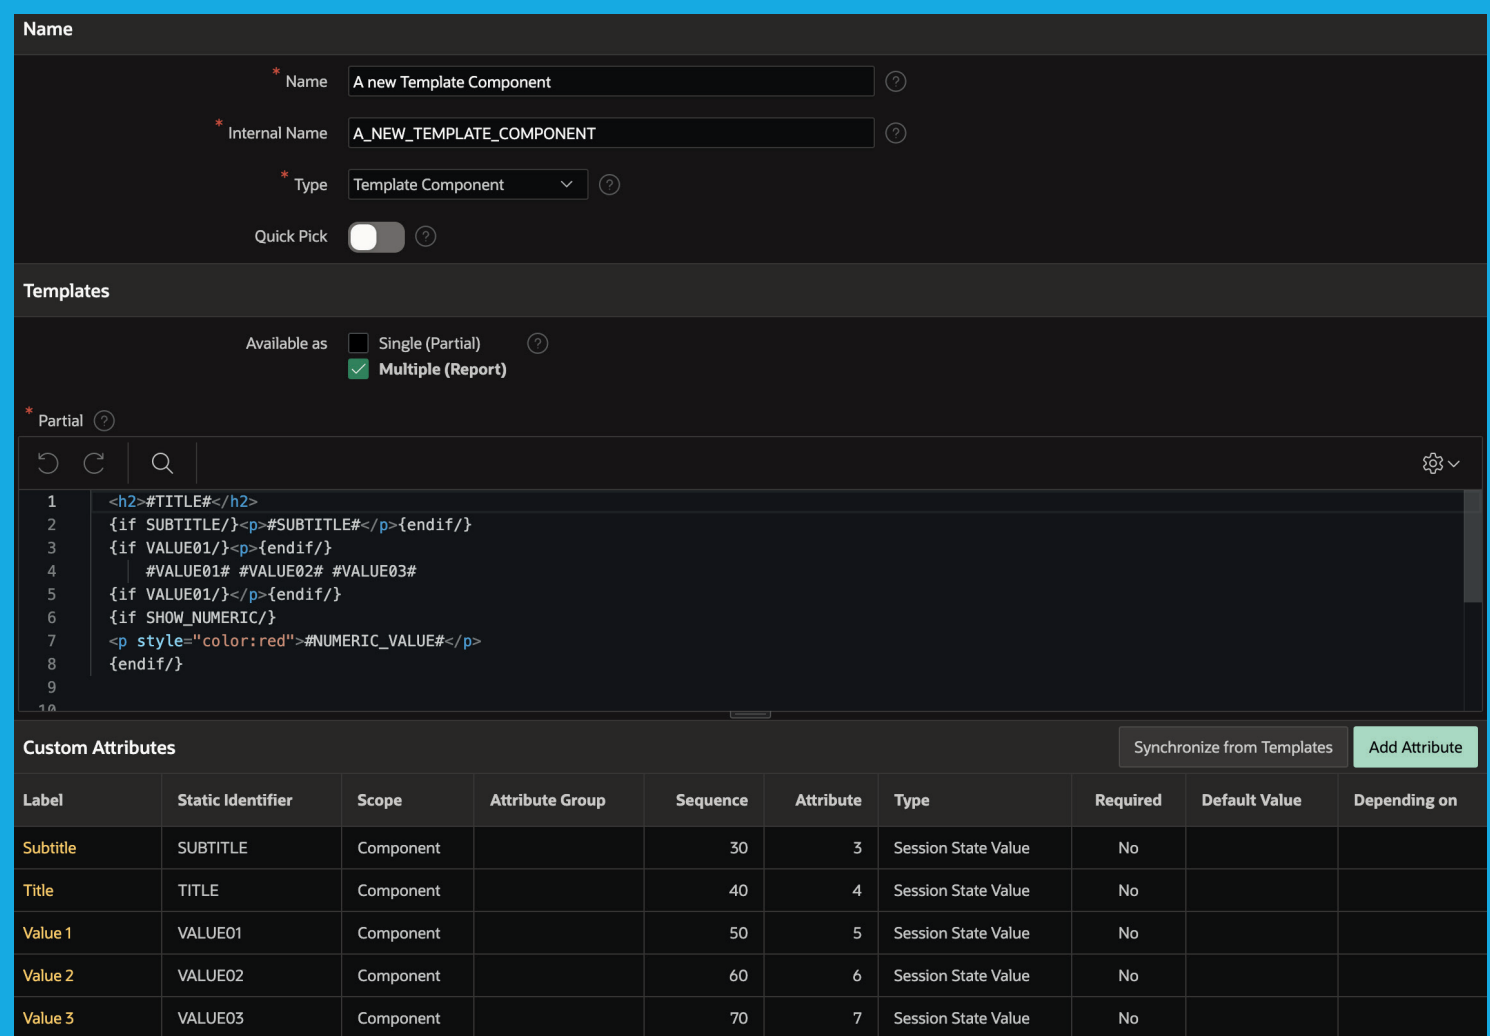


Fig. 3: Defining a new Template Component as a Plug-In.

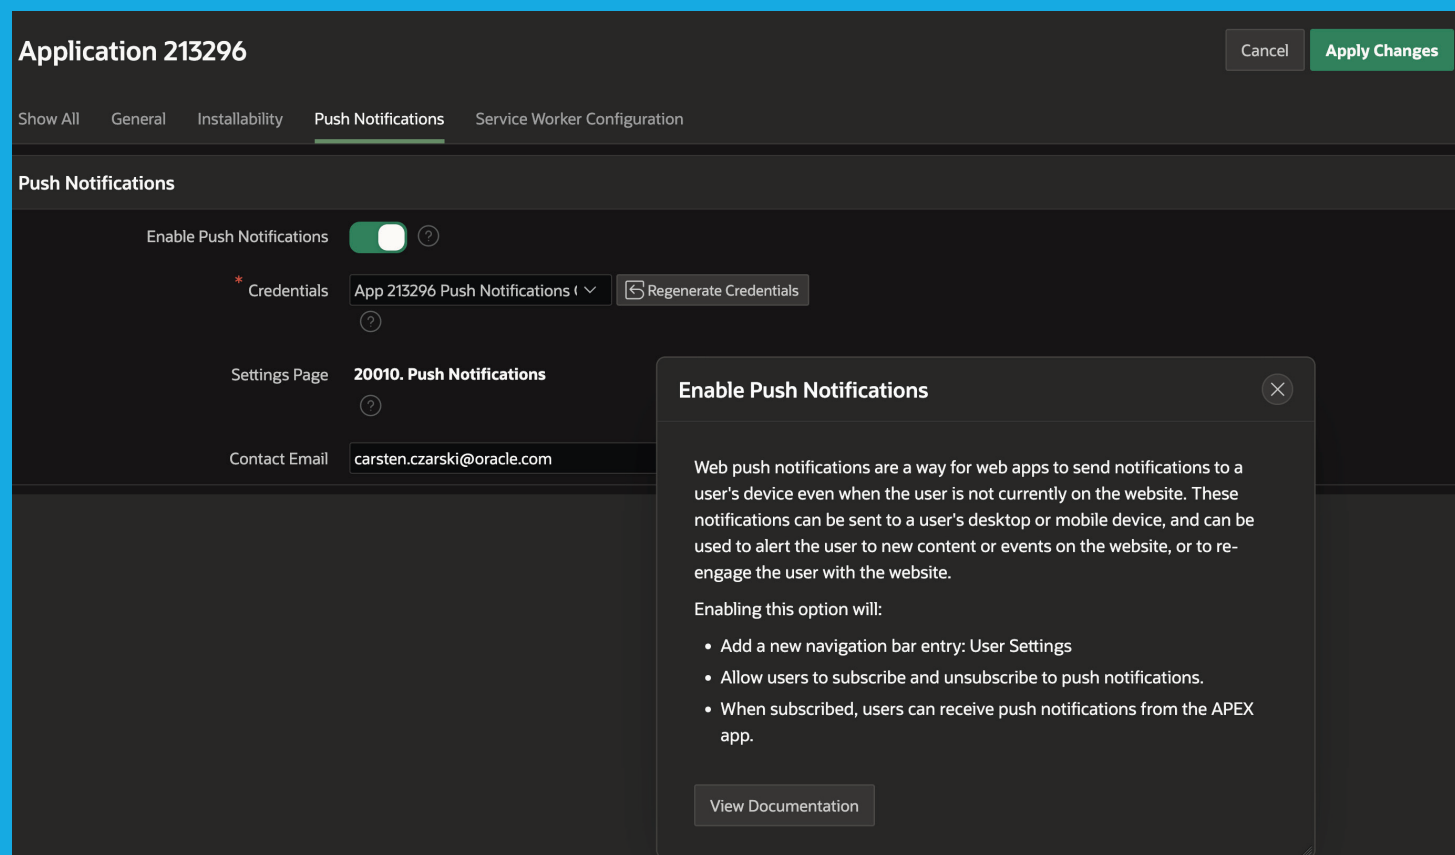


Fig. 4: Enabling Push Notifications for an existing application.

Template Components do not contain PL/SQL code, so building Template Components is way easier than building traditional Plug-Ins. The most important bit is the Partial, and the Custom Plug-In attributes, which act as template parameters. Once the Template Component is ready, it can be used on any APEX page, with any data source.

PWA Push Notifications

APEX 23.1 makes sending *Push Notifications* to end user devices super-easy. All infrastructure is provided by the APEX instance - developers can just enable and use the feature in a declarative manner. For Push Notifications to work, the application must be enabled as a PWA (Progressive Web App). When creating a new application, this is done by simply clicking a checkbox in the Create App wizard. the APEX wizard. For an existing application, this can be done in

Application Attributes, Progressive Web App (Fig. 4).

After toggling the **Enable Push Notification** switch, it's required to set up *credentials* for secure notification delivery to the provider notification endpoints (Google, Apple, and others). APEX can auto-generate these credentials on button click. Notifications are sent by the APEX instance, so the APEX database must be able to connect to the provider notification endpoints on the internet. On the public apex.oracle.com instance, or on Oracle Autonomous Database, this is already set up properly, and things will just work.

Also, the application needs a *Set Up Page* where end users can enable or disable Push Notifications for their very devices. APEX adds this Settings Page simply by clicking a button. Fig. 5 shows this Settings page in action.

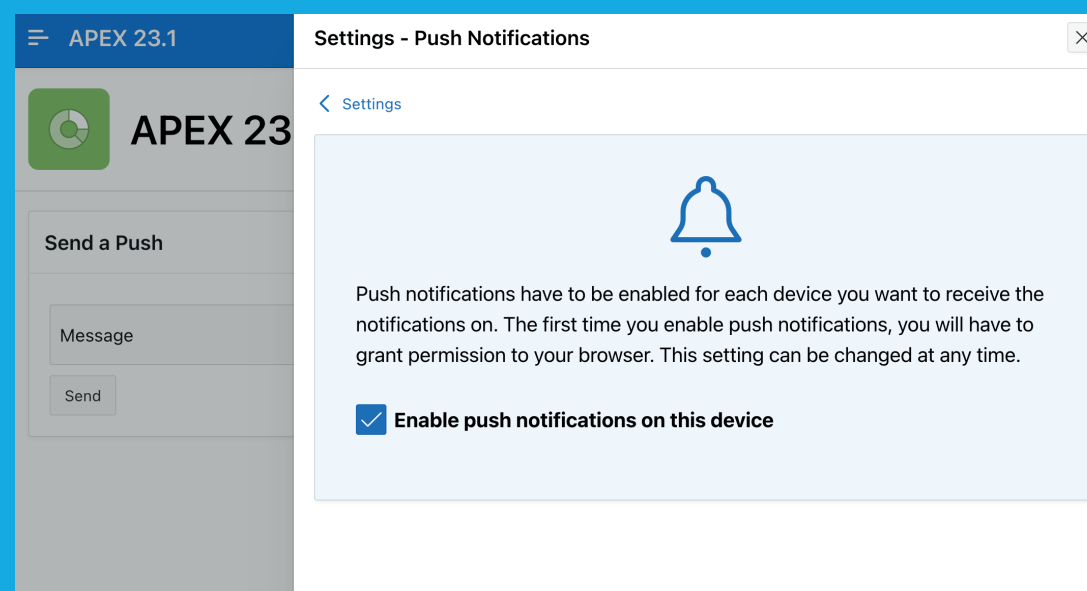


Fig. 5: End users can enable or disable Push Notifications for their device.

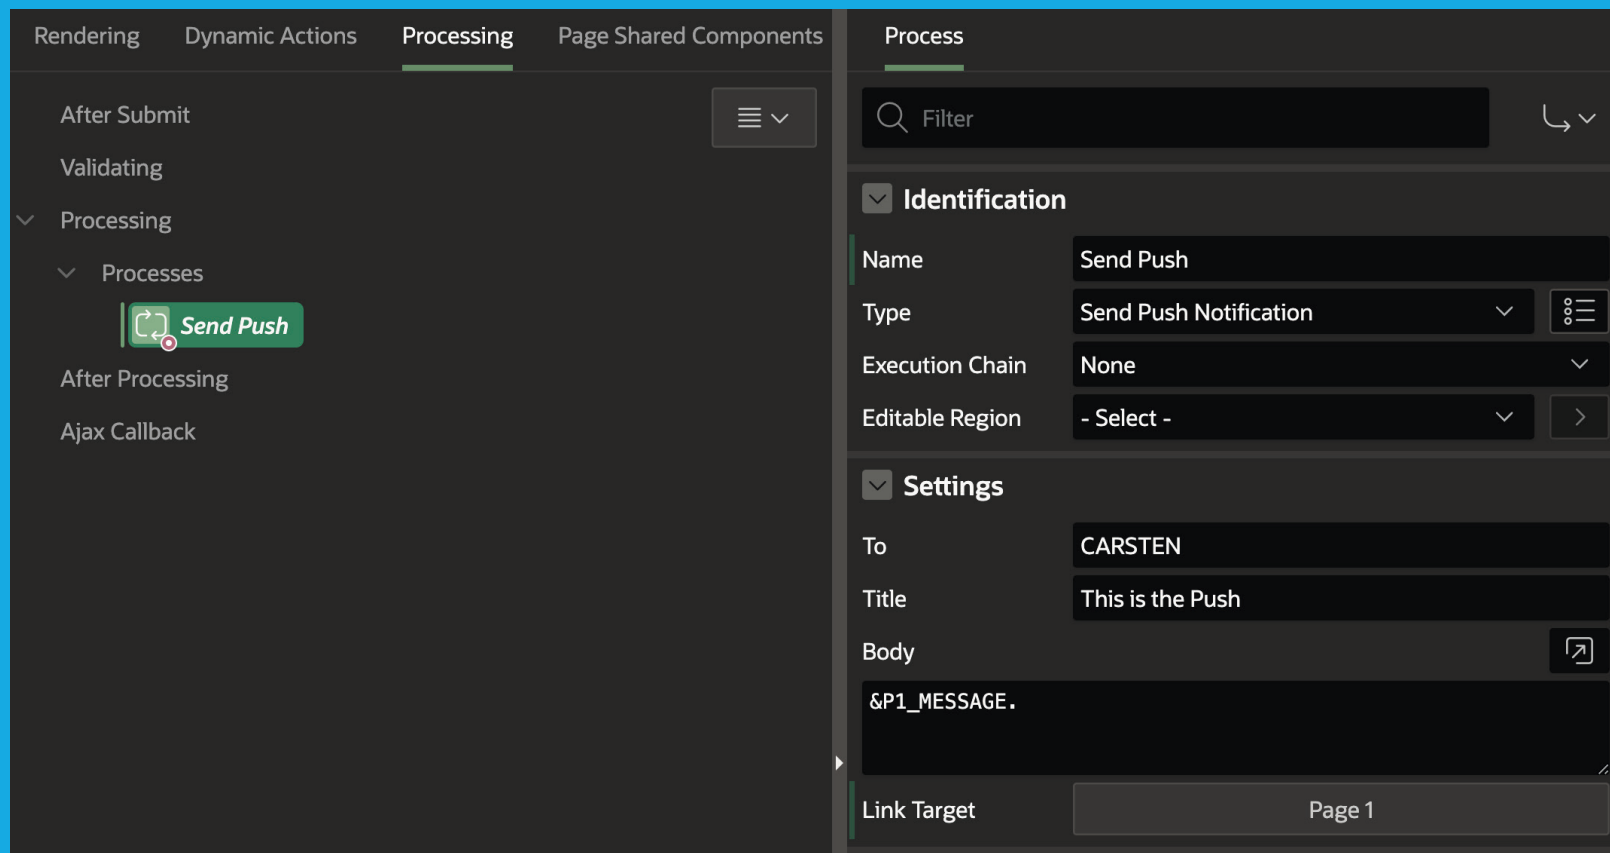


Fig. 6: Declarative "Send Push Notification" Process Type.

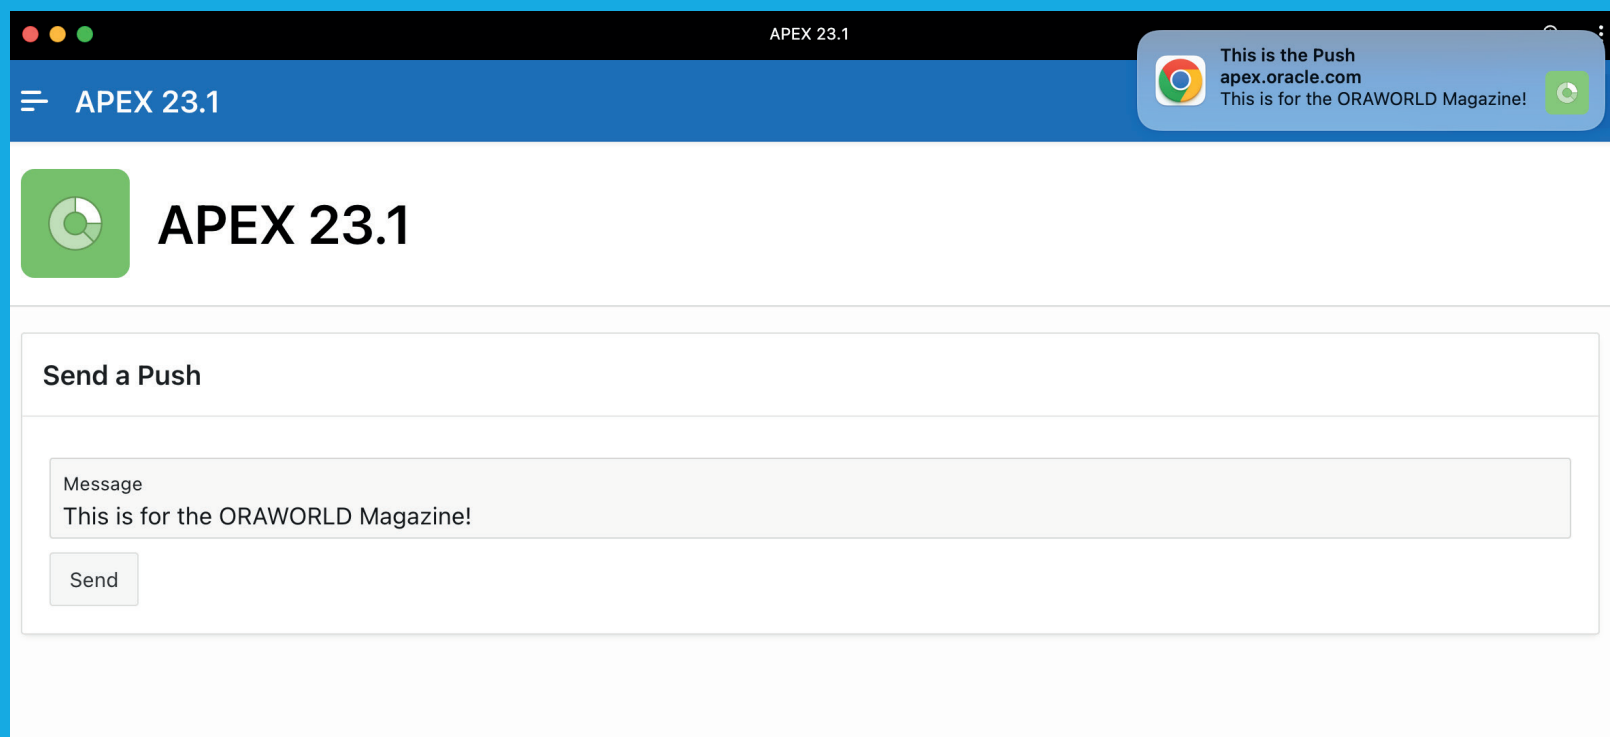


Fig. 7: A push notification has just arrived.

Once that is all in place, developers can push out notifications from everywhere in the application. The most straightforward way is to use the declarative **Send Push Notification** Page Process, which allows to simply configure the *application user* to send notifications to, the title, and the body of the message. Of course, Page Item substitutions are supported. Fig. 6 shows how this looks in Page Designer, Fig. 7 shows the running APEX application, together with a received Push Notification in the Chrome browser on Mac OS.

With the **APEX_PWA** package, developers can also send out push notifications from within PL/SQL. The package also allows to programmatically subscribe or unsubscribe devices or users from push notifications.

Background Page Processing

Page Processes are used by every APEX developer, to perform data transformation, business, or technical logic within an application. Page processes are typically executed when the end user submits the page by clicking a button. The browser will wait until processing completes and is then redirected to the next page within the application flow.

If this waiting time remains within a few seconds, things are fine. However, if processing requires more time, for instance a few minutes, things become problematic: End users only see the browser's loading indicator all the time and don't know whether this is expected, whether there is a networking issue, or a severe error occurred. And if web server timeout thresholds are exceeded, processing is even terminated with very technical *Request Timed Out* messages.

APEX 23.1 makes it super-easy to send execution of long-running processes to the background, so that the end user

sees the browser “coming back” immediately. Background execution is now a *declarative attribute* - no manual coding with the DBMS_SCHEDULER package is needed at all.

The new **Execution Chain** page process type acts as a *group for page processes*. Processes within an execution chain are executed one after the other, and the **Execute in Background** attribute controls whether processing is done in the foreground or sent to the background.

Fig. 8 shows, how an execution chain to run in the background is configured in Page Designer. The process sequence is to run some longer running code first, then use the **Push Notification** process type to send a success message to the end user. Most declarative process types can be executed in the background.

When the page is processed, the background execution task will be submitted to an internal execution queue, and processing does not block. The APEX engine will process this execution queue asynchronously in the background. The **Return Execution ID into Item** allows the developer to retrieve the ID of the background execution, which is important to query its status later.

The **APEX_APPL_PAGE_BG_PROC_STATUS** view contains information about background executions and their current

state and progress. The process code, which runs in the background, can use the **APEX_BACKGROUND_PROCESS** package to report its status and progress to APEX.

The **Developer Toolbar, Session** page has been extended to show information about background executions as well (**Fig. 9**). Clicking the **Abort** button terminates a background execution. This is useful for developers to kill runaway background executions, for instance accidentally submitted endless loops.

Background executions can also be terminated with the **ABORT** procedure in **APEX_BACKGROUND_PROCESS** package, as well as within workspace administration or in the INTERNAL workspace, for APEX instance administrators.

Invoke API for REST

The Invoke API page process type was introduced in APEX 22.2 and allows to declaratively call a PL/SQL function or procedure in a page process. All the mappings of page items or SQL expressions to procedure parameters are done declaratively in Page Designer.

APEX 23.1 extends that process type to support *REST Data Sources*. After picking a *REST Data Source*, and an *Operation*, REST parameters are being mapped to page items or SQL Expressions, like working with a PL/SQL procedure (**Fig. 10**).

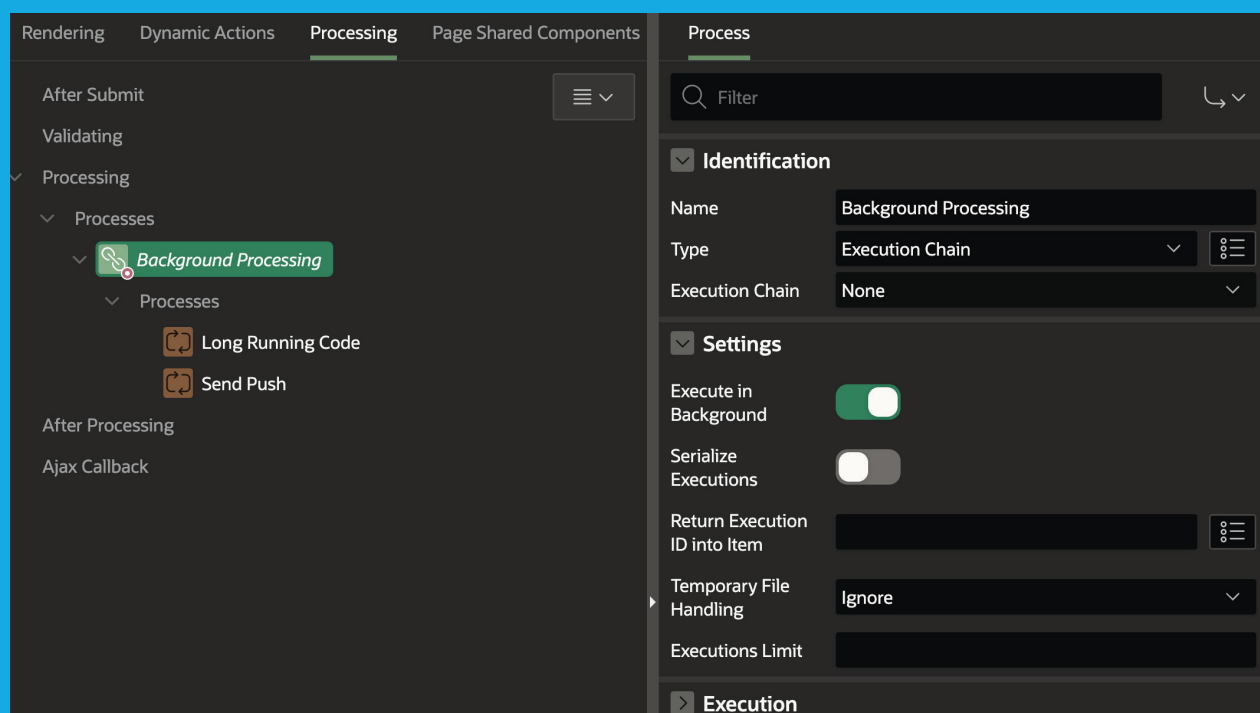


Fig. 8: Configuring an Execution Chain to run in the background.

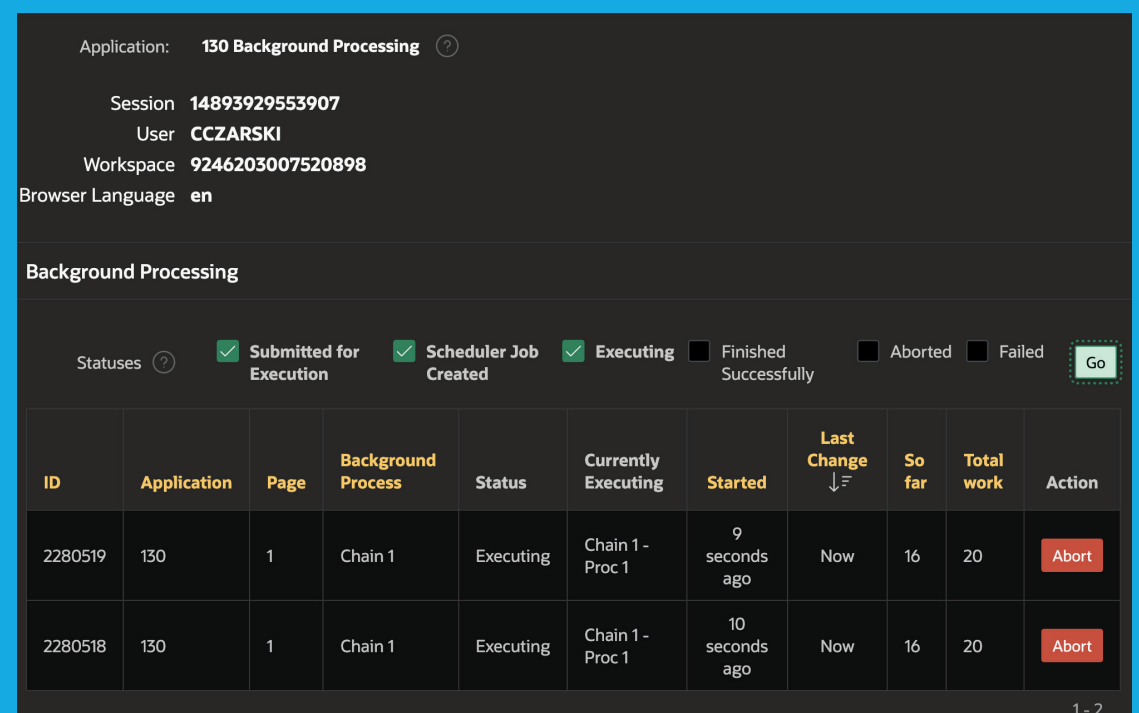


Fig. 9: The Developer Toolbar, Session page allows to manage background executions.

When the Invoke API process is executed, the APEX engine will invoke the REST operation with the correct parameter values.

SQL Workshop and Application Builder

SQL Workshop Object Browser (Fig. 11) is totally revamped and modernized in APEX 23.1. The whole developer experience is improved, all the wizards to create new objects were cleaned up and redesigned to support the most common tasks in an optimal manner.

APEX Application Builder was extended with several smaller improvements, including:

- The Create Page wizard now allows to copy a page as well. No need to navigate to Page Designer first.
- The code editors in Page Designer now have a **Save and Run Page** button, which immediately saves all changes and runs the page. That will save quite a few mouse clicks.
- When running on Oracle Database 23c, Page Designer will automatically allow to create APEX components on top of **Property Graphs**. Property Graphs are a new feature in Oracle Database 23c.
- The Application Builder now also uses friendly URLs.

Other new Features

The *Administrator Digest*, a dashboard summarizing activity on the APEX instance, was added to the **Monitor Activity** area in the **INTERNAL** workspace (Fig. 12).

The new **APEX_APPLICATION_ADMIN** package allows to programmatically set application attributes like Availability Status, Application Version, Global Notification, or others.

The *TinyMCE* JavaScript Library has been added and is now the base for the Rich Text Editor as CKEditor has been deprecated.

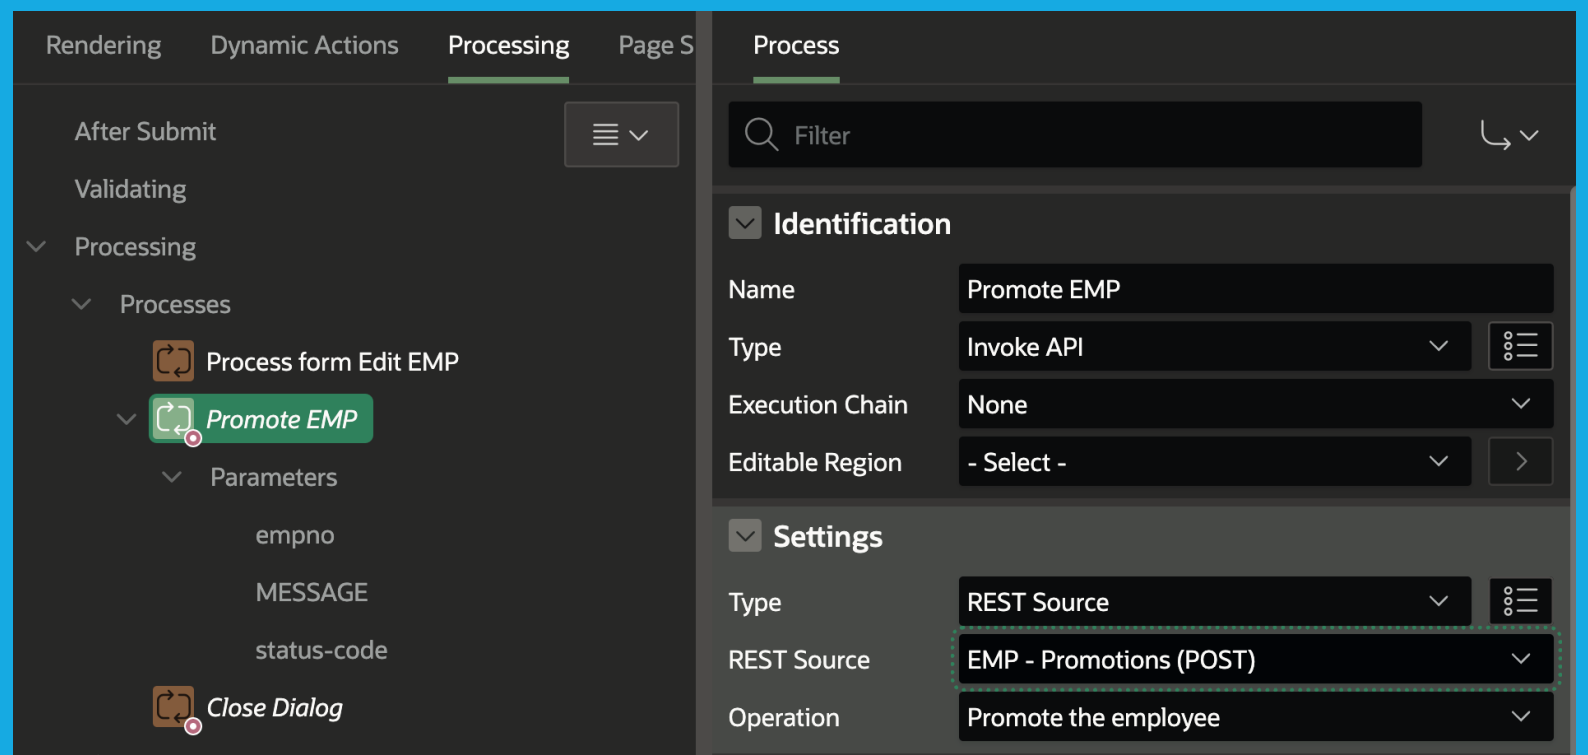


Fig. 10: The Invoke API process type now supports REST Data Sources.

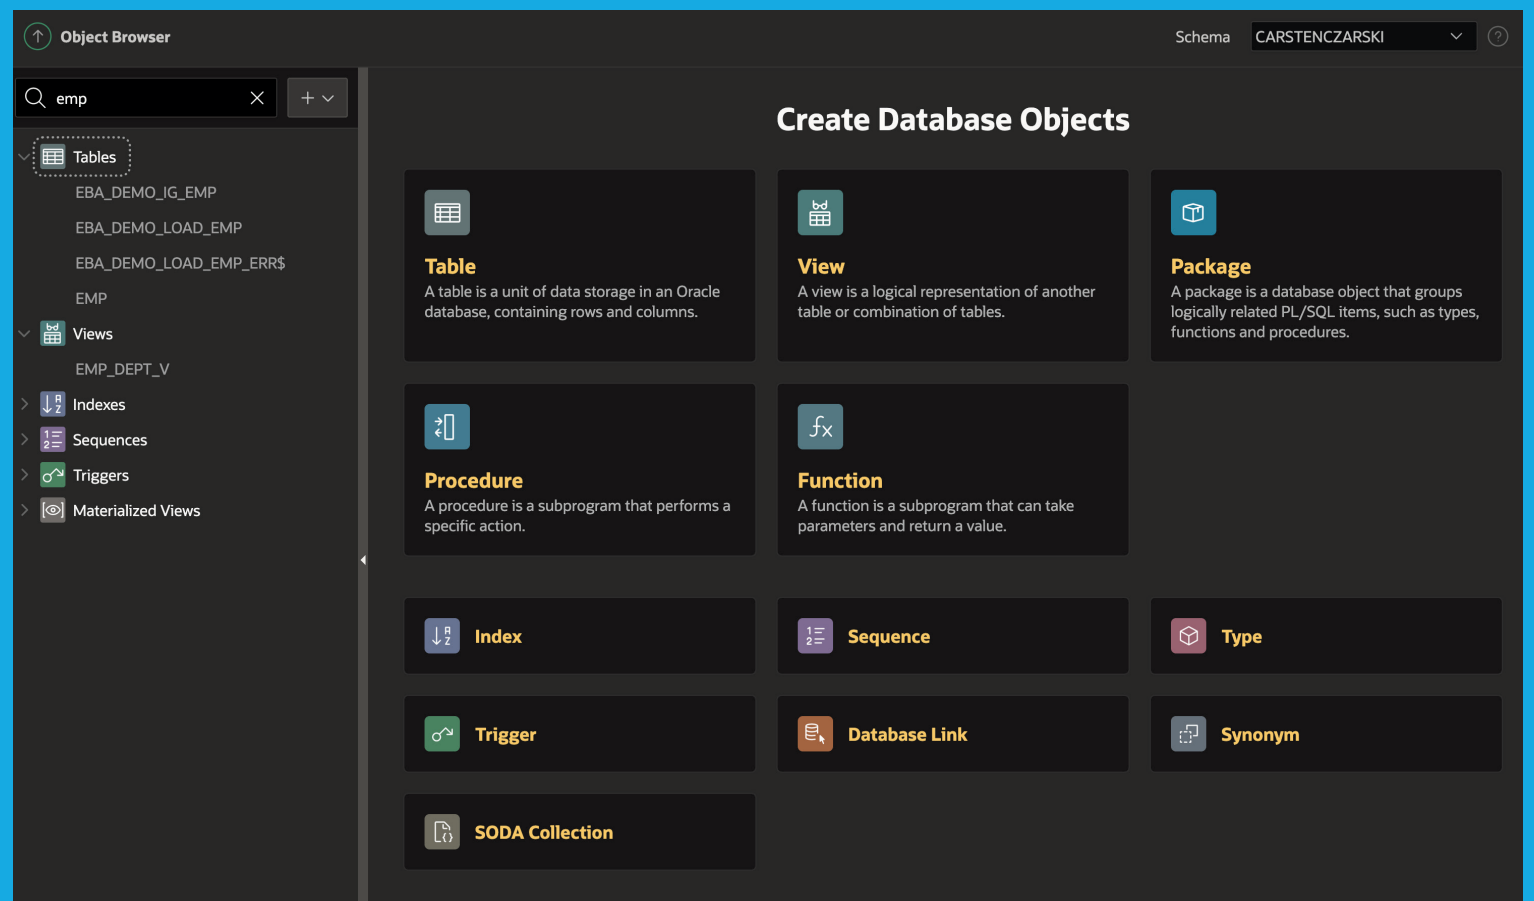


Fig. 11: New Object Browser in APEX 23.1 SQL Workshop.

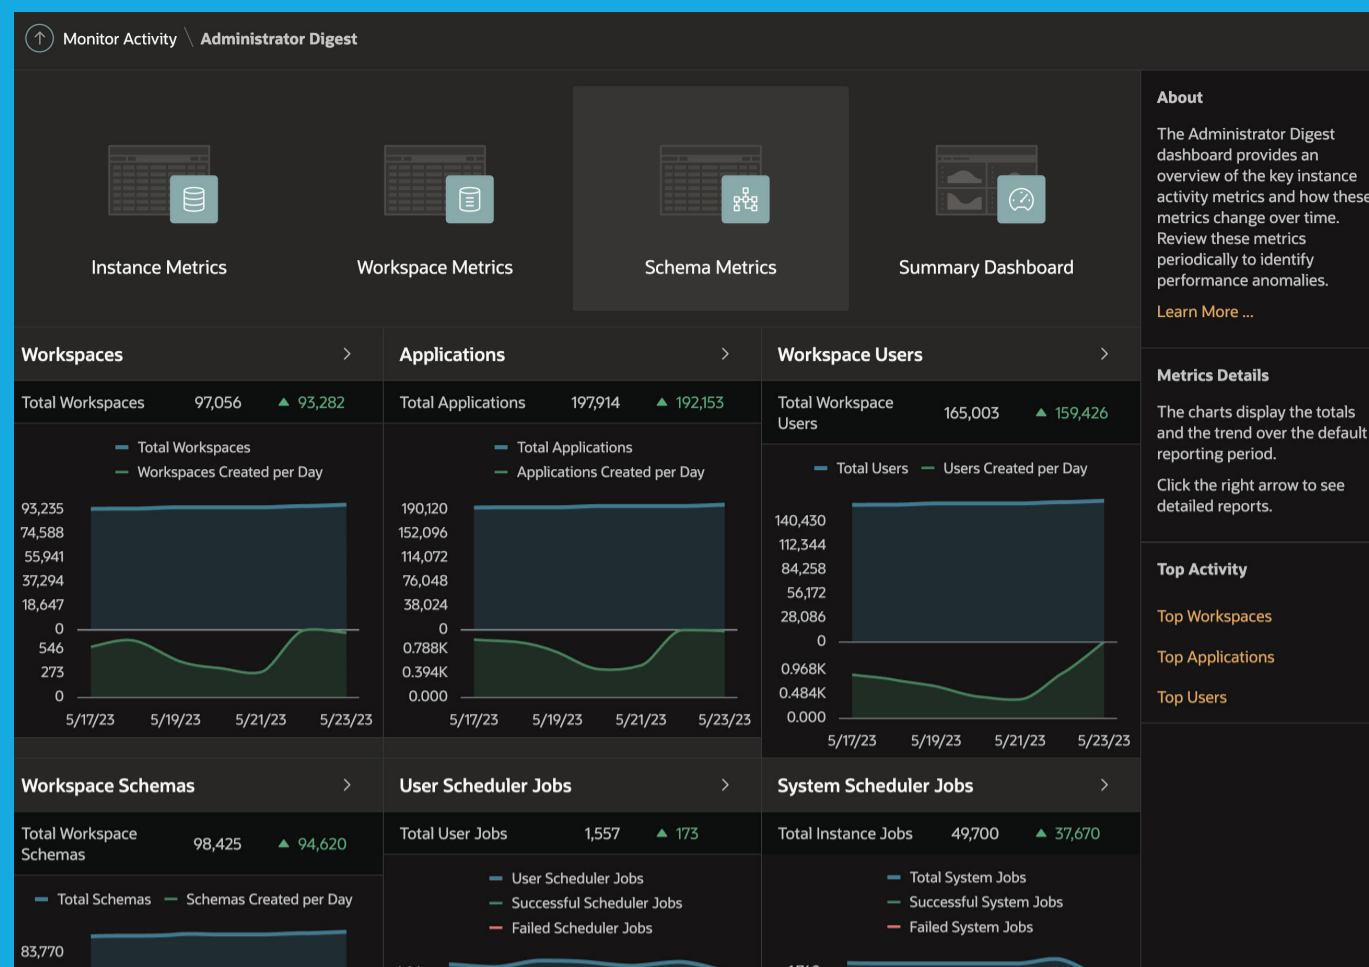


Fig. 12: APEX Administrator Digest

Besides that, several JavaScript libraries were upgraded to newer releases including Oracle JET 14.0.0, FullCalendar 5.11.3, PrismJS 1.29.0, MarkedJS 4.2.5, DOMPurify 2.4.3, CKEditor5 36.0.0, Terser 5.16.1, CSSO 5.0.5, Cropper.js 1.5.13, and MapLibre 2.4.0.

Summary

APEX 23.1 contains several new features and enhancements. The marquee features are Template Components and Push Notifications. Template Components are a big leap towards more reusable UI components, without having to code PL/SQL for a region plug-in. Push Notifications further complement PWA capabilities in APEX – and fit perfectly well to the new background execution for page processes: Once the long-running process has completed in the background, simply push a message to the end user’s device.

A lot of work went into the modernization of the aged object browser. Developers now find a modern and accessible environment to maintain tables, views, and other database objects. Numerous smaller enhancements to the Application builder will help developers to be more productive.

The best way to find out about everything is to sign up for a workspace on apex.oracle.com, or to get a subscription to Oracle Cloud “Always Free” Services, configure the APEX workspace and get rolling.

More Information

- Information and environment to test-drive APEX: <http://apex.oracle.com/en>

- Overview: What’s new in APEX 23.1 <https://apex.oracle.com/en/platform/features/whats-new-231/>
- Download APEX 23.1 <https://www.oracle.com/tools/downloads/apex-downloads/>



About Carsten Czarski

Carsten works for Oracle in Germany since 2001. He started in the Presales organization helping customers and partners regarding database-centric application development. Since March 2016, Carsten is a member of the Application Express development team. Focus of his work is on the support for REST services in APEX, Data Loading, SQL Execution and a lot of other areas. He is a frequent speaker at international user group conferences.



Mohit Jain and Megha Gupta

Powerful Intercompany Eliminations of Oracle EPM Financial Consolidation and Close Services (FCCS)



Why let Intercompany Eliminations increase Corporate Risk

While there is some indication that Intercompany Elimination, if not happening in right manner, can increase the risk over time, there continues to be noticeable room for improving how organizations should identify, manage, use technology and keep their eyes on risks that may arise and extensively impact their ability to achieve strategic goals.

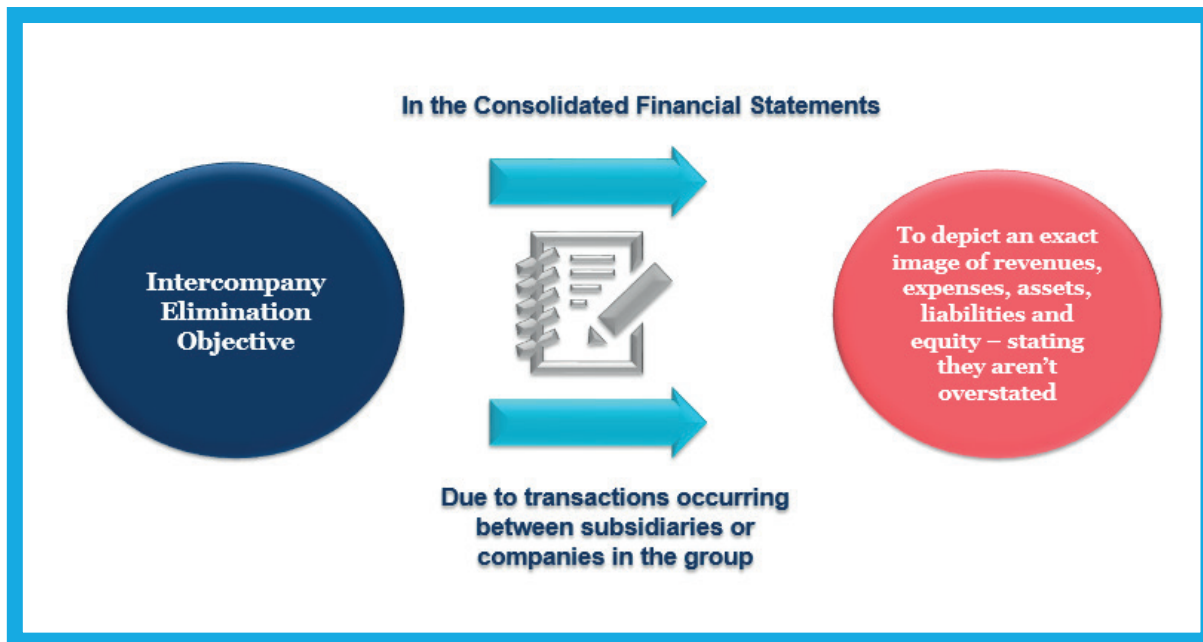


Fig. 1: Intercompany Elimination Objective

When IC entities are used, users usually lose overall visibility into the individual intercompany balances and see the total numbers in the IC entities only. It's a **high-risk activity** and becomes very difficult to understand where, the out of balance condition resides, and what needs to be addressed and by what party!

Key Observations/Gaps observed in various Organizations

In many organizations comprehensive IC Policy detailing IC events/transactions type, timelines etc. are not available. In absence of policy guidelines around cost & revenue accruals for intercompany transactions, some of pass-through accounts are not reconciled for a long time and results in significant mismatch between IC partners across Business Units and regions. Intercompany Invoices are posted in silos across Business Units which leads a limited alignment between the IC Partners regarding month end accrual that leads to month end mismatches. Platforms/infrastructure which are used to align with upstream processes for invoicing are not adequate and intercompany accruals are posted by both the parties manually as one-sided entry and that results in transaction link breakage.

Missing Controls in IC observed in Organizations

It is essential to have some controls on intercompany transactions within organizations. Many organizations are still defining elimination rules manually in excel or ERP. The training and process documentation is not available across IC activities. There is no resolution mechanism or platform to report, manage, and resolve the IC transactions. For IC processes, the adequate reporting/dashboard, and detailed responsibilities metrics - across the upstream and downstream sub-processes does not exist. If there are no such controls, it could affect the organization performance. There will not be any control for timely closure of disputes, escalations, and review of open items. Business will not be able to track the aging on IC mismatch reports and even drill down/transactional level details. This inconsistent usage of trading partners while posting accruals leads to duplicate adjustments at consolidation level.



Multinationals need to treat their internal business with as much rigor and control as they treat their external business.

Intercompany accounting issues aren't limited to large multinationals. I've seen companies with 10 or fewer legal entities that have major problems.

- Kyle Cheney, CPA, a partner in Deloitte's advisory practice who focuses on governance of accounting processes and activities.



Where we've come from and where we're going

Enterprise Performance Management (EPM) is more understood and more valued today than it was 10 years ago. Every global company should now address the mess of intercompany eliminations – with technology. This will lighten the growing pains of acquisitions as well as global expansions. As today companies have expanded their global footprints and generating an enormous volume of intercompany transactions. Operating in multiple countries, exposes companies to greater risk if they fail to streamline their intercompany transactions, even the smooth month end close depends on this. The challenges of performing efficient intercompany elimination are substantial.

How Intercompany Eliminations mess can be cleaned up effortlessly with Oracle Financial Consolidation and Close Cloud Services (FCCS)

Financial Consolidation and Close cloud services (FCCS) is a Cloud-based 'Software-as-a-Service (SaaS)' delivery with a Tablet user interface and it offers powerful Inter Company eliminations.

Fig. 2 shows an organization hierarchy of company named "G Tech". It shows how different base level company codes are aggregated to sub parents and sub parents are rolling to top parent "G Tech Worldwide before Acquisition".

In this organization, there are intercompany transactions among different companies.

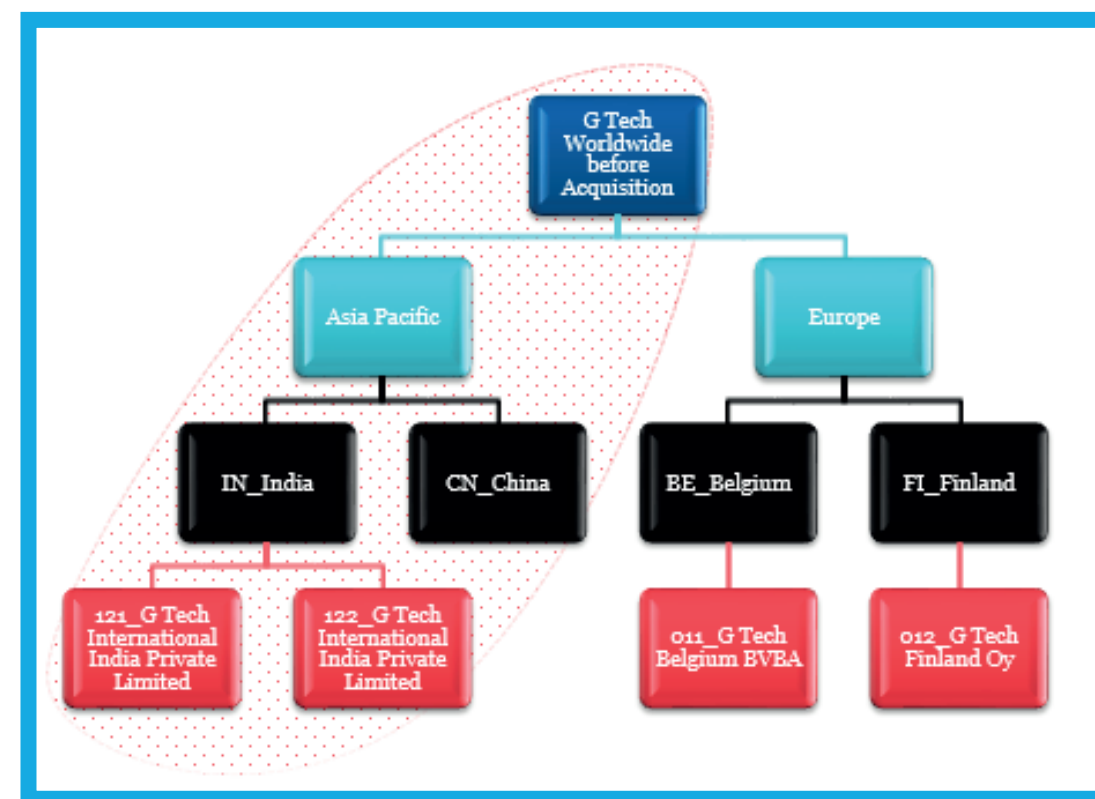


Figure 2: Organization hierarchy of a company named "G Tech".



Intercompany (IC) eliminations **automatically occur at every first common parent** in every alternate hierarchy without complex rules.

More than half of the mess of IC eliminations is resolved with its design itself. FCCS provides a "Intercompany dimension", captures all party counterparty balances and provides complete visibility into what is in and out of the balance. This allows business users to identify and correct the discrepancies quickly and easily.

How it Works in FCCS

It's an out of the box functionality of FCCS. Now the question arises, how it works. As mentioned above, **elimination takes place at first common parent**. Here are some examples of from above structure of organization:

- 1st IC elimination takes place between "**121_G Tech**" and "**122_G Tech**" at **parent "IN_India"** in the currency of "IN_India"
- 2nd Elimination takes place between member "**IN_India**" and "**CN_Chain**" at **parent "Asia Pacific"** in the currency of "Asia Pacific". Note: "IN_India" includes its children "121_G Tech" and "122_G Tech"
- Elimination takes place between member "**Asia Pacific**" and "**Europe**" will takes place at parent "**G Tech Worldwide before Acquisition**" in the currency of parent "G Tech Worldwide before Acquisition."

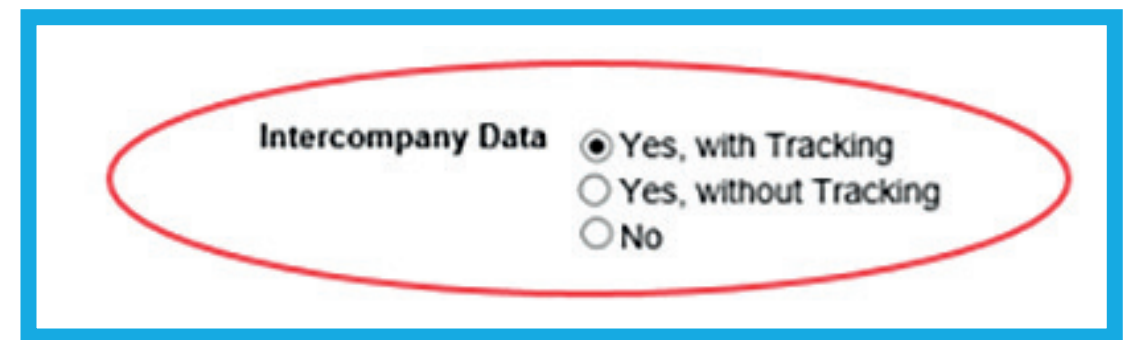
Setups in FCCS

To enable the intercompany elimination, FCCS administrator needs to enable some settings. These are one time setting defined in application. Business user does not need to perform any extra activities for IC elimination. Below are the

details of these steps -

Settings at Application Level

At the time of creation of FCCS application, Admin should select one of the following options to enable IC eliminations functionality:



- Select 'Yes with Tracking' or 'Yes, without tracking' to enable and
- 'No' in case you do not want to enable IC eliminations.

Here we are using '**Yes, with Tracking feature**'.

Once intercompany tracking has been enabled, FCCS will create a system defined dimension named "**Intercompany Dimension**". This dimension has 2 default members as "**FCCS_No Intercompany**" and "**FCCS_Intercompany Entities**". All trading partner entities will be children of "**FCCS_Intercompany Entities**".

Moreover, below settings need to be defined by administrator:

1. **Intercompany Entity** is created with "**ICP_Entity_Yes**" as default member and this will be used to enable Entities

as IC Entities. This dimension has a single member “ICP_Entity_Yes”.

2. Intercompany Account is created with “IC_Acc_Yes” as a default member and will be used to enable accounts as IC accounts. This dimension has a single member “ICP_Account_Yes”.

3. A Plug Account If there are any differences between the intercompany account balances in FCCS, these differences in intercompany accounts are stored at PLUG accounts as per organization policies.

“Plug Account attribute” will have list of custom plug accounts based on organization policy. Admin can also use a flat file to load plug accounts in FCCS.

4. Is Plug Account

FCCS provides a predefined member as “Plug_Acc_Yes”. Assign this member to account to enable account as Plug account. This dimension has a single member “Plug_Acc_Yes”.

These setting can be divided as settings at **Entity and Account dimensions** level as below.

Settings at Entity Level

Admin should define IC settings in Entity Dimension to incorporate all trading partners in FCCS intercompany entities.

Enable entity as Intercompany Entity

Once any entity member will be assigned “ICP_Entity_Yes” as attribute value, that entity member will be added into “Intercompany(IC) Dimension” with Prefix of “ICP_”. For example - entity “121_G Tech” will be assigned to “ICP_

Entity_Yes” as attribute value, below member will be added into Intercompany dimension - “ICP_121_G Tech”.

If entity member is **not assigned** “ICP_Entity_Yes” as attribute value, that entity member **will not participate in intercompany transaction** and will not be added into “Intercompany (IC) Dimension”.

Settings at Account Level

The admin should define below settings in Account Dimension to enable intercompany elimination:

- Enable the account as Intercompany account.
- Define the Plug Account
- Enable the Plug Account
- Assign the Plug Account to IC Account

Enable the account as Intercompany account

To enable an account as Intercompany account the admin must assign default property “IC_Acc_Yes” and assign “Plug account”.

If any account member has not been assigned “ICP_Acc_Yes” as attribute value, that account can't **participate in intercompany transactions**.

Define the Plug Accounts

If there are any differences between the intercompany account balances in FCCS, these differences in intercompany accounts are stored at PLUG accounts as per organization policies. Plug account members are created in “Plug Account” dimension.

Enable Plug Account

Once, plug account has been added into “Plug Account” dimension, it is mandatory to assign default property “Plug_Acc_Yes” to the account to enable it as Plug Account.



Assign Plug Account to IC Account

Once, Plug Account has been enabled, it can be added to the relevant Intercompany Account to define where delta is supposed to move during the intercompany elimination.

Supporting Dimensions for IC Elimination

Along with Account and Entity dimensions - 'Data Source' and 'Consolidation' dimensions are two additional prebuilt dimensions available in FCCS which play an important role for IC Elimination to work.

Data Source dimension - It is a prebuilt dimension available in FCCS. This dimension provides information about the source of data available in FCCS.

- Member "FCCS_Intercompany Elimination" is used to see eliminated data.

Consolidation dimension - It is a prebuilt dimension available in FCCS. This dimension provides complete audit trail of data. Like, what data was loaded, what is translated value, what will be the eliminated amount, consolidated data after elimination and other details.

- Member "FCCS_Elimination" is used to see eliminated data at first common parent.
- Member "FCCS_Contribution" is used to view what data will aggregate from a child to its parent.

Quick retrieval of data using Smart View (Excel Add-ins)

FCCS provides quick retrieval of data using Smart View (Excel Add-ins) functionality (Fig. 3). You can drill and view your eliminations details using the same.

Entity	IC Entity	Account	Currency	Entity Currency	Parent Currency	Entity Currency
121_G Tech International India P	ICP_122_G Tech India Private Limited	AC_2012	FCCS_Entity Input	5,209,150,120.15	5,209,150,120.15	0
122_G Tech India Private Limited	ICP_121_G Tech International India Private Limited	AC_2012	FCCS_Entity Input	(3,158,462,552.36)	(3,158,462,552.36)	0
IN_India	FCCS_No Intercompany	Intercompany RecPay Diff	FCCS_Contribution	0	0	8,367,612,672.51
IN_India	FCCS_No Intercompany	Intercompany RecPay Diff	FCCS_Elimination	0	0	-
IN_India	FCCS_No Intercompany	Intercompany RecPay Diff	FCCS_Proportion	0	0	8,367,612,672.51
IN_India	FCCS_No Intercompany	Intercompany RecPay Diff	FCCS_Entity Total	0	0	8,367,612,672.51
IN_India	FCCS_No Intercompany	Intercompany RecPay Diff	FCCS_Entity Input	0	0	-
IN_India	FCCS_No Intercompany	Intercompany RecPay Diff	FCCS_Entity Consolid.	0	0	8,367,612,672.51

Fig. 3: Example of Intercompany Elimination

Conclusion - The powerful IC elimination features provided by FCCS!

FCCS provides a single version of truth to track intercompany elimination across different geographies and business units. It provides complete audit trail of Entered, Calculated, translated, eliminated, and consolidated data.



About Mohit Jain

Mohit is an Oracle ACE Pro and ODTUG EPM community Lead with 15+ years of finance & IT experience and a focus on the Oracle EPM technology stack and Business Intelligence.

He specializes in architecting, developing, and training of EPM set of tools for Controllershship and FP&A applications, primarily in Oracle EPM Cloud tools (ARCS, FCCS, PBCS), Essbase, Planning, FDMEE, DRM and Reporting. He is a Hyperion certified consultant, an avid reader, and data enthusiast with a passion for open source, Machine Learning and AI ecosystem.

Mohit was a presenter at ODTUG Kscope21, Kscope20, UKOUG20, AUSOUG22, AUSOUG21 and INNOVATE 21 and has presented webinars and written articles for ODTUG on a couple of topics. In his free time, he loves to spend time with his wife and 2 kids.

Customized reports/dashboards are available to view IC process. Users can track IC transactions for any historical period. It requires minimum IT support, there is no need to write any formula for IC elimination and no requirement of accountant intervention.



About Megha Gupta

Megha is an Oracle ACE Pro, ODTUG EPM community Lead, having 16 years of experience, with proven mixture of business intelligence and technical knowledge on consolidation and Reporting tools and Oracle EBS to successfully capture requirements from the business perspective and translate those into designs that are efficient, effective, and logical.

Megha was a presenter at ODTUG Kscope21, Kscope20, Kscope19, UKOUG20, UKOUG22, AUSOUG22, AUSOUG21 and INNOVATE 21 and has presented webinars and written articles for ODTUG on a couple of topics. In her free time, she loves to play with her six years old daughter and go for walk.



One Goal, Different Approaches:

Dani Schnider and Martin Berger Performance Tuning in Practice



Dani and Martin have worked together for years to analyze and fix performance problems for a customer. In this work, they have experienced how the common goals were achieved through different ways. In this article they explain some of the similarities and differences, their backgrounds and approaches.

For more than two years we worked together with other colleagues as a team and did performance analyses and optimizations of various Oracle databases and application systems for one customer. The special situation here was that this was not a continuous work for each member of the team, but a sporadic one, at different moments of time: the client had comprehensive coverage, but each team member had no contact with the project for weeks.

In the process, we quickly realized that we often cannot solve specific problems alone but have to hand them over to colleagues due to time or organizational limitations. This made us realize that although all team members are very good specialists with many years of experience, they each have their own technique, approach and methods.

This made us not only address the actual problems in the analyses, but also exchange our respective methods, specific analyses and results. In this exchange, we were all able to learn from each other, expand our knowledge and also our own toolbox. We shared our experiences on some of the topics.

Communication with the customer

Dani: Our colleagues from the 2nd Level Support have a list of standard questions they ask for every ticket when a performance problem is reported (see Fig. 1). Martin, what do you do with the answers to these questions?

Martin: For one thing, it's a good practice from the 2nd level to ask the questions; this shows the customer that someone is working on the ticket. There is nothing worse for the customer than having no one to take care of their problem that is important to them. At the same time, there are several details that are asked that may be important for our analysis.

1. How long it is since you've started experiencing the performance degradation?
2. Does performance degradation affect the entire application or only one sql statement?
3. Please provide timestamp when query worked faster/slower and if it's the case, a database where it works fine.
4. Please make sure to provide database name, username and a sql id or sql statement as well as a timestamp of when it ran.
5. Have you recently changed something on the application side or at the data level?
6. Have you gathered statistics recently? - If not, please run gather stats on schema and let us know how the query performs afterwards

Fig. 1

Many are generally valid, others can be customer specific. A general list of initial questions is certainly useful to establish a baseline.

Dani: I always smile at question no. 5: "What have you changed?" The answer is always that they haven't changed anything. Question 6 about the statistics I think it's good, but in my experience outdated or incomplete statistics are rarely the cause of performance problems. For most clients, the statistics are calculated with the default job and therefore are usually up to date. But a quick check of when the statistics were last updated is certainly a good idea.

Martin: ... and then write in the ticket that we have checked the statistics. In general, it is important to document as

well as possible what we have examined. Most important is to document failures and dead ends. This saves time and prevents colleagues from having to do the same checks again.

Dani: Speaking of communication - I often ask myself the question: "How do I tell the customer our conclusion?". In some tickets I wanted to write: "Your application is crap" - but that doesn't go down so well. How do I tell a developer that his data model or his lack of SQL knowledge is the main reason for the performance problems?

Martin: This ticket system has the advantage of having two levels: one for us and one for the customer. On the level that's visible to us we can and should write quite clearly what we see and document our technical activities and results.

On the level that is visible for the customer, of course, we have to choose a different language, adapted to the particular client, and as detailed as necessary in terms of content. We know some of the "regular clients" and we can adapt our language to their skills and manners. The goal is to formulate our resolution in a friendly but clear way and to show that it is not enough to simply change a few parameters on the database. It is also important to point out possible solutions or alternative implementations.

Identification of the Performance Problem

Dani: Now we are already at the solutions. But the first challenge for us often is to find out where exactly the cause of a performance problem lies. How can we identify where most time is being lost?

I find it useful that 2nd level support creates an AWR report for the critical time period for each performance ticket. I have

to admit that I only need these AWR reports to get a first overview. For the detailed information, I checked directly into the database. What is your approach? Do you often work with the AWR report or do you have other variants?

Martin: I completely agree with you: an AWR report is always just an average over a certain period of time between the initial and final snapshot. It gives us basic information about the system like database version, host or cluster name. It gives us also average values over the report period like the total load or the various TOP-reports. The AWR report alone does not tell us whether the specific problem is to be found among these.

Dani: Does that mean that, besides the AWR report, you use other sources of information?

Martin: There are tools that collect more information than AWR, for example SQLdb360 [1]. Here, there are hundreds of tests that collect detailed information, for example the number of archive logs on an hourly basis, grouped over days. This gives a good view of the transaction load. A big advantage of this tool is that for each individual check, the SQL with which it was performed is also specified. So, if you want to do a deeper analysis based on this, you have a very good starting point. Which tools do you use, to go into more detail?

Dani: Occasionally I use the troubleshooting scripts from Tanel Pöder [2]. I do not know all the scripts from his library, but very often I use `ashtop.sql` and `dashtop.sql`. Using these scripts, information from the Active Session History (ASH) can be displayed and grouped according to various criteria.

If I already know which SQL is affected, then I can use the



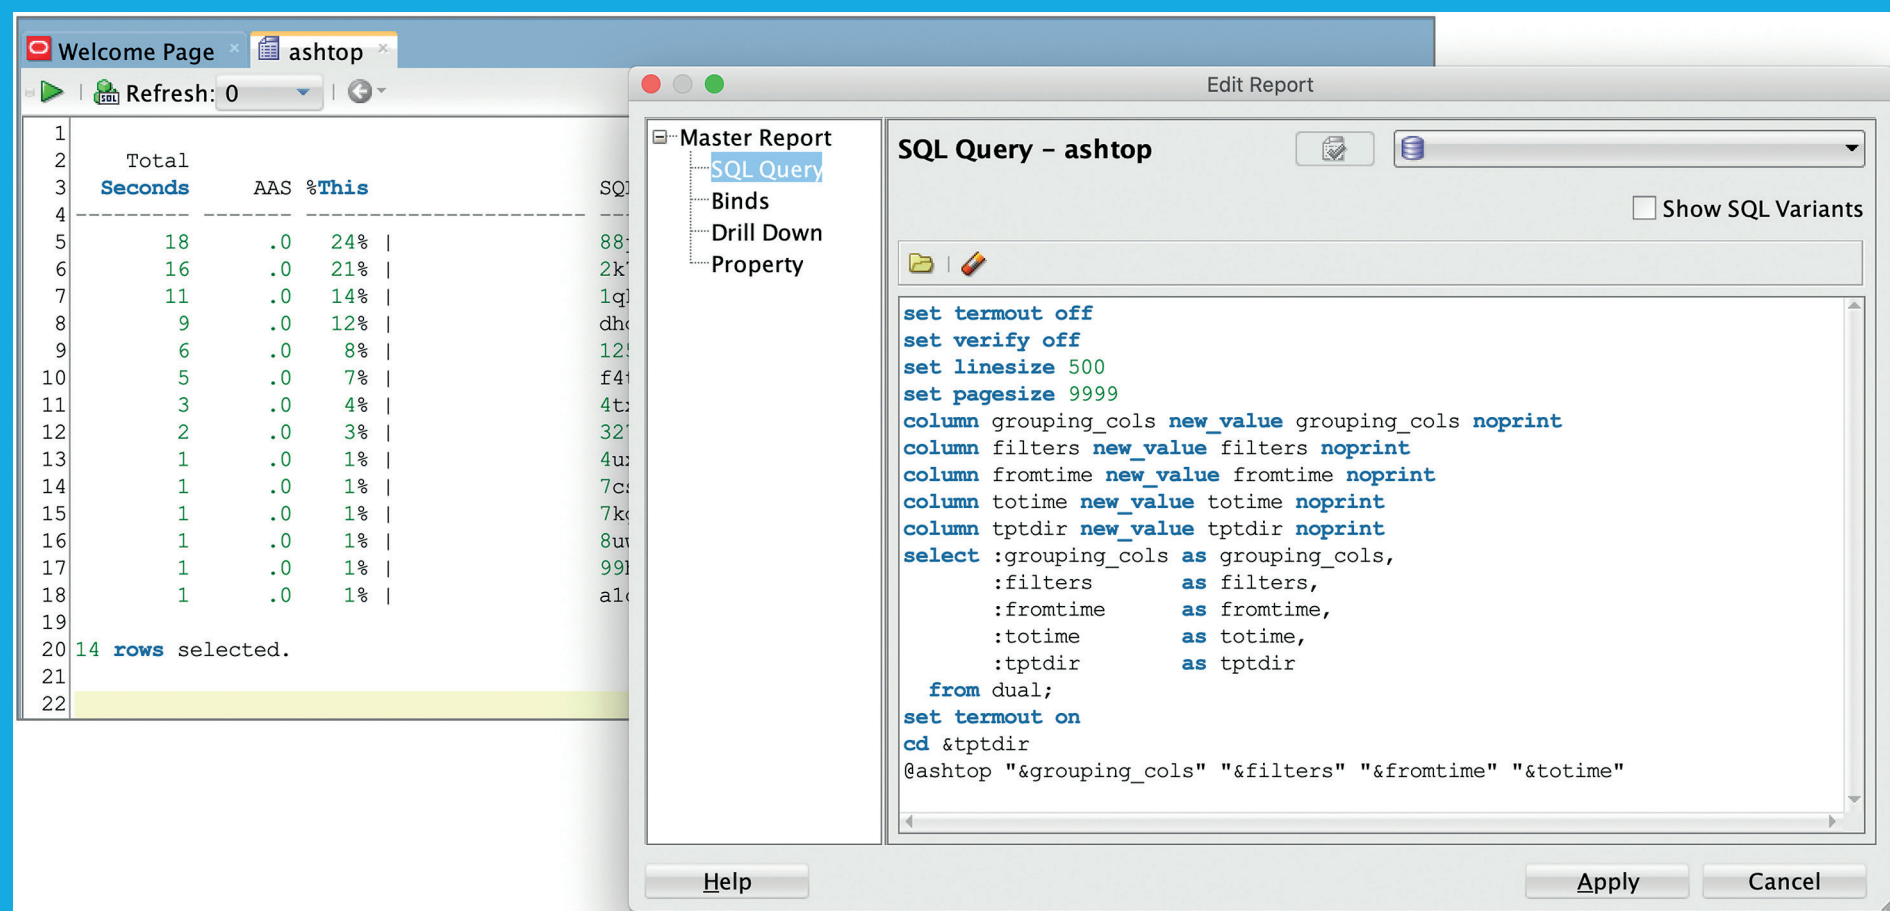


Fig. 2

script to determine the most frequent wait events for this SQL. Or I can determine the slowest SQL statements that were executed in the period in question under a certain user name. Most of this information can also be found in the AWR report, of course, but here I have more possibilities, to narrow down my questions by parameters. more precisely.

As a typical developer, I usually use tools like SQL Developer or Toad. However, the scripts from Tanel Pöder's scripts, which I use regularly. I can run them as "User Defined Reports" in SQL Developer (see Fig. 2).

I have applied the method, in the same way, in my environment at the customer for Tanel Pöder's scripts, which I use regularly. I can run them as "User Defined Reports" in SQL Developer (see Fig. 2).

Martin: I can see that, somehow, we are on different paths with our choice of tools. But with all these tools, the first step is usually to identify the performance problem, i.e. the "problematic" SQL_ID, the PL/SQL block or the business process.

Dani: Sometimes it's even easier. If it is already stated in the ticket which SQL is causing the performance problem, we can find the SQL_ID in the shared pool (gv\$sql) or via the AWR history views (dba_hist_sql_text) (see Listing 1).

```

SELECT sql_id, sql_text
FROM gv$sql
WHERE UPPER(sql_text) LIKE '%MV_CURIOUS%';

SELECT sql_id, sql_text
FROM dba_hist_sqltext
WHERE UPPER(sql_text) LIKE '%MV_CURIOUS%';

```

Listing. 1

Instrumentation

Martin: But the tools we have so far do not directly help us to distinguish at once the problem within all other concurrent activities. This often requires some manual work.

A successful approach to this is generally instrumentation. This involves clearly marking a specific process or flow and thereby making it identifiable and measurable. In Oracle,

this is very easy and even free of charge: Instrumentation is available in all versions and editions, already integrated in many tools and directly usable: in v\$session, in SQL Trace, in AWR or in Statspack. Enterprise Manager also uses it. Additionally, there is no performance overhead, it only needs to be enabled programmatically. In principle, this is possible in all environments: be it in PL/SQL, clients like jdbc thin/thick, OCI, .NET, php. Even in the cloud, instrumentation is used. So, from Oracle's point of view, it is the optimal preparation to be able to identify problems of any kind.

Dani: As a developer I also try to log information about the executions of procedures and functions, for example with the OraOpenSource PL/SQL Logger [4]. This works well in case of PL/SQL or any other programming language. But many tools have no logging mechanisms or proprietary ones. Our customer uses Informatica PowerCenter as an ETL tool. As far as I know, we don't have access to Informatica's log files? That could have been useful?

Martin: Informatica does not provide anything for instrumentation out of the box. This is a challenge for everyone involved: the DBAs, developers and architects need to be creative. You mentioned a logging function. With this you can add instrumentation at the beginning and at the end of each process, for example with a trigger or directly in the package.

For this customer, we chose a different way: Informatica itself offers a so-called transaction environment. With this we can execute an additional SQL at the beginning of the transaction and define the workflow name and the workflow run instance – the concrete execution – as modules and actions. This makes it very easy to identify the entire workflow that the customer considers to be a business process.

View Execution Plans

Martin: We have identified the place in the application that is responsible for the performance problems – either through instrumentation or using one of the methods discussed earlier. The next question would be: How do you look at the execution plans in each case?

Dani: Most of the time, I use the classic way with package `dbms_xplan`. If I can execute the SQL directly, I use the hint `/*+ gather_plan_statistics */` and then look at the executed Execution Plan with `dbms_xplan.display_cursor`. The advantage here is that I not only see the estimated cardinality of the optimizer ("E-Rows" column), but also the actual number of rows executed in each line ("ARows" column). For statements that run longer and which I can't execute directly, I use `dbms_xplan.display_awr`.

Martin: It is important to add here that the Adaptive Plans should also be displayed – for `dbms_xplan` with the parameter `format => '+adaptive'`. If you don't want to use the command line method, you can of course use Enterprise Manager or SQL Developer.

Dani: Yes exactly. When I am testing interactively, I also often choose the convenient solution and show the execution plan with "Autotrace" (F6) in SQL Developer. As with `dbms_xplan.display_cursor`, the actual execution plan is displayed, but unfortunately not with the information about the number of rows per line by default. However, the configuration of SQL Developer can be changed so that the additional `LAST_OUTPUT_ROWS` column (corresponding to "A-Rows" in `dbms_xplan`) is also displayed (see Fig. 3).

Martin: In SQL Developer, we also get a delta of the session statistics right away with the autotrace function. This



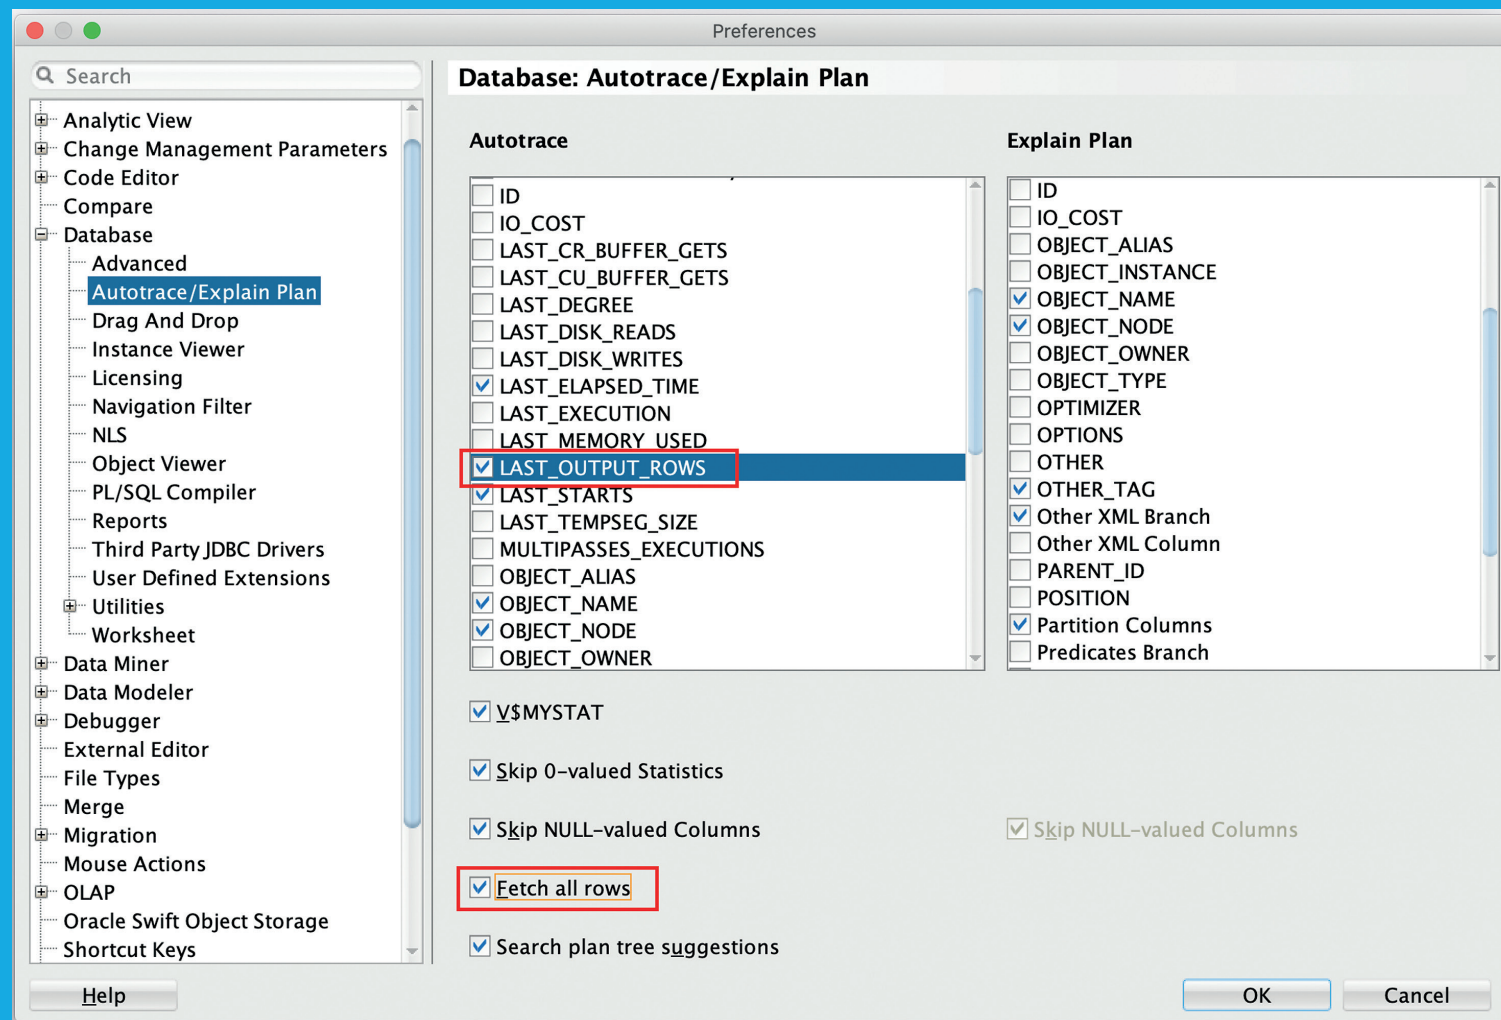


Fig. 3

additional information shows what happened between the beginning and the end of the execution of the SQL in the session. Unfortunately, "Fetch all rows" is not enabled by default and only the first 50 rows are read and evaluated. In order for the runtime statistics to refer to the entire execution, "Fetch all rows" must also be enabled in the settings (see Fig. 3). But now a general question: You mentioned different ways how the execution plan is displayed, but I still don't know where exactly in the plan the time is lost.

Dani: Yes, that's not always easy to see with complex plans. If it's a SQL that doesn't exactly run for hours, I also like to use SQL Real-Time Monitoring - a very cool way, in my opinion, to display the most important information about the execution of a SQL statement. Here you can see graphically how much time was taken in each line of the execution plan.

Martin: I'm also a big fan of SQL Real-Time Monitoring, especially the interactive HTML report. Here most of the information is available at a glance and also easily explained in a conversation. If we can't record the SQL with SQL-Monitoring, we still get to most of this information. In the view `gv$active_session_history`, there is a column `sql_plan_line_id` that exactly matches the line number in the Execution Plan. If we aggregate the information as shown in Listing 2, we can see where proportionally the most time was spent via the information `count(*)`. This is a bit more elaborate than the report, but we get the same information.

```
select sql_full_plan_hash_value, sql_plan_hash_value,
       sql_plan_line_id, count(*) cnt
from gv$active_session_history ash
where ash.sql_id='gdtppzfkmlw'
and sample_time between '2022-01-25 02:27:26'
                    and '2022-01-31 02:27:26'
group by sql_full_plan_hash_value,
         sql_plan_hash_value, sql_plan_line_id
order by sql_full_plan_hash_value,
         sql_plan_hash_value, sql_plan_line_id;
```

Listing. 2

Dani: In the past, I often used SQL Trace and TKPROF, because they show the number of processed rows and the execution times of parse, execute and fetch phases directly. But with many customers this is quite tedious, as I usually don't have OS-level access to the database server.

Martin: In cloud-like or very restrictive environments, we as DBAs often do not have access to the operating system or the diagnostic repository. But we have our own views for this purpose: In `V$DIAG_TRACE_FILE` we see a list of trace file names - an equivalent to `ls` or `dir`. In `V$DIAG_TRACE_FILE_CONTENTS` we can see the content of the respective trace file: Each line of the file is a line of the view, nicely sorted with a line number. In this case I was once so lazy and wrote a solution for the SQL Developer: A small report that gives me the trace file name and in a second step outputs the information as a LOB, which I can simply then save locally with a right mouse click (see [5]).

Why do we have different plans?

Dani: Often we have to deal with situations where different execution plans are used for the same SQL. There are different reasons ...

Martin: Yes, for example new statistics after major changes in data volumes or data distribution. Or new database objects, for example indexes, were created ...

Dani: ... but the customers always answer that they haven't changed anything. New features after database upgrades or patches can also be a reason. Here I often have a hard time figuring out which feature exactly is responsible for a plan "toppling".

Martin: Yes, thank God, in most areas there is already the preliminary work of other people, for example the tool Pathfinder [6]. The idea behind it is to simply go through all possibilities of the optimizer to find out which parameter settings force which concrete path. All optimizer parameters, relevant session parameters and `_fix_` control numbers are tried through to deduce which settings lead to which concrete plan. On a modern database version this is over 2,000 tests. Depending on the parse or execution time, this is not for the impatient.

Dani: We often hear the statement from the customer: "Until yesterday the query was fast, today it takes much longer". To verify - or disprove - such statements, I usually use a query (based on the `dba_hist_sqlstat` view) that I have stored as "SQL History" in my Code Snippets in SQL Developer (see Listing 3). This allows me to quickly determine if and when the Execution Plan has changed in the last 30 days. This allows me to narrow down if we are dealing with a general problem in the SQL or if the plan has changed due to one of the reasons mentioned.

```
SELECT stat.sql_id
       , stat.plan_hash_value
       , TO_CHAR(snap.end_interval_time, 'dd.mm.yyyy') exec_date
       , ROUND(SUM(elapsed_time_delta)/1000000) elapsed
       , ROUND(SUM(cpu_time_delta)/1000000) cputime
       , ROUND(SUM(iowait_delta)/1000000) iotime
       , SUM(stat.rows_processed_delta) num_rows
       , SUM(stat.executions_delta) execs
       , ROUND(SUM(elapsed_time_delta)
              / GREATEST(SUM(stat.executions_delta), 1)/1000000, 2) sec_per_
exec
FROM   dba_hist_sqlstat stat
JOIN   dba_hist_snapshot snap
      ON snap.dbid = stat.dbid
      AND snap.instance_number = stat.instance_number
      AND snap.snap_id = stat.snap_id
WHERE  stat.sql_id = '&sql_id'
      AND stat.elapsed_time_delta > 0
GROUP BY stat.sql_id
       , stat.plan_hash_value
       , TO_CHAR(snap.end_interval_time, 'dd.mm.yyyy')
ORDER BY TO_DATE(exec_date, 'dd.mm.yyyy');
```

Listing. 3

Martin: I notice that we are doing similar things, although we are using different tools and methods for this.

In the end, we have a wide variety of methods at hand to achieve our common goal: a high performing system.

Sources:

- [1] <https://github.com/sqlldb360/sqlldb360>
- [2] <https://github.com/tanelpoder/tpt-oracle>
- [3] <https://www.salvis.com/blog/2019/10/24/integrate-sqlplus-scripts-in-sql-developer/>
- [4] <https://github.com/OraOpenSource/Logger>
- [5] https://github.com/berx/oracle_scripts/blob/master/getTracefile.xml
- [6] <https://mauro-pagano.com/2015/10/26/introducing-pathfinder-is-there-a-better-plan-for-my-sql>



 This article was first published in Red Stack magazine issue 04/2023.



About Dani Schneider

Dani works as a Principal Oracle & DWH Consultant at Callista, a Swiss IT consulting company. He has been working in database consulting for over 20 years, mainly in the area of data warehousing. His tasks include DWH architecture reviews, data modeling, ETL development and performance optimization in Oracle databases. Dani is Oracle ACE Director and co-author of the books "Data Warehousing with Oracle" (2011) and "Data Warehouse Blueprints" (2016), both published by Hanser-Verlag.



About Martin Berger

Martin is in direct and constant contact with Oracle databases since 2000 - always with a great desire to understand complex technical aspects - hunting for performance secrets. Martin is a Symposium42 and Oracle ACE member. He loves family, people, technical talks, and beer.

Back to OUGN –

Creating a Conference Post Pandemic

Ann-Sofie Vikström Often



In April, we invited our community for a two-day conference at Rebel.no. The theme this year was “Back to OUGN”, as a connection to the movie “Back to the Future”. Rebel owns one of the DeLoreans from the movie. This was our first bigger conference after the pandemic. Normally we would have invited for a conference cruise to Germany and back to Oslo. Since we have seen symptoms in other physical conferences with less people than before, we were a bit scared to arrange such a big conference.



The choice of Rebel seemed to be a great choice. Rebel is a community venue with several conference rooms in different sizes and themes. So even if we would not get as much people as we wanted, the venue looked full. Our event partner Macsimum made a huge effort to help us with everything and even the marketing part.

We gathered around 270 Oracle enthusiasts and professionals for this 2-day event. Our conferencier Matt C Smith kickstarted the first day rolling in as Marty McFly on his skateboard.

There were 7 tracks with topics for Analytics, Oracle APEX, Oracle APPS, Development, Database, Oracle Cloud Infrastructure and Strategy and More. We have seen a big interest in Oracle APEX and wanted to test out a split from development. This was a great success and we even had attendees from Sweden for the Oracle APEX track. After dinner on the first day, we had an afterparty with DJ and band. For call for papers, we used Papercall.io which was very easy to use and to communicate with speakers. We had a luxury problem with a lot of good speakers and topics, and we hope the ones we rejected will come back for OUGN2024. For the conference App we used an Oracle APEX application provided by Smart4Solution.nl. A big thank you to Richard

Martens, Sergei Martens and team for the great effort in supporting us. Oracle community support was really helping out getting Oracle speakers and we were grateful to have both Chris Saxon and Connor Mc Donald. Thanks to Jennifer Nicholson and Mary Beth Pierantoni for their huge support.

For a local Oracle user conference, we cannot do anything without the huge support from our Oracle office in Norway. They say our event is the most important event, and we are humbled for their enthusiasm in helping us out. This year we had to sign new partners which was not so easy. Having a good sales team is crucial, and we got 5 sponsors in the end, in addition to Oracle. It is important to have an attractive package for the partners, but we will also use the success for this year's conference to attract new partners for OUGN2024 where we will stay at Rebel.



About Ann-Sofie Vikström Often

Ann-Sofie is a consultant at Cegal AS in Norway. The last couple of years working as Oracle APEX developer for different Oracle customers. Outside of work she is active as the board leader of Oracle User Group Norway (ougn.no) and part of the board in EOUC, an umbrella organization for user groups in EMEA. In April awarded as Oracle ACE associate during the OUGN2023 conference. In 2020 she was awarded as one of the 50 leading IT women in Norway, for her contribution to the digitalization of the public sector amongst all. Having a dog, a lot of time goes to keep him happy with walks and meaningful training. The rest of the time she is spending with her family as much as possible skiing, trekking, biking, swimming etc.



Call for Papers

Michigan Oracle Users Summit

September 1, 2023
<https://www.mous.us/present.html>

Analytics and Data Summit 2024

December 1, 2023
<https://sessionize.com/andsummit2024>



Follow us on



@EOUC @ORAWORLD_Mag #ORAWORLD



@emeaoracleusergroups @ORAWORLDMag



www.oraworld.org

Events

LAOUC Community Tour 2023

July 28 - August 16, 2023

Cali, Colombia: July 28 | San José, Costa Rica: July 31 | Ciudad Panamá, Panamá: August 2 | Guadalajara, México: August 4 | Ciudad Guatemala, Guatemala: August 7 | Santiago, Chile: August 9 | São Paulo, Brazil: August 12 | Montevideo, Uruguay: August 14 | Buenos Aires, Argentina: August 15 | Asunción, Paraguay: August 16

<https://www.laouc.org/laouc-community-tour-2023>

SOUG Day

September 5, 2023

Winterthur, Switzerland

<https://soug.ch/event/soug-day-5-sept-2023>

Oracle CloudWorld 2023

September 18 - September 21, 2023

Las Vegas, NV, USA

<https://www.oracle.com/cloudworld>

EMEA Community Tour

October 2 - October 19, 2023

Copenhagen, Denmark: October 2 | Espoo, Finland: October 3 | Utrecht, Netherlands: October 5
Bishkek, Kyrgyzstan: October 9 | Tashkent, Uzbekistan: October 10 | Dushanbe, Tajikistan: October 12
Baku, Azerbaijan: October 14 | Madrid, Spain: October 17 | Rovinj, Croatia: October 19

<http://ogbemea.com>

HrOUG 2023

October 17 - October 20, 2023

Rovinj, Croatia

<https://2023.hroug.hr/eng>

AUSOUG CONNECT 2023 (in-person and online)

November 13 - November 15, 2023

Perth, Australia: November 13

Melbourne, Australia: November 15

<https://www.ausoug.org.au/connect-2023>





UKOUG Conference & Exhibition '23

November 15 - November 16, 2023
Reading, UK
<https://ukoug.org>

POUG 2023 November

November 17 - November 18, 2023
Warsaw, Poland
<https://poug.org/en>

DOAG 2023 Conference + Exhibition

November 21 - November 24, 2023
Nuremberg, Germany
<https://anwenderkonferenz.doag.org/en/home>

AI Navigator

November 22 - November 23, 2023
Nuremberg, Germany
<https://ki-navigator.doag.org/en/home>

IT-Tage 2023

December 11 - December 14, 2023
Frankfurt, Germany
<https://www.ittage.informatik-aktuell.de/konferenz/einfuehrung.html>

East Coast Oracle Users Conference 2023

December 12 - December 13, 2023
Raleigh, NC, USA
<https://www.eastcoastoracle.org>



ORAWORLD is a publication of the EOUC – EMEA ORACLE USERGROUP COMMUNITY

The following user groups belong to EOUC:

Angola Oracle User Group, Oracle User Group Armenia, Austrian Oracle User Group, Azerbaijan Oracle User Group, Bulgarian Association of Software Developer, Bulgarian Oracle User Group, Hrvatska udruuga Oracle korisnika, Czech Oracle Applications User Group, Danish Oracle User Group, Egypt Oracle Users Group, Oracle User Group Estonia, Oracle User Group Finland, Club Français des Utilisateurs JD Edwards, Association des Utilisateurs Francophones d'Oracle, Club des Utilisateurs PeopleSoft, Oracle User Group Georgia, Deutsche Oracle Anwendergruppe, PeopleSoft Germany, Hungarian Oracle User Group, Israel Oracle User Group, Taranta Valley Oracle User Group, Italian Oracle User Group, Jordan Amman Oracle User Group, Latvian Oracle Users Group, Lithuanian Oracle Users Group, Mauritius Oracle User Group, Oracle Gebruikersclub Holland, Oracle Benelux User Group, Oracle User Group Norway, Polish Oracle Users Group, Oracle Users Group Portugal, Romanian Oracle User Group, Russian Oracle User Group, EBS Finance Special Interest Group Russia, Arab Oracle User Group, Serbia and Montenegro, Slovenian Oracle User Group, South African Oracle User Group, Spanish Oracle User Group, Swedish Oracle User Group, Swiss Oracle User Group, Tajikistan Oracle User Group, Turkey Oracle Users Group, Ukraine Oracle User Group, Middle East Oracle User Group, United Kingdom Oracle User Group, Zimbabwe Oracle User Group.

Editorial board:

Registered office: DOAG Dienstleistungen GmbH
Tempelhofer Weg 64, 12347 Berlin, Germany
www.doag.org,
Director Fried Saacke,
AG Berlin Charlottenburg HRB 95694B,
VAT ID DE240700058
Contact: redaktion@doag.org
Editor-in-chief (ViSDP): Mirela Ardelean

Editorial team:

Mirela Ardelean, Jean-Jacques Camps, Heli Helskyaho,
Ann-Sofie Vikström Often, Andrejs Vorobjovs

Other editors: Christian Luda

Other authors are credited by name with their article. They are subject to the terms and conditions for authors: www.oraworld.org/terms-and-conditions-for-authors

Graphic design:

Alexander Kermas, DOAG Dienstleistungen GmbH,
Tempelhofer Weg 64, 12347 Berlin, Germany

Disclaimer:

All rights reserved. Duplication or retransmission in whatever form or whatever medium either in whole or in part requires written permission, to the extent the content is not made available for duplication or retransmission.

The information in this publication has been duly researched and is correct to the best of our knowledge and belief. The use of this information is at your own risk. No liability for the accuracy of the information is accepted and, in particular, for its practical application in individual cases. Opinions represent the views of the individual author and do not necessarily represent the view of the publisher.

The ORAWORLD e-magazine reports on events in the Oracle and IT world. The publication covers current topics in the international user group network, as well as news items on products and technologies and their use. The purpose of the magazine is to foster the sharing of knowledge and experience among readers. ORAWORLD is independent of Oracle and does not represent its commercial interests either directly or indirectly.

ORAWORLD is published by DOAG Dienstleistungen GmbH, Tempelhofer Weg 64, 12347 Berlin, Germany, legally represented by director Fried Saacke, the nature and purpose of whose business is management of the group, organising events and publishing.

DOAG Deutsche Oracle Anwendergruppe e.V. holds 100 percent of the capital invested in DOAG Dienstleistungen GmbH. DOAG Deutsche Oracle Anwendergruppe e.V. is legally represented by the managing committee; Chair: Björn Bröhl.

Article submission:

If you are interested in submitting an article, please e-mail us your article via the online form at www.oraworld.org. Deadline for the next edition: **September 22, 2023**.

Photo credits:

Title: © Designed by fullvector | <https://freepik.com>
P. 4: © floralset | <https://123rf.com>
P. 5: © Comic | www.commitstrip.com
P. 6 © Created by bing | <https://bing.com>
P. 7: © Irina Strelnikova | <https://stock.adobe.com>
P. 12: © Somchai | <https://stock.adobe.com>
P. 17: © Roman | <https://stock.adobe.com>
P. 24: © Designed by vecstock | <https://freepik.com>
P. 31: © schuetz-mediendesign | <https://www.pixabay.com>
P. 39+40: © Ann-Sofie Vikström Often
P. 41: © Designed by pch.vector | <https://freepik.com>
P. 42: © stnazkuls | <https://123rf.com>
P. 43: © Designed by stories | <https://freepik.com>