

DATA SHEET **EXASOL V7** 

EXASOL V7

YOUR GUIDE TO WHAT'S NEW
AND HOW YOU'LL BENEFIT

Exasol
The analytics database

CONTENTS

INTRODUCTION – HELPING YOU UNCOVER BETTER INSIGHTS AND ADD MORE VALUE	3
BETTER PERFORMANCE – HELPING YOU BECOME MORE AGILE	4
MORE FLEXIBILITY TO DEAL WITH STRUCTURED AND SEMI-STRUCTURED DATA	6
STAYING COMPETITIVE WITH DATA SCIENCE	7
WHAT ELSE HAS IMPROVED FOR YOU?	8
WE'D LOVE TO HEAR YOUR FEEDBACK	8



INTRODUCTION

HELPING YOU UNCOVER BETTER INSIGHTS AND ADD MORE VALUE

OUR ANALYTICS DATABASE HAS ALWAYS BEEN DESIGNED TO FIT YOUR DATA STRATEGY – POWERFUL ENOUGH TO HANDLE WHATEVER YOU THROW AT IT AND GIVING YOU THE FREEDOM TO USE DATA TO ACHIEVE YOUR UNIQUE GOALS.

But we're always striving to make things better for you. By improving our **performance, flexibility and data science capabilities** in key areas we've now made it even easier for you to:

- › Run your analytical models on even larger data volumes to find those business-critical answers and increase visibility.
- › Do it at speed, so you don't miss important business opportunities.
- › Transform your employees into a data-driven team with an infrastructure that helps everyone extract more value from your data



FIND OUT WHAT'S NEW

DATA VAULT BETTER PERFORMANCE – HELPING YOU BECOME MORE AGILE

We've improved performance in three key areas of Data Vault, helping you optimize your data modelling approach, become a more agile organization and meet the twin challenges of:

- Increased global uncertainty due to rapid technological and business change
- The growing volume, complexity, and distribution of your data



HASHES MEAN BETTER PERFORMANCE FOR YOUR ANALYSES AND REPORTING THANKS TO **IMPROVED JOINS** BETWEEN TABLES AND SCHEMAS.

DATA VAULT - HASH DATA TYPE IMPROVES QUERY PERFORMANCE

WHAT'S NEW?

In Data Vault (2.0) all keys are stored as hashes. These hashes can now be stored with our Hash Data Type – enabling you to combine data from multiple tables and schemas much faster than before.

HOW WILL IT HELP YOU?


Hashes mean better performance for your analyses and reporting thanks to improved joins between tables and schemas. These joins use hashes and run approximately four times faster than joins based on strings.

Using the hash data type in your data model also gives your data engineers (and anyone involved in the data preparation processes) even more options for bringing additional tables and data sources into the model, adding granularity and opportunities for more detailed analyses while improving performance further.

To start benefiting from the new Hash Data Types in version 7.0, all you need to do is use the 'alter table' command.

DATA VAULT BETTER PERFORMANCE – HELPING YOU BECOME MORE AGILE

To answer key business questions, you often need to join large tables and schemas that are distributed across your cluster. When joins are not local on a single node, complex queries can impact performance as your data volumes continue to grow.



YOU CAN IMPORT DATA INTO EXASOL EVEN FASTER AND YOUR QUERIES ON PARTITIONED TABLES WILL RUN MORE EFFICIENTLY.

DATA VAULT - GLOBAL JOIN OPTIMIZATION IMPROVES PERFORMANCE FOR QUERIES ACROSS THE CLUSTER

WHAT'S NEW?

We've now fully optimized our global joins.

HOW WILL IT HELP YOU?

Global joins across your database cluster will perform much better, which means you don't need to force these joins to be local anymore. You can benefit from our global join optimization whether you model your data in a star-schema, Data Vault or using other approaches. This is an out-of-the-box solution, so you will see the benefits immediately.

DATA VAULT – FASTER PARTITIONING: IMPORT DATA FASTER THAN EVER BEFORE

WHAT'S NEW?

We've optimized data inserts into and queries on partitioned tables.

HOW WILL IT HELP YOU?


You can import data into Exasol even faster and your queries on partitioned tables will run more efficiently. Our improvements to how Exasol handles partitioned tables mean you can run bulk inserts into your tables, even when you need to do this in real time or near-real time.

JSON

MORE FLEXIBILITY TO DEAL WITH STRUCTURED AND SEMI-STRUCTURED DATA

Although trickier to analyze than structured data, semi and unstructured data is far more prolific in the enterprise. Indeed, IDC has projected that 80% of worldwide data will be unstructured by 2025, with the vast majority of this influenced by the rise of IoT.

To stay competitive, we're giving you greater flexibility to handle semi-structured data directly in the database. Version 7.0 natively supports both structured and semi-structured data formats, bringing you the benefits of scale and performance to more types of data, while also making it easy to work with your data directly in SQL.



DOCUMENTS CAN BE PROCESSED IN THE DATABASE MUCH FASTER THAN BEFORE.

NATIVE JSON FUNCTIONS SAVE YOU TIME TRANSFORMING DATA

JSON FILES ARE COMMONLY USED FOR SEMI-STRUCTURED DATA BUT MOST RELATIONAL DATABASES STORE DATA IN TABLES.

WHAT'S NEW?

To save you time when importing and converting your data, all JSON functions are now natively integrated in SQL and can be executed directly in the database. We removed the previous need for using UDFs and you don't have to build complex ETL processes or do any up-front data modeling.

HOW WILL IT HELP YOU?

- ▶ Documents can be processed in the database much faster than before.
- ▶ You'll have more flexibility to work with data sources that produce JSON files as outputs.
- ▶ You'll save time preparing data for analysis by being able to quickly and easily transform your semi-structured data from JSON files into structured data in Exasol.

GPUs

STAYING COMPETITIVE WITH DATA SCIENCE

AI and ML are essential for your ability to drive innovation. But the key is staying competitive when 53% of enterprises are already using AI (according to a 2018 survey of data and analytics decision makers by Forrester) - and adoption is only set to increase.

To help you do this, we've made it easier for you to build or utilize a data infrastructure that's scalable and flexible enough to meet quickly evolving and expanding AI workloads.



YOU'LL SEE A **PERFORMANCE BOOST** WHEN TRAINING YOUR MODELS USING GPUs.

GPUs (GRAPHICS PROCESSING UNITS) – POWER YOUR DEEP LEARNING AND AI APPLICATIONS

WHAT'S NEW?

We're bringing together the power of parallelization from our database software with the parallel architecture of GPUs. Using our UDF framework in combination with GPUs means you can bring any data science language into our database and work with it faster than ever before.

HOW WILL IT HELP YOU?

- ▶ You'll see a performance boost when training your models using GPUs. They can be deployed in the cloud or on premises to power your machine learning applications. Running through multiple epochs in shorter amount of time will improve the time to value of your data science projects.
- ▶ You'll be able to run deep learning applications directly in the database, making it more economically viable, especially as you utilize the flexibility of the cloud.

WHAT ELSE HAS IMPROVED FOR YOU?

HELPING YOU STAY GDPR COMPLIANT – WITH HIGHLY PERFORMANT SYSTEMS

According to a 2019 survey of security professionals just over half, 59% of companies believe they are GDPR-compliant.

With the latest improvements to our product, we've gone beyond GDPR compliance and built features directly into the database that securely process and manage your data. We've developed a fully compliant GDPR plug-in with our partner, Protegrity, a commercial vendor of data protection and privacy solutions. The plug-in equips you with row level security and helps you automate the protection, masking and encryption of your business users' information - without compromising performance across large data volumes.



59% OF COMPANIES BELIEVE THEY ARE **GDPR-COMPLIANT**.

USAGE PRIORITIZATION GIVES YOUR DATABASE ADMINISTRATORS MORE CONTROL

You can now prioritize more than just CPU workloads and keep in step with demands from the business with more fine-grained control. This gives greater control to allocate resources in line with your organization's needs – such as controlling Temp DB RAM at a global level, for individual users and individual sessions.

USE OUR PREVIEW FEATURE TO SEE WHAT'S COMING UP – AND BUILD YOUR CASE FOR UPGRADES

Our new preview feature allows your users and administrators to see upcoming features and understand how they would improve your existing processes. For instance, you'll be able to show your stakeholder whether certain operations and manipulations will run faster and more efficiently by upgrading to the latest version.

WE'D LOVE TO HEAR YOUR FEEDBACK

We're really keen to know how you get on with the new updates. If you have any feedback, please get in touch with your account manager.

Or simply log on to our [user portal](#) - to share your insights, exchange ideas and get support from your community of fellow Exasol customers.

EXASOL - THE ANALYTICS DATABASE



[CLICK HERE](#) TO VISIT
EXASOL.COM

TAKE THE NEXT STEP AT EXASOL.COM

GERMANY

Exasol AG
Neumeyerstr. 22 – 26
90411 Nuremberg
German

T: +49 911 23991-0
F: +49 911 239 91 24
E: info@exasol.com

UNITED KINGDOM

Exasol UK Limited
Parkshot House
5 Kew Road
Richmond
London
Greater London
TW9 2PR
United Kingdom

T: +44 20 3813 8310
E: info@exasol.com

USA

Exasol, INC.
c/o WeWork
1372 Peachtree Street
Atlanta, GA, 30309
USA

T: +1 415 363 5500
E: info@exasol.com

Exasol
The analytics database