


# AN INTRODUCTION TO DATABASE RELIABILITY



Data is arguably the most important asset that your organization has. Your database is a key part of your organization's infrastructure, and if it goes down, it can create major problems. Data is key for handling daily activities and for making short and long-term decisions. For an organization to run effectively and efficiently, it needs to have a reliable database.

Given the essential role that data plays, many organizations will avoid any sort of data-related innovations or changes out of fear of creating problems. As you can imagine, this can actually lead to issues with the database, making it difficult for some organizations to maintain a competitive database. The good news, however, is that adequate database reliability ensures that organizational data is protected, that engineers are free to innovate, and that necessary work is done efficiently.

## What is Database Reliability?

Database reliability is defined broadly to mean that the database performs consistently without causing problems. More specifically, it means that there is accuracy and consistency of data. For data to be considered reliable, there must be:

- **Data integrity**, which means that all data in the database is accurate and that there is consistency throughout data. Data consistency is defined broadly to include the type and amount of data.
- **Data safety**, which means that only authorized individuals access the database. Data security also includes preventing any type of data corruption and ensuring that data is always

accessible. When it comes to data safety, engineers must ensure that data is accessible even in the event of unforeseen circumstances, such as emergencies or natural disasters.

- **Data recoverability**, which means there are effective procedures in place to recover any lost data. This is a key to database reliability, ensuring that even if other safety measures fail, there is a system for recovering data.

## The Importance of Database Reliability

Organizational databases store a broad range of information, including customer information, sales information, financial transactions, vendor information, and employee records. This information is essential for maintaining the health of organizations and plays a central role in everything from competitive strategy to daily logistics. In many ways, data works as the eyes and ears of the organization, and without it, organizations lack the necessary information to make informed decisions. It's the database that makes this information accessible and usable.

If an organization's database is not reliable, consistent, or accurate, it can lead to making bad or misinformed decisions. Further, as the database is a central part of organizational infrastructure, if it goes down, it can lead to substantial issues throughout the organization. This means that database reliability is and must remain a central concern for businesses.

Yet, in the current environment, data problems are increasingly complex, making it continually difficult to create, manage, and manipulate databases. Given the importance of database reliability and the increased demands that come with database management, it's important that organizations have advanced and innovative approaches to ensuring database reliability. A couple of such approaches that organizations should consider implementing are database reliability engineering and the use of effective database management systems.

## Database Reliability Engineering

Database reliability engineering is an effective way for organizations to ensure database reliability and to ensure that organizations are able to effectively use data. Database reliability engineering is generally driven by the database reliability engineer, a data administrator that works to ensure that data is protected and accessible. Among other things, this increased reliability provides the safety and support needed to enable innovation and facilitate work.

It's worth noting that the phrase "database reliability engineering" was first coined by Laine Campbell and Charity Majors in their book, *Database Reliability Engineering*. This book can serve as a good resource for organizations looking to strengthen their database or wanting to implement new systems for promoting database reliability.

## The Role of the Database Reliability Engineer (DBRE)

First and foremost, the Database Reliability Engineer, or DBRE, is an enabler that allows other data and software engineers to work efficiently without causing problems. The DBRE allows engineers to work within data shares while also ensuring that all data is protected, reliable and accessible. In addition to this central role as an enabler, database reliability engineers:

- **Utilize automation.** A big part of database reliability engineering is automating tasks. Particularly important is automating safety operations, including failovers, backups, and back-

pressure mechanisms. It's this critical automation that lets engineers work quickly and efficiently without having to worry about losing or messing up data. These measures help to protect data and to encourage innovation among engineers.

- **Conduct risk analysis.** Whenever considering automation, database management, or utilizing new tools, it's important to conduct a thorough risk analysis. It's a DBRE's role to consider potential risks and then to make informed decisions.
- **Make decisions about scaling.** It's the role of the DBRE to anticipate capacity needs and to make timely decisions about scaling. Doing so helps to maintain database reliability, ensuring that the database is meeting organizational needs.
- **Educate other engineers.** Part of the DBRE role is knowledge sharing and educating other data software engineers on everything from the database to the organization's domain to best practices.

## Database Management Systems

An additional way to promote database reliability is through the use of an effective database management system. A database management system is a type of software that is used to create and manage databases. Additionally, database management systems retrieve, manipulate, and define data.

Utilizing an effective database management system can help to ensure that your database is always accessible, that it is safe from corruption, and that your data is accurate and consistent. When determining whether to use a database management system and which system to use, some key features to look for are high availability, corruption and debugging prevention, clustering, and type-safe API.

## Causes of Database Failures Can Be the Result of Many Factors

As you work to ensure that your organization's database is reliable, it's important to consider some of the issues that lead to database problems. As you conduct this analysis, keep in mind that it's often organizational infrastructure or hardware that leads to issues. So when something goes wrong, it can be helpful to look for causes among hardware or infrastructure. Some common sources of problems are issues with disks, RAM, or the motherboard.

Regardless of the specific issue, when looking to address issues with database reliability, keep in mind that problems are not always caused by software but can also be caused by infrastructure or hardware.

## Summary

Data is at the heart of everything your organization does. It helps you keep up with customers, financial transactions, employee information, vendor information, sales, and supply chain information. In addition to being central to short and long-term goals, data can be used to enable leaders to make high-level decisions that significantly impact how organizations grow and thrive.

As a result, it's critical that your organization's data is reliable. As data becomes increasingly complex, more and more organizations are changing their approach to database reliability to keep up with the changing environment. Database management systems and database reliability engineering are two effective ways to meet this need.