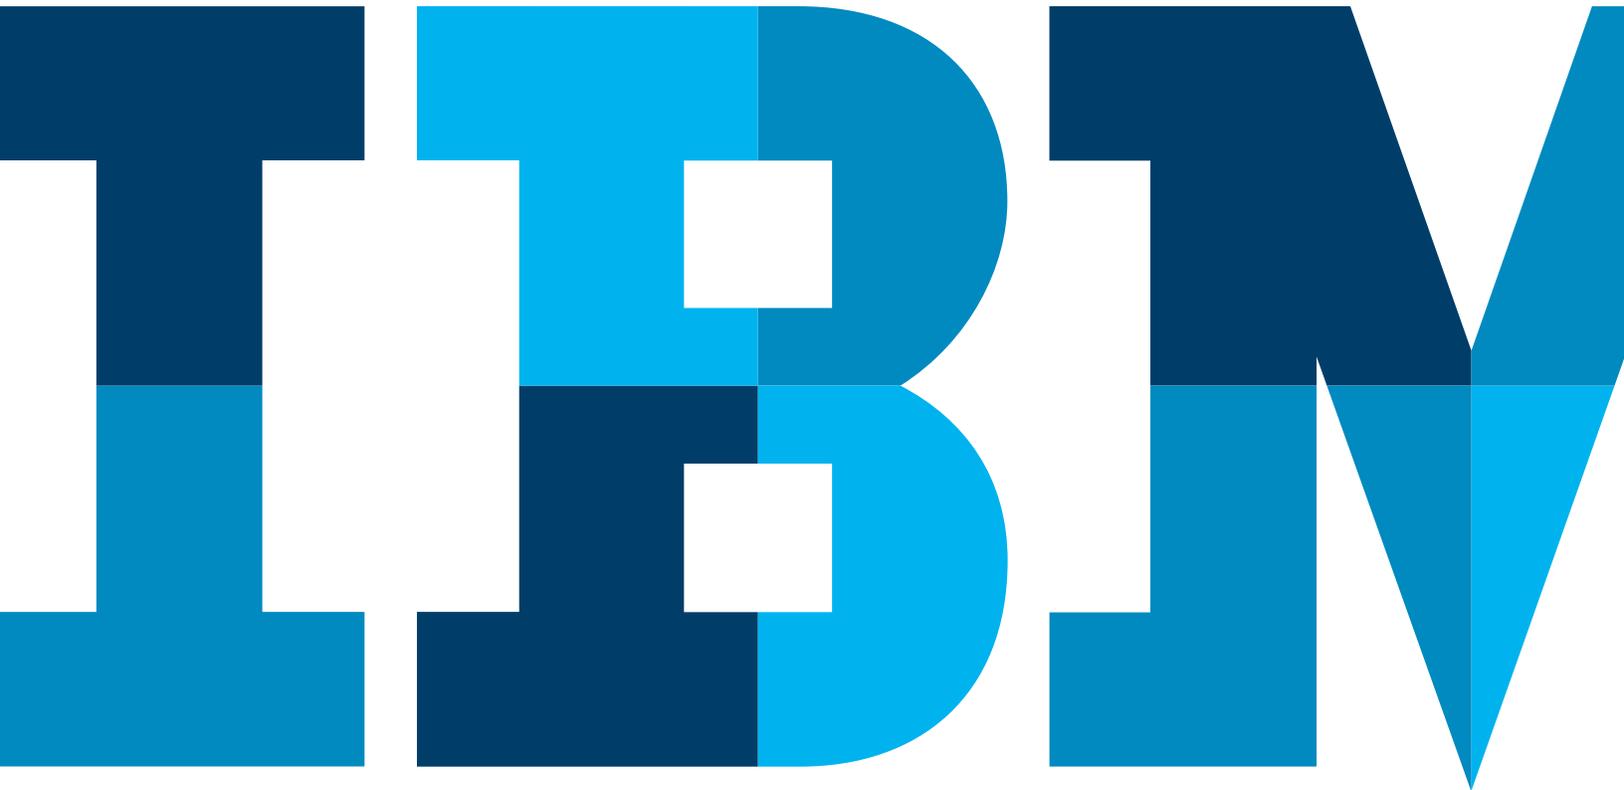# Achieving hybrid IT:
Five database must-haves

IBM

Introduction

According to IDC, by 2017, 80 percent of enterprises will have adopted a hybrid IT strategy composed of on-premises, private and public cloud components. In an increasingly cloud-first era, an agile, scalable infrastructure matters more than ever, and the need for a database that scales, up, out and into the cloud is evident.

Five key considerations for evaluating database options

The following five key considerations can help evaluate database options for tomorrow's enterprise, including a brief look at the changing nature of business and the growing importance of data to all parts of the enterprise.

1. Support for next-generation applications.

We are in an increasingly application-driven economy, where "there's an app for that" has gone from internet hype to truism. In the new digital era, businesses are increasingly relying on technology to make better-informed decisions as soon as possible. Development platforms that combine sophisticated database technology with decision-support tools provide the foundation for cognitive applications that can learn and react to customer preferences using natural language processing (NLP).

This can present significant challenges when it comes to application development. With the introduction of private and public clouds, mobile, big data and bring-your-own-device (BYOD) policies, modern applications must be able to integrate with numerous platforms, devices and data, and offer a seamless user experience across a variety of platforms and mobile devices.

2. Enabling big data for the enterprise.

Big data has brought major shifts to the way data is acquired, analyzed, stored and used. It offers enormous potential for businesses, but requires them to rethink how they access and use information, how they interact with consumers holding vital data and how they use the skills and technologies needed to take advantage of it. Big data means nothing unless it can provide big insights.

Two types of systems dominate the big data landscape: those that provide operational capabilities for real-time, interactive workloads where data is primarily captured and stored; and systems that provide analytical capabilities that can be used to perform complex, retrospective analysis of large volumes of data. Since these two systems complement each other, they are frequently deployed together.

Consequently, the database used must be able to handle traditional structured and non- or semi-structured data. It must also take advantage of the latest in-memory acceleration techniques to speed the processing of columnar data. Finally, it must use threading and vectorization to get the most out of today's multicore, multiprocessor servers and server clusters.

3. Support for true hybrid IT.

When it comes to cloud computing, there are usually two lines of thought: that the benefits of cloud computing—availability, flexibility, scalability, automatic maintenance, and pay-per-use pricing—can only be provided by a public cloud. The second perspective is that the potential risks of the public cloud outweigh its benefits; yet, internal infrastructure can benefit from being more cloud-like.

Both perspectives are right, in their own way. There are some benefits that the public cloud offers that cannot be replicated in a private cloud environment. And, in some situations or for some applications, an on-premises infrastructure is preferable to a cloud offering, regardless of whether the cloud is public or private.

As a result, the concept of a hybrid IT infrastructure has been introduced as a way for businesses to get the best of both worlds. The idea behind a hybrid IT infrastructure is exactly as the name implies: a mix between public and private cloud services and on-premises computing resources. With the rise of hybrid IT, it is clear that having data locked in silos and unavailable to other applications is an impediment to delivering the data-driven insights that applications and users demand.

Databases used in a hybrid IT environment must support multiple sources of information and connect traditional enterprise data to new web or mobile applications, regardless of whether they are created on premises or are born in the cloud.

4. Supporting cloud economics.
Because CIOs are constantly challenged to deliver IT services with the greatest value for the business, they must determine specifically how the principles, costs and benefits of cloud computing will affect their IT budget and staffing needs. Typically, when assessing cloud economics, CIOs and IT leaders weigh the costs associated with infrastructure, management, security and support to determine if moving to the cloud makes sense for a given situation.

As business demands and cloud options change, on-premises application workloads may be delivered more economically from the cloud. Or, born-in-the-cloud offerings may need to shift to on premises for regulatory or compliance reasons. Therefore, the database chosen must offer both power and agility, and cloud-ready features like self-service cloud deployment, subscription licensing and worldwide reach, to meet hosting requirements regardless of where it physically resides—whether in the US, the EU or somewhere else in the world.

5. Database security and uptime.
There is no time for downtime in our 24x7 economy. The penalties for not meeting data protection guidelines outlined in specifications such as ISO 27001, HIPAA, and PCI-DSS (Payment Card Industry Data Security Specification) can be severe.

Since the database is integral to the enterprise, it must use security mechanisms like authentication, authorities and privileges, and other built-in controls to ensure that a consistent set of rules for limiting data access based on an individual's need-to-know basis are strictly enforced. Equally important, the database chosen should use SSL encryption technology to protect sensitive information stored on media (data at rest) or that is transmitted through untrusted communication channels (data in transit).

Finally, the database must contain high-availability and disaster-recovery mechanisms to ensure that an organization can continue to operate in the event of a serious incident or disaster. These mechanisms need to be able to recover the organization to an operational state within a reasonably short period of time.

Making the right choice

IBM® DB2® database software delivers on all these points and more on a solid, versatile platform that grows with your future needs, whether on premises, in the cloud or both.

For more information

To learn more about IBM DB2 database software, please contact your IBM representative or IBM Business Partner, or visit: **ibm.com**/db2