How Organizations Can Turn Explosive Data Growth into an Advantage

bmc

A Data Deluge

Organizations are drowning in the very data they need to survive.
With so much data created every second, IT leaders often find themselves pulled between two critical, yet competing, demands: managing this ceaseless flood of information, and serving it up fast enough to meet customers' rising performance expectations. And as business leaders see answering—and exceeding—these demands as a key enabler of growth, the pressure to handle more data and transactions will only increase.

The trouble is, there's no simple, single answer to both demands. Changing an application or index structure won't eliminate the bind. Adding hardware may help mitigate outages, but the costs can quickly spiral out of control. Of course, it's always possible to meet demanding SLAs by skipping important data quality and integrity steps or creating unsustainable workarounds—but that just means more problems in the future.

That none of these approaches are ideal reveals a core truth about this era of exponential data growth: simply dedicating more resources to the problem will do nothing to minimize it.

In fact, adding more people, hardware, and infrastructure into the mix only serves to increase risk. Costs will continue to rise. An IT environment that continually grows in size and complexity eventually buckles under its own weight, leading to more downtime and more dissatisfied customers. And all the while, the data onslaught rolls on unabated.

Instead of piling resources on top of resources, businesses need a new way to manage what they have available today.

Getting there isn't easy. To keep pace with growing data volumes—while providing best-in-class service—it's imperative to first address:

The explosion of new data and its impact on critical systems and services

The pressure to prove the business can rely on its current backup and recovery measures

The constant demands from customers for continuous uptime

Growing Data, Overtaxed Systems

The rise of digital business practices has dramatically increased the amount of data flowing through today's mainframe environments. Not only have tens of thousands of transactions per minute become commonplace, new types of data—video files, voice recordings, large documents, and high-res images—must now be treated as business-critical assets. However, protecting and managing them is decidedly more difficult than traditional, structured data—especially when they're stored as large objects.

Large objects have long been ungainly and difficult to use. All too often, though, they're the only option. Because Db2 data types weren't designed with capacity in mind, protecting and ensuring the validity and integrity of large volumes—and newer types—of data becomes an increasingly complex problem.

Despite any difficulties this situation poses, business users are eager—and incentivized—to accumulate more data, analyze it, and extract insights that inform and enrich their decision-making processes. Of course, accommodating these massive, complex volumes of data threatens to slow the mainframe to a halt, at the precise time business users need it more than ever.

Managing mountains of data—especially when it takes on new, non-traditional forms—is not without its challenges. But there's not an organization in existence that couldn't benefit from more data. That's why it's so vital to find ways to turn the size and rapid growth of data into an advantage.

To do that, businesses need systems that can check, validate, and ensure the integrity of all types of data, including both structured and unstructured data. They need to be able to reorganize this data in a timely manner, so they can meet customers' performance expectations. And they must find ways to load and unload millions of rows of data as quickly, efficiently, and seamlessly as possible. This way, they'll be able to capitalize on the speed of today's digital business environment.

Requisite Recoverability

No matter what events befall an organization, customers expect their data to be protected and secured at all times. Whether it's an unforeseen outage or downtime due to routine maintenance, the data simply must be intact and available without delay.
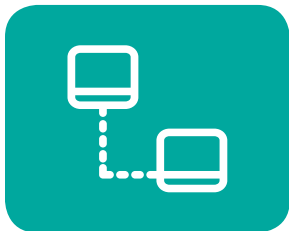
With the tolerance for downtime—as it applies to both the ability to access data and process transactions—effectively nonexistent, companies find themselves facing a harsh truth: lose the data, lose the brand reputation, and, eventually, lose the business.

But the pressure doesn't end there. With governments beginning to regulate data backup and recovery commitments, businesses now bear the burden of proving that every aspect of their operations can trust the systems that safeguard their data. Fail to provide this proof, and the resulting penalties can be crippling.

Data issues can—and will—happen. Mitigating their effects, and bringing critical applications and services back online as quickly as possible, can't be achieved by simply adding more hardware or mirrors. Instead, businesses need a new way to visualize, estimate, and report on their recoveries.

Getting there requires intelligent systems that back up changed data only, making it possible to execute recoveries far faster than traditional measures. To add a much-needed layer of confidence to the process, organizations need to run estimations and simulations that can show whether or not their efforts will uphold established RTOs and RPOs. And, they need to tightly coordinate all of the resources required to manage the process, so they can recover quickly at a lower total cost. Together, these capabilities form the foundation of an effective—and, more importantly, defensible—data backup and recovery program.

The Uptime Imperative

Customer expectations for data availability have never been more exacting. Simply ensuring data is online and uncompromised doesn't automatically create happy customers. They'll still demand access to this data when and where they want—and won't hesitate to find a more responsive experience if they don't get what they desire the first time.

Few organizations are able to overcome the negative effects customer dissatisfaction can have on their reputation and bottom line. All too often, however, business leaders' demands for constant uptime find themselves at odds with the pressures and responsibilities shouldered by IT.

Routine maintenance gives IT the opportunity to tune critical applications and services, while helping the business avoid inefficiencies and rising costs. But no matter how brief, carefully planned, or necessary the downtime may be, any drop in availability could lead to serious consequences. In the end, it becomes a painful paradox: the very same maintenance efforts that are so vital to driving the business forward often end up doing irrevocable damage to customer satisfaction, brand value, and revenue generation.

To avoid unneeded downtime—while minimizing the potential for outages and other unforeseen issues—organizations need to execute data reorganizations and other critical maintenance tasks without taking key systems offline, or making the data unavailable to the applications that need it.

This way, they'll avoid performance degradations, better uphold their SLA commitments, and enhance the customer experience. And by extracting key insights from their applications, businesses will be able to better plan for future performance demands, enabling them to anticipate customer needs before they even arise.

Managing Data, Maximizing Performance

That data onslaught isn't slowing down. To get out ahead of rapidly rising data volumes—as well as exacting customer expectations—companies must reevaluate their existing processes to see what may be restricting agility and hindering the customer experience. They need to make data and applications available around the clock, without neglecting maintenance schedules. And they must protect and secure that data, while making sure it can be quickly and completely recovered following an incident.

And that's why companies are turning to BMC.

Our intelligent automation and efficient processing empowers organizations to tap into the unused capabilities and power that already reside in their infrastructures, so they can deliver maximum performance and lasting availability—without adding resources or increasing complexity.

WITH BMC, COMPANIES CAN:

Turn the data deluge into an advantage:

By ensuring the integrity of structured and unstructured data, and using analytics to future-proof their operations, organizations can transform ever-growing data volumes into new sources of differentiation and revenue generation.
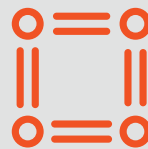
Provide auditable proof of recoverability and recovery times:

By simulating recovery tests and validating the results with their auditors, companies can design a recovery plan that delivers fast, predictable, SLA-compliant results.

Take the downtime out of performance management:

By diagnosing known and unknown issues, and tuning systems to account for these problems, businesses can realize the full value of performance management without ever bringing down their applications.

To learn more, please visit **our web site**.