

Enterprise data is your edge

*A security-rich, resilient environment to drive
analytic insights from your enterprise data*



Contents

- 2 What's your advantage?
- 3 Your enterprise resource systems and their wealth of data
- 4 Generate personalized customer interactions
- 5 Every copy of data has its own latency, cost and risk
- 6 Leverage your unique data to deliver an optimal customer experience
- 7 Better. Faster. Cheaper.
- 11 A new alternative
- 11 For more information
- 11 About the authors and contributors

What's your advantage?

Everyone wants an advantage. But, what exactly does that mean? Does an advantage imply an edge over your competitors, your partners, your customers, all of them? An edge is differentiation. An edge is insight. An edge is data, *your* data.

Enterprises spend billions of dollars to extract competitive advantages from publically available internet and social media data. These efforts have led to incredible insights and growth in new technology. They're fueling new ideas, new products, new companies and even completely new industries. Probably, your competitors, your partners and to an extent your customers use some of the same data resources you use. But, you also have unique data. You have data and potential insights that no one else has—or can have.

Every organization has its exclusive product portfolio, marketing, processes, personnel, relationships and corporate personality. This combination of corporate traits attracts a unique clientele. These relationships generate exclusive data about your clients that only you possess.

“The value of an idea lies in the using of it.”

— Thomas Edison

According to IBM studies, only 20 percent of the world's data is searchable. Anyone can get to that 20 percent. The real value lies in the 80 percent of the world's data behind the firewall that's *not* publically searchable.¹ Hidden value lies in decades of underwriting, pricing, customer experience and loan risk data that you have already collected. Leveraging that data gives you your advantage.

Your unique data is your unique advantage.

Your enterprise resource systems and their wealth of data

Today, there are tens to hundreds of operational application data sources and complex environments where data is sent back and forth between multiple systems at lightning speed. Many of us forget, however, that much of the data used in these systems originates within a centralized system. These systems support the core sales, inventory management and general ledgers of an organization. For many successful enterprises, these operational applications run on a mainframe computer.



Mainframe computers are used primarily by large organizations and originally performed centralized, critical bulk data processing for industries and governments. Today's mainframes have advanced and include exceptionally high reliability and stability, input and output alternatives, transaction processing, security and many features that make them essential to modern enterprises. Modern mainframes support new strategies to ensure even greater availability and reliability, no matter what the circumstances. Their secure, consistent, stable and reliable environments have enabled them to run uninterrupted for years.

Virtualization is key to mainframe capabilities and provides them with the ability to share physical resources and run many applications simultaneously, ensuring that mainframes are leveraged cost-effectively and efficiently. They also play a central role within analytic applications that require real-time insight. Mainframes can be viewed as an entire data center's worth of computing within a single *box*.

Operational applications run on the mainframe because of the mainframe's exemplary security, reliability and performance. It continues to be the most cost-effective, capable and reliable option available. For many enterprises, their mainframe computers collect vital operational application data and have done so for decades.

The applications that run the business have been revised, updated and extended to add more value. They have been modified to support new routes to market, such as the internet in the nineties and mobile applications today.

The mainframe is the most security rich and resilient platform available. And it's still collecting data, *your* data. Mainframe data is easily accessible for use in analytics and can help personalize interactions with customers.

What data is being collected, governed and secured on your mainframe? Think of how much value it has.

Generate personalized customer interactions

Forrester Research emphasizes that today's customers want celebrity treatment. They don't want to be treated just like everyone else; they want you to understand their needs, likes and dislikes.² A company's ability to do so depends on insights generated from the data collected on its customers. Leveraging insights generated from publically available internet and social media data can be valuable. However, using insights generated from your unique data, and combining those insights with new, engaging applications, helps to further personalize each interaction.

Engagement applications, those applications that the customers interact with, that behave the same way for everyone, do not deliver on the promise of personalized engagements. New engagement applications behave differently based on insights about the customer. These new engagement applications are often driven behind the scenes by output from machine learning applications.

Machine learning, combined with your unique data on IBM® Z® mainframes, can give you an advantage. Machine learning is a type of data-driven analytics that enables a computer to learn without being explicitly programmed. It helps data scientists:

- Identify patterns in historical data
- Build behavioral models from those patterns
- Use those models to predict behavior

Today's machine learning technology can analyze bigger, more complex data and deliver faster, more accurate results on a very large scale.

According to an Accenture study, seven out of 10 customers are willing to take banking, financial services and insurance advice from machine learning-based robo-advisors.³

Are you driving personalized experiences from your unique data?

Every copy of data has its own latency, cost and risk

If, as Thomas Edison suggested, the true value of an idea lies in the *using* of it, then to get the true value of data, it must be used, as well. And organizations are doing just that.

New engagement and analytics teams have expanded the use of this extremely valuable and sensitive data that originates on the mainframe. Organizations have spent a great deal of time and effort governing, managing and curating that operational application system data to make sure it's usable, accurate and secure. They have also spent a great deal of time and effort moving that data to other less secure, less resilient platforms for analysis and customer engagement.

The amount of data being moved to other platforms is growing exponentially. At the same time, the cost, latency and security-breach risk of moving and protecting this valuable data has proven untenable. These concerns are encouraging many organizations to consider alternative approaches and keep the data and new engagement applications on a security-rich, resilient mainframe platform—IBM Z.

According to data breach statistics at breachlevelindex.com, of the more than 9 billion records lost or stolen since 2013, only 4 percent of breaches were secure breaches where encryption was used and the stolen data was rendered useless.⁴

*How much data are you moving from your mainframe?
What risks do you take?*

Leverage your unique data to deliver an optimal customer experience

Real-time analytics and engagement applications now demand the same security and resiliency as the operational applications they complement. These safeguards help ensure that opportunities are not lost.

IBM Z is a purpose-built mainframe environment that includes hardware, software and an operating system. Running analytic applications on the same platform as the operational applications helps ensure the same security-rich environment and resiliency for both. Resiliency is essential to deliver real-time insight and value within every single customer transaction.

IBM Z is extremely resilient. Resiliency implies the capacity to recover quickly from difficulties. IBM Z excels at meeting service-level agreements (SLAs) under the most demanding circumstances. The mainframe can also adapt to the changing marketplace with new demands, capabilities and workloads. One of those new workloads is cognitive services.



In the financial industry, customer service agents typically spend their time on a small group of premier clients. This process is referred to as *wealth management*. One of the activities of wealth management is to provide asset allocation services.

Today, applications can automate this now routine recommendation process and offer those same services more broadly at significantly lower cost, driving greater productivity. Cognitive applications now handle even the less-routine situations. This method allows customer service agents to provide superior personalized experiences, so customers feel like celebrities.

The unique data residing on an IBM Z mainframe can be accessed in real time in support of cognitive applications. There's an overall trend to move compute to the data that's referred to as *data gravity*. IBM makes it possible for customers to cost-effectively keep data in place while also using it for analytics and machine learning to deliver a unique customer experience.

Do you use your core data and mainframe's resiliency to provide a unique customer experience?

Better. Faster. Cheaper.

Many of us have heard the old IT saying, "You can have it cheap, fast or good... chose any two." But is it true? When building analytics applications, can you possibly have all three qualities? Can an application be less expensive, faster and better? We briefly discussed resiliency earlier. Here are some of the other factors that should be considered when platforming an analytics application. Besides resiliency, an organization should consider data quality and governance, IT server sprawl, complexity, energy consumption, security, meeting SLAs, data movement and consistency, and cost.

IBM experience shows that many companies see thousands of peak transactions per second. A reasonably conservative number would be 5,000 transactions per second. Today's mainframes provide five nines availability, 99.999 percent, versus distributed approaches that may deliver as much as four nines availability, 99.99 percent. That's a potential difference of 38,880 missed opportunities between these two approaches within a single day.

Customers often experience data quality and governance issues associated with downstream data applications. These issues are associated with, among other factors, the large number of server platforms and the resulting system complexity. Moving data from platform to platform, often without validation check points, can greatly affect data quality. Traditional analytic application implementations require many servers that perform separate functions for the analytics application.

For example, there might be an extract, transform and load (ETL) server, an application server, an HTTP server, a statistical platform or data mart. The resulting infrastructure is far more complex than *just a couple of servers*. Each category of server—ETL, application, HTTP, database, analytical engine and more—requires several server platforms. Each category of server may also require a production, development, test, quality assurance and disaster recovery environment. The result is IT complexity and server sprawl. As data is moved from its original platform, costs are incurred, latency and security-breach risks are introduced, and the insights become out of date.

IBM Z is a self-contained system where vital, sensitive data originates. Data can be leveraged in a security-rich environment in place to deliver the most current insight at the point of customer interaction.

Whether stand-alone, within a grid or virtualized, servers must have their own power supplies, memory, CPU, I/O and network interfaces. Each of these servers must communicate with one another; they are not self-contained systems. They need routers, switches and external hubs.

The cooperative systems within a networked, distributed environment can't compete with a mainframe when it comes to energy efficiency. The mainframe has a single, but redundant, power supply. All the components are designed for energy efficiency and virtualization. The energy consumption of a mainframe is significantly less than corresponding distributed architectures, providing comparable millions of instructions per second (MIPS). Mainframes use energy-efficient water cooling to cool their CPUs. Less energy consumed means less energy wasted and less heat generated. This architecture also means less data center air conditioning. All these factors contribute to a smaller carbon footprint for analytic applications deployed on an IBM Z mainframe. For organizations concerned about their environmental impact, the IBM Z mainframe delivers a distinct advantage.

With a large number of servers, rolling out new applications requires significant testing. The more complex an environment is the more expensive it is to test. Testing bills can reach into the millions.

All this complexity leads to the employment of tens to hundreds of people just to keep the distributed systems up and running. This amount of manpower used solely on ensuring systems uptime and not adding organizational value wastes significant IT funds.

Security is one of the most important factors to address when considering a platform for applications. Mainframes are known for their security. This is not only an industry belief. It's backed up with decades of IBM investments in hardware and software, security control, auditing and encryption. Most banks, financial services and insurance companies depend on the mainframe's security to guard their most sensitive data assets.

A recent study conducted by Ponemon Institute and sponsored by IBM estimated the cost of data breaches to be USD 172 per record.⁵ For more information on security and regulatory compliance read the [white paper](#), "Is your business risking a USD 100 million loss? How can new technologies help protect your data and your bottom line?"

“According to many media reports, [company name removed] will soon deplete its USD 100 million cyber-insurance coverage just to notify the victims and provide free identity-theft and credit monitoring. Ponemon Research conducts annual studies on the cost of a data breach, which consistently hovers around USD 200 per record. But that number doesn't include the hard-to-calculate costs like reputational repercussions, business distraction, class-action lawsuits and regulatory fines.”⁶

— CFO magazine

Meeting analytic application SLAs is a matter of corporate responsibility and good business. Not meeting SLAs can result in unexpected penalties. Application server and data repository performance are essential to ensuring mobile and internet applications load within SLAs. IBM Z workload management features help ensure that SLAs can be met—even under demanding circumstances.

According to Kissmetrics, 47 percent of visitors expect a website to load in less than two seconds, and 40 percent of visitors will leave the website if the loading process takes more than three seconds.⁷

Analytic repositories cannot be updated on a weekly, daily or even an hourly basis. Today's expectations are that decisions are being made with the freshest data. But data freshness is only one part of a bigger picture. Data consistency is a challenge, as well—and core to making appropriate decisions. The concept of eventual consistency has recently been promoted for use with analytic applications. Eventual consistency is a common term used with open source distributed servers. Since open source distributed servers scale horizontally instead of vertically, they must synchronize data among themselves over time. This process implies that an update, an insert or a delete might not be visible on the other servers immediately.

This delay would certainly be disastrous for some applications, such as a banking application. But with analytics applications, the argument is that there's little impact. But what happens if your analytic application depends on data that could be coming from more than one server? Will you eventually be impacted by data inconsistency? Will it be costly?

An IBM study shows that over a one-year period, each terabyte of data per day that's extracted, transferred off platform and loaded on external systems costs USD 2.5 million per year or USD 10 million over a four-year period.⁸

Traditional approaches attempt to deliver resiliency by duplicating the physical infrastructure. This approach may reduce the potential for outages but duplicates the cost and significantly increases IT complexity. Data duplication can result in other challenges, such as greater security risk and data inconsistency.

Moving data off the original platform is not necessarily better, faster or cheaper.

A new alternative

When your high-value, sensitive data originates on an IBM Z mainframe, it makes sense to leverage that data in place. You minimize costly data movement and maintain a high level of data governance, encryption, and resilience in a security-rich environment.

Deploying enterprise analytics and machine learning on IBM Z, enables you to access that data in place and combine it with other sources of information, such as structured and unstructured data from other systems. Using IBM Z for real-time insight allows you to benefit from a security-rich environment with resiliency that's on par with your business-critical transactional applications.

Using real-time data to make decisions, you can:

- Identify risk more quickly
- Improve your customers' experience
- Increase revenue by better targeting cross-sell and upsell opportunities
- Reduce operational costs
- Deliver greater organizational productivity that results in greater agility

Greater agility is all about making the right decision at the right time. Use IBM z Analytics and machine learning to provide real-time insights that convert your unique enterprise data into more business opportunities.

You *have* an edge. It's time to *use* it.

For more information

To learn more about IBM z Analytics solutions, including IBM Machine Learning for z/OS®, IBM Db2® for z/OS, IBM Db2 Analytics Accelerator for z/OS, IBM QMF™, IBM Data Virtualization Manager for z/OS and Db2 Tools, please contact your IBM representative or IBM Business Partner, or visit: ibm.com/analytics/us/en/enterprise.

About the authors and contributors

The following people contributed to this thought leadership white paper:

Jonathan Sloan, IBM Analytics Solution Architect, z Analytics

Annette Zawacki, IBM Portfolio Marketing Manager, z Analytics



© Copyright IBM Corporation 2017

IBM Corporation
IBM Analytics
Route 100
Somers, NY 10589

Produced in the United States of America
November 2017

IBM, the IBM logo, ibm.com, Db2, IBM Z, QFM, Redbooks, and z/OS are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at www.ibm.com/legal/copytrade.shtml.

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED “AS IS” WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

Statement of Good Security Practices: IT system security involves protecting systems and information through prevention, detection and response to improper access from within and outside your enterprise. Improper access can result in information being altered, destroyed, misappropriated or misused or can result in damage to or misuse of your systems, including for use in attacks on others. No IT system or product should be considered completely secure and no single product, service or security measure can be completely effective in preventing improper use or access. IBM systems, products and services are designed to be part of a lawful, comprehensive security approach, which will necessarily involve additional operational procedures, and may require other systems, products or services to be most effective. IBM DOES NOT WARRANT THAT ANY SYSTEMS, PRODUCTS OR SERVICES ARE IMMUNE FROM, OR WILL MAKE YOUR ENTERPRISE IMMUNE FROM, THE MALICIOUS OR ILLEGAL CONDUCT OF ANY PARTY.

- 1 CNBC Transcript, *IBM Chairman, President & CEO Ginni Rometty on CNBC's "Mad Money with Jim Cramer,"* 21 June 2017, <https://www.cnbc.com/2017/06/21/cnbc.html>
- 2 IBM Webcast: *Machine Learning, Analytics and Open Source for the Enterprise—An Analysts Perspective from Forrester,* 20 September 2017, <https://event.on24.com/wcc/r/1479375/EAC5643953CCD8C7128AC317AB190D07>
- 3 *The Rise of Robo-Advice: Changing the Concept of Wealth Management,* Accenture, https://www.accenture.com/_acnmedia/PDF-2/Accenture-Wealth-Management-Rise-of-Robo-Advice.pdf
- 4 Breach Level Index, <http://breachlevelindex.com/>
- 5 Ponemon Institute's *2017 Cost of Data Breach Study: Global Overview,* <https://www.ibm.com/security/infographics/data-breach/>
- 6 *Calculating the Colossal Cost of a Data Breach,* Mary A. Chatput, CFO Magazine, 24 March 2014, <http://www2.cfo.com/data-security/2015/03/calculating-colossal-cost-data-breach/>
- 7 *Kissmetrics Blog: A blog about analytics, marketing and testing,* <https://blog.kissmetrics.com/loading-time/>
- 8 *Reducing Data Movement Costs in z Systems Environments:* An IBM Redbooks® Point-of-View publication, <http://www.redbooks.ibm.com/abstracts/redp5217.html?Open>



Please Recycle