

September 2020



Lower the cost of new business opportunity with IBM Db2 for z/OS Data Gate



New and existing business demands timely data delivery

For almost any company, building demand for new business opportunity is a key to success. This means rapid, iterative application development accessing and presenting data from enterprise data sources that customers need, want and find compelling. This can help an organization to differentiate themselves from competitive alternatives. Data delivery can be essential to new business development.

Whether displaying product inventory or checking account balances, providing timely and accurate data, on demand, is key to supporting customer expectations and can enable organizations to attract new clientele. Today's commerce requires that lookups or queries can be done anytime, and from anywhere, on mobile phones and other devices.

This type of activity can consume significant IT resources, resulting in unpredictable and at times, substantial charges. To offset the financial impact of rising inquiry activity, IT organizations are seeking lower-cost data access solutions.

IBM Db2® for z/OS® Data Gate

IBM Db2 for z/OS Data Gate (Db2 Data Gate) allows users with existing Db2 for z/OS systems to synchronize a copy of their data in real time for read-only access within IBM Cloud Pak for Data on another platform. This allows users to readily access the most current data and allows IT organizations to minimize the workload and potential financial impact to IBM Z®. Furthermore, Db2 Data Gate's Integrated Synchronization leverages IBM z Integrated Information Processor (zIIP) technology to offload Central Processor (CP) workload. Db2 Data Gate delivers real-time data to support existing and new business opportunities.

Db2 Data Gate's ability to provide the most current data facilitates new business opportunity, is easier to implement and is more cost-effective than Do It Yourself (DIY) approaches. Db2 Data Gate is a fully integrated solution, simplifying deployment and IT operations. This paper examines the cost of Db2 Data Gate versus a DIY data extraction and synchronization approach to quantify the benefits of using Db2 Data Gate technology.

Comparing Db2 Data Gate to DIY solutions

Although DIY solutions may use commercially available products, they still require an initial effort for production use. The DIY approach will typically require customization to meet an organization's needs, which will entail testing, integration with other workloads, deployment, and maintenance of the overall solution once deployed.

Developing and deploying a new solution involves considerable IT programming and administrative effort. It also involves commitment of valuable IT resources (software, hardware, data center infrastructure) associated with the solution. At times, the programming and administrative effort and associated IT resource costs can exceed the anticipated benefit of the solution. Db2 Data Gate addresses these issues by providing a complete, prepackaged solution that is simple to deploy and maintain.

To examine the cost/benefits of Db2 Data Gate versus a DIY solution, two Db2 Data Gate use cases, a transactional and an analytical use case, were evaluated. For the comparison, the models assumed incremental data load methods from an existing Db2 for z/OS database only. Bulk data load methods are not considered as they would generate large capacity requirements and would impede data currency.

The IBM IT Economics team used IBM pricing estimates and IT data, from client engagements associated with previous IBM IT Economics assessments, to calculate five-year total cost of ownership (TCO) models for each use case. The models included data such as software, labor, hardware, networking, space, and energy costs.

Transactional use case

The transactional use case examines the costs of a scenario in which users access their transactional data within a database copy optimized for a read-only transactional like workload. Examples include viewing checking account balances, insurance policy expiration dates, shipping-delivery status, etc.

For the Db2 Data Gate transactional TCO model, we used Db2 Data Gate, its pre-requisite IBM Cloud Pak™ for Data Enterprise Edition (which includes Db2 Warehouse), and IBM Cloud Pak for Data Db2, which provides a relational database optimized for transactional type applications. The hardware was comprised of two Linux® x86 servers that connect to Db2 v12 on z/OS v2.3 installed on an existing IBM Z system with 4,500 Million Instructions Per Second (MIPS) base capacity. To process the Db2 Data Gate integrated synchronization load, we included the purchase of an zIIP.

Both IBM Cloud Pak for Data and Db2 Data Gate use licensing models based on Virtual Processor Cores (VPC) and the x86 environment was modeled for a sample installation with vCPUs for Db2 Data Gate and 19 vCPUs for IBM Cloud Pak for Data¹. Software license list price costs for both products were included in year one, and subscription and support were included for five years. The purchase cost of x86 hardware was based on vendor list prices with a 30% discount. Both environments assumed no workload growth over the five-year period.

Analytical use case

The analytical use case examines costs of a data warehouse in which an organization can perform analytical evaluations of their data (improve customer satisfaction, adapt to market and risk patterns, report package shipping/orders/payments, etc.).

Db2 Data Gate: For the Db2 Data Gate analytical TCO model, we used Db2 Data Gate, and its pre-requisite IBM Cloud Pak for Data Enterprise Edition (which includes Db2 Warehouse), on two Linux x86 servers that connect to Db2 v12 on z/OS v2.3 installed on an existing IBM Z

¹ Db2 Data Gate uses RHEL CoreOS core provided by IBM Cloud Pak for Data. The DIY solution is also containerized with licensing that includes a Linux operating system. The TCO model assumes list prices for the IBM Db2 for z/OS Data Gate, IBM Cloud Pak for Data Enterprise Edition (which includes Db2 Warehouse), and IBM Cloud Pak for Data Db2, and the competitive replication product.

system with 4,500 Million Instructions Per Second (MIPS) base capacity. To process the Db2 Data Gate integrated synchronization load, we included the purchase of an additional zIIP.

Both IBM Cloud Pak for Data and Db2 Data Gate use licensing models based on Virtual Processor Cores (VPC) and the x86 environment was modeled for a sample installation with vCPUs for Db2 Data Gate and 17 vCPUs for IBM Cloud Pak for Data² (two vCPUs less than the transactional use case). Software license list price costs for both products were included in year one, and subscription and support were included for five years. The purchase cost of x86 hardware was based on vendor list prices with a 30% discount. Both environments assumed no workload growth over the five-year period.

DIY: For the DIY scenario in the cost case we used a competitive replication product with the associated license costs. We modeled a comparable DIY environment to the Db2 Data Gate environment by including IBM Cloud Pak for Data (which includes Db2 Warehouse) and the same IBM Z environment (Db2 v12 running on z/OS v2.3 on an IBM Z system with 4,500 MIPS base capacity). To process the replication load, we estimated an additional 400 MIPS of Central Processor (CP) capacity instead of purchasing an additional zIIP. The competitive replication product on x86 was modeled to use the same capacity as the Db2 Data Gate x86 environment. No other additional commercial product costs were included in the DIY model in order to calculate with minimal DIY costs.

Db2 Data Gate brings savings

Cost analysis for both use cases showed that Db2 Data Gate can lower IT costs.

In the modeled read-only transactional scenario, Db2 Data Gate delivers a 37% lower total cost of ownership than the comparable DIY solution over five years³.

In the modeled read-only analytical scenario, Db2 Data Gate delivers a 60% lower total cost of ownership than the comparable DIY solution over five years⁴.

Software cost for both models was the most significant cost driver. For the DIY approach, another notable cost in both use case models was the implementation and maintenance (labor) cost. In contrast, other cost components such as energy, space, and networking were relatively inconsequential.

² Db2 Data Gate uses RHEL CoreOS core provided by IBM Cloud Pak for Data. The DIY solution is also containerized with licensing that includes a Linux operating system. The TCO model assumes list prices for the IBM Db2 for z/OS Data Gate, IBM Cloud Pak for Data Enterprise Edition (which includes Db2 Warehouse), and the competitive replication product.

³ An IBM IT Economics model was used to examine the cost of IBM Db2 for z/OS Data Gate versus a DIY change data capture solution. The model compared total cost of ownership over five years for hardware, software, people, networking, floor space, and energy costs for a transactional (read only) use case. Partial data load methods execute to and from an z/OS Db2 database using Db2 v12 on z/OS v2.3 installed on an existing IBM Z system with 4,500 MIPS base capacity and one zIIP to process the Db2 Data Gate integrated synchronization load. The model uses the same IBM Z environment for the CDC DIY solution with an additional 400 MIPS of CP capacity was calculated to process the CDC load for the DIY solution. Both scenarios include 19 vCPUs for IBM Cloud Pak for Data Db2 on two 28-core x86 servers. The model uses list pricing for IBM Db2 for z/OS Data Gate, IBM Cloud Pak for Data Db2, and the competitive CDC product from site URL. Results will vary depending on the size of the IBM Z base MIPS environment, the size of the additional CDC MIPS load, and hardware and software pricing. Contact the IBM IT Economics team at IT.Economics@us.ibm.com for additional information.

⁴ An IBM IT Economics model was used to examine the cost of IBM Db2 for z/OS Data Gate versus a DIY change data capture solution. The model compared total cost of ownership over five years for hardware, software, people, networking, floor space, and energy costs for a data warehouse (simple queries) use case. Partial data load methods execute to and from an z/OS Db2 database using Db2 v12 on z/OS v2.3 installed on an existing IBM Z system with 4,500 MIPS base capacity and one zIIP to process the Db2 Data Gate integrated synchronization load. The model uses the same IBM Z environment for the CDC DIY solution with an additional 400 MIPS of CP capacity was calculated to process the CDC load for the DIY solution. Both scenarios include 17 vCPUs for IBM Cloud Pak for Data on two 28-core x86 servers. The model uses list pricing for IBM Db2 for z/OS Data Gate, IBM Cloud Pak for Data, and the competitive CDC product. Results will vary depending on the size of the IBM Z base MIPS environment, the size of the additional CDC MIPS load, and hardware and software pricing. Contact the IBM IT Economics team at IT.Economics@us.ibm.com for additional information.

The TCO model for the transactional use case below (Figure 1) shows significant savings for Db2 Data Gate, 37%, as compared to the DIY approach. The difference in cost is primarily a result of the additional labor and software costs associated with the DIY approach. Labor costs in the model are driven not only by the initial development effort of the DIY solution but also by ongoing maintenance (software integration, patches, upgrades, etc. described below).

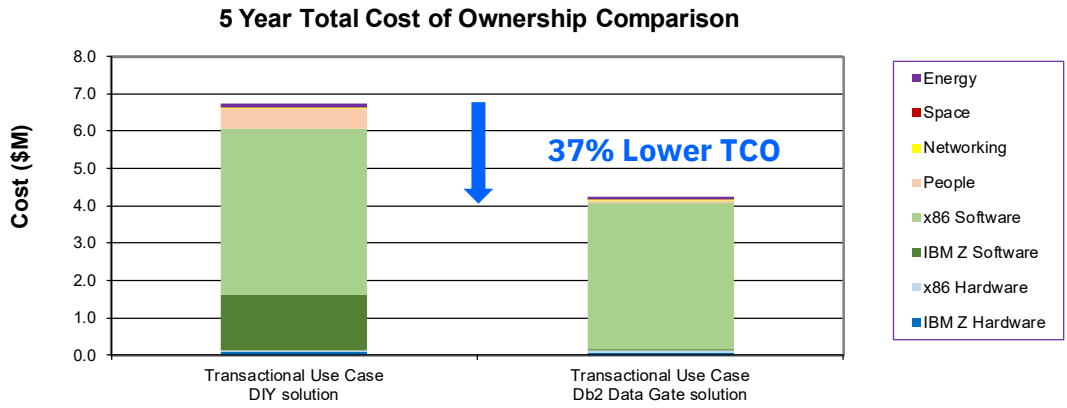


Figure 1: Five year total cost of ownership comparison by cost components between IBM Db2 for z/OS Data Gate and DIY solution for the transactional use case.

Figure 2 below shows the breakout of costs for the TCO analytical use case. The cost difference between the DIY and Db2 Data Gate solutions is mainly driven by the reduction of associated software charges and labor costs, resulting in a 60% lower TCO for Db2 Data Gate when deployed in support of the analytical use case.

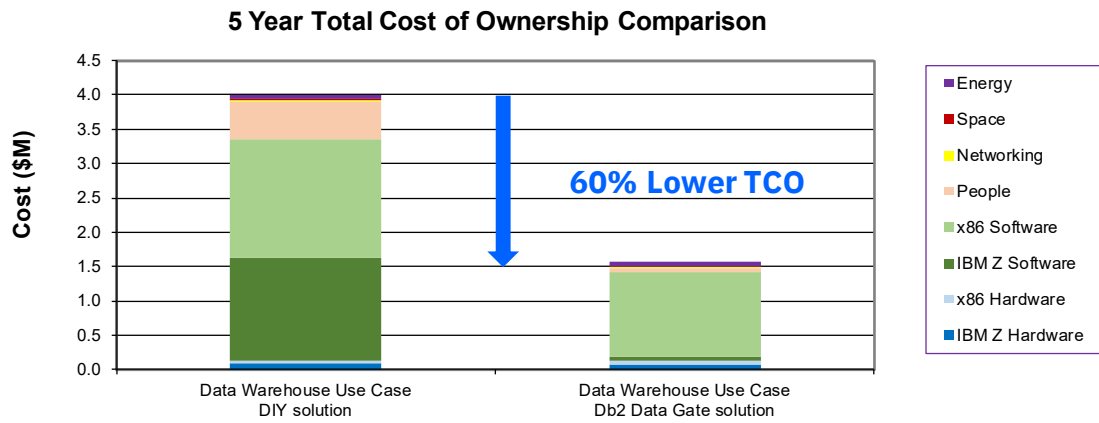


Figure 2: Five year total cost of ownership comparison by cost components between IBM Db2 for z/OS Data Gate and DIY solution for the analytical use case

Figure 3 shows the impact of cost over time for the analytical use case. In this accumulated TCO graph the DIY solution not only foresaw higher investment in year one for software licenses and implementation, but also exhibited a steeper cost slope for subsequent years, due to significant software costs and administration (people) costs for maintenance of the solution.⁵

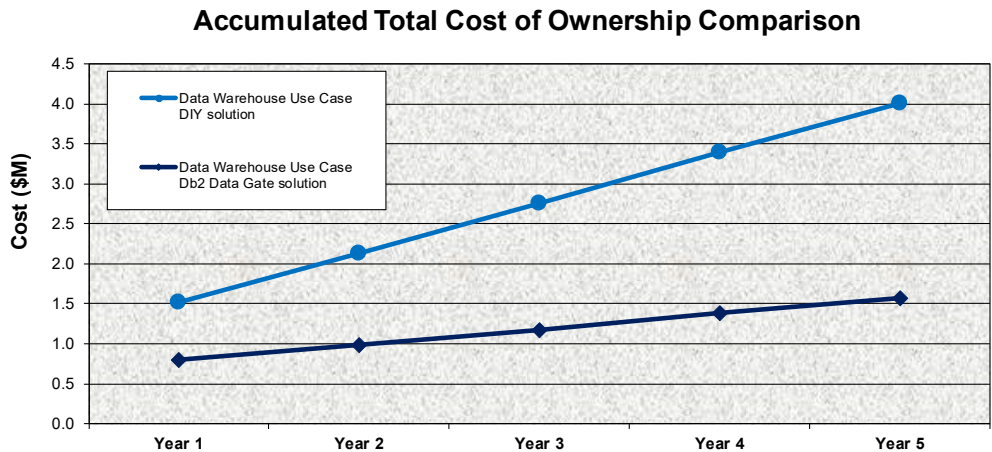


Figure 3: Comparison of accumulated costs over five years for IBM Db2 for z/OS Data Gate versus DIY solution for the analytical use case

Lower implementation costs and maintenance for Db2 Data Gate

DIY solutions require considerable effort to integrate and test different components, and they typically require more complex application logic to achieve data currency. They also require significant maintenance effort post deployment since they lack the benefit of being a pre-tested, integrated solution by design. Using labor estimates based on similar DIY projects, the model forecasted 120 full time equivalent (FTE) people days for coding, implementation and deployment, and one FTE per year for maintenance once the solution was deployed (see Table 2 for details on FTE efforts).

For the integrated Db2 Data Gate solution the effort was found to be considerably lower, with an estimate of five FTE days for implementation based on Db2 Data Gate installation user experiences⁶. For maintenance the model estimates an effort of one tenth of an FTE per year for Db2 Data Gate based on administration effort for comparable database solutions.

⁵ People effort estimates are based on early user feedback and comparable internal development projects. People cost is based on an effort of 120 FTE days (6 months) and 1 FTE per year for maintenance for the DIY implementation versus 5 FTE days for implementation and 1/10 FTE per year for maintenance for the Db2 Data Gate solution using an annual FTE rate of \$100,000.

⁶ During early user program clients reported an average of 5 FTE days to install Db2 Data Gate.

| | DIY Solution | | IBM Db2 for z/OS Data Gate | |
|------------------------|----------------------------|-------------------------------|----------------------------|-------------------------------|
| | One-time effort (FTE days) | Ongoing effort (FTE per year) | One-time effort (FTE days) | Ongoing effort (FTE per year) |
| Planning/Design | 20 | | NA | |
| Code development | 40 | | NA | |
| Testing | 20 | | NA | |
| Integration/Deployment | 40 | | 5 | |
| Code Updates | | 0.4 | | NA |
| Error handling/fixes | | 0.3 | | NA |
| User support | | 0.2 | | 0.1 |
| Monitoring | | 0.1 | | NA |
| Total effort | 120 FTE days | 1 FTE per year | 5 FTE days | 0.1 FTE per year |

Table 1: Full time equivalent people effort for implementation and maintenance of a DIY versus a Db2 Data Gate solution

Fewer software charges for Db2 Data Gate

DIY solutions lack the advantage of Db2 Data Gate's integrated synchronization feature that enables general processor capacity to be offloaded to zIIPs. DIY solutions consume standard IBM Z processor capacity, thereby increasing MIPS consumption on z/OS, which results in higher monthly license charges (MLC).

Integrated synchronization technology in Db2 Data Gate enables up to 96% of log read work to be zIIP eligible⁷. This allows Db2 Data Gate to offload most of its extraction and synchronization work to zIIPs so that general processor capacity and MLC can be minimally impacted.⁸

Db2 Data Gate brings business value

Db2 Data Gate offers an integrated and cost-effective solution for the extraction and synchronization of z/OS Db2 data so that businesses can easily and affordably give access to essential data. Offloading eligible log read work to zIIPs allows IT organizations to respond to increasing data access requests while minimizing software charges.

With Db2 Data Gate, business applications can access data originating in their systems of record on IBM Z without having to access to the core Db2 for z/OS system. This broadens the aperture for both new and existing client opportunities through the use of many read-only data use cases.

- Mobile device queries (insurance policy information, price comparisons)
- Seasonal activity peaks (vacation planning, holiday shopping)
- Annual/monthly payments (utility bills, loan status)
- 24x7 virtual support (how to information, help documents)
- Unlimited data checks (bank accounts, financial reports)

⁷ Results extrapolated from measurements done on z14 with 6 GCPs and 2 ZIIPs, 96G Memory, 10Gbit network to Accelerator with z/OS 02.02.00, Db2 for z/OS v12, CDC for z/OS V10R2M1. Test execution was comprised of reading a Db2 transaction log pre-populated with inserts and updates against 500 tables. The average row length reported by Db2 was 59 bytes compressed. zIIP eligibility is based on CPU consumption of all address spaces that took part in the log read activity (primarily DDFWORK and DIST were available for offload) during the test run. Result may vary.

⁸ MLC charges can also be minimized through the use of Tailored Fit Pricing. <https://www.ibm.com/it-infrastructure/z/software/pricing-tailored-fit>

Client satisfaction and increased market share

Enabling your clients to access the data they request on their terms increases client satisfaction. Clients can decide how, what, and when to access their data. Simple requests that might have required interaction with a service center can be resolved on-line, at your client's convenience.

Unlimited access to data can also accelerate business opportunity. Enabling your clients to access their data simplifies business operations. It allows your personnel to focus on more complex requests and to pursue new client acquisition for business revenue growth.

Reduce data query costs

If your IT organization is interested in exploring how to increase business opportunity while minimizing cost related to data queries, contact the IBM IT Economics team at IT.Economics@us.ibm.com . Ask for a no-charge, customized Db2 Data Gate assessment to help you quantify costs and benefits for your business.

About the authors



Ingo Aller is an IBM IT Economics Consultant working with clients primarily in Europe, the Middle East and Africa to optimize IT solutions. Ingo has 23 years of IBM experience in the areas of microprocessor development, technical sales, and IT Economics consulting. He is an author and coauthor of numerous technical papers in different technical fields, and holds two patents. During his career in IBM he worked for two years in the USA, and for three years in Bangalore, India, where he established a microprocessor design team.



Tilman Wagner is a consultant on the IT Economics Consulting and Research team, focusing on total cost of ownership and business value assessments for clients in Europe, the Middle East and Africa. Tilman joined IBM in 1997 as Field Design Engineer in Chip Design and spent two years in the IBM Silicon Valley Office at Mountain View, CA, USA where he supported numerous major IBM customers on network processor design projects. In 1999 he returned to Europe and became the Engagement Manager and Chip Timing Expert for development projects in Europe, the Middle East and Africa where he led multi-national teams across North America, Europe and Asia Pacific. In 2008 Tilman became Chip Timing Lead in IBM Boeblingen Lab within IBM Processor Hardware Development for IBM Z Processor Development before changing to his current role in 2012.



Susan Proietti Conti, PMP® is an IBM Executive Project Manager and Program Director for the WW IBM IT Economics Consulting and Research organization. She manages IT Economics projects and helps clients leverage IT Economics to increase IT efficiencies and reduce costs.



^(c) Copyright IBM Corporation 2020
IBM Corporation
New Orchard Road
Armonk, NY 10504
U.S.A.
09/20

IBM, the IBM logo and IBM Z, IBM Cloud Paks, Db2 and z/OS are trademarks or registered trademarks of the International Business Machines Corporation. A current list of IBM trademarks is available on the Web at <https://www.ibm.com/legal/us/en/copytrade.shtml>, and select third party trademarks that might be referenced in this document is available at https://www.ibm.com/legal/us/en/copytrade.shtml#section_4.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

InfiniBand and InfiniBand Trade Association are registered trademarks of the InfiniBand Trade Association.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

The registered trademark Linux[®] is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

OpenStack is a trademark of OpenStack LLC. The OpenStack trademark policy is available on the [OpenStack website](#).

Red Hat[®], JBoss[®], OpenShift[®], Fedora[®], Hibernate[®], Ansible[®], CloudForms[®], RHCA[®], RHCE[®], RHCSA[®], Ceph[®], and Gluster[®] are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

RStudio[®], the RStudio logo and Shiny[®] are registered trademarks of RStudio, Inc.

TEALEAF is a registered trademark of Tealeaf, an IBM Company.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Worklight is a trademark or registered trademark of Worklight, an IBM Company.

Zowe[™], the Zowe[™] logo and the Open Mainframe Project[™] are trademarks of The Linux Foundation.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice and represent goals and objectives only.

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. It is the user's responsibility to evaluate and verify the operation of any other products or programs with IBM products and programs.

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.

34032134USEN-01