

# Strategies for **Effective Data Management** on Db2 for i and IBM Power Systems Data Migration

- IBM's Sarah Mackenzie on the importance of quality data and effective security for Db2 for i data management
- Richie Palma of iTech Solutions Group on IBM Power Systems migration considerations, and what migration strategies to use based on your business needs

**TechChannel**



---

**3** Simple Steps to Ensure Effective Data Management on Db2 for i

---

**8** Why Having a Backup and Recovery Plan Improves IBM i Vitality

---

**9** IBM Power Systems Data Migration Considerations and Strategies

---

**15** Accelerating Clinical Trials With Big Data Analytics

## Effective Data Management and Migration Planning



As a valuable proprietary asset that can be increasingly leveraged for direct business insights, data is the lifeblood of many organizations. With this in mind, effective data management and data migration strategies are of the utmost importance. While there isn't a "one-size-fits-all" approach here, knowing your business needs and system demands, and leveraging tools like Db2 for i will help with the planning process.

When planning for an IBM Power Systems update or migration, for example, you should take inventory of existing physical and virtual environments you will migrate, and what new environments you might need. A solid backup plan is also critical—whether you migrate with replication, LPM or SAN to SAN replication.

Db2 for i also offers several built-in solutions to help streamline data management, e.g., Db2 Web Query, the security audit journal, authority collection, etc.

In this TechChannel e-book, IBM's Sarah Mackenzie explains the importance of quality data and effective security for Db2 for i data management; and iTech Solution Group's Richie Palma highlights IBM Power Systems data migration best practices to keep in mind for successful migration projects.

**Keelia Estrada Moeller, Senior Editor**

# Simple Steps to Ensure Effective Data Management on Db2 for i

IBM's Sarah Mackenzie on the importance of quality data and effective security for Db2 for i data management

BY SARAH MACKENZIE

**D**ata is the lifeblood of the database. However, because of the large amount of data received every day, managing your data can be difficult. Successfully managing your data goes beyond just acquiring and storing large amounts of data. Effective data management also requires you to ensure your data is both of the highest quality and ever secure.

It's easy to overlook some of the built-in features that Db2 for i offers to help you with data management. Here are a few easily implemented features that you already own to help you better manage your data.

## Ensuring Quality Data

The first step to successful data management is ensuring your data is of high quality. Quality data is complete, accurate, consistent and normalized.

### *Constraints*


Implementing constraints to enforce your business rules is a simple solution for enforcing consistency and accuracy.

Primary key constraints can be added to ensure the uniqueness of a business element, such as an order number or a customer identification number. For example, adding a primary key constraint to a customer identification number in a CUSTOMER table ensures you don't have customers with the same identification number leading to errors like billing the wrong customer.

Referential constraints, on the other hand, allow you to enforce that relationship between tables. For example, say you have an ORDERS table containing a row for each order placed and this table has a column containing the customer identification number of the customer who placed the order. You need to make sure any customer ID inserted into the ORDERS table references a valid customer in the CUSTOMER table. A referential constraint relationship between the customer ID column in the ORDERS table and the customer ID column in the CUSTOMER table will establish this protection.

The third type of constraint, check constraints, can be added to ensure data inserted or updated in a column meets your business logic. An example of logic that could be enforced using a check constraint is guaranteeing a customer rating is between one and ten.

The beauty of adding constraints is the database itself is enforcing your business logic. If the value isn't valid according to the rules of the constraint, the insert or update will simply fail. You don't have to try to enforce





the rule across all your different interfaces—the database will do it for you.

### *Data types*

Another simple solution for ensuring data accuracy is confirming that your data is represented using the correct data types. For example, are you using a character data type to represent a timestamp instead of using the timestamp data type? If you're using a character data type, any characters can be inserted, including an invalid timestamp. By using the correct data type, you will put the burden on the database to ensure that a real timestamp is inserted. Preventing data type errors can ensure your data is accurate.

### *Triggers*

It's also important to analyze how you implement your business rules. Does every application have to ensure data inserted into the database conforms to a certain format? For example, you may require a column always contain



## TechChannel Webinars

### **Free, Expert-Led Education; Live or On-Demand**

Want to keep up on the latest IT trends? Wondering what your peers are doing? Need a solution for a business challenge? With many in-person conferences and expos on hold, there's no better way than complimentary, expert-led TechChannel Webinars to get the latest on all things tech. Attend now and interact with the technologists; ask the questions that matter most to you.

[View all live and on-demand webinars](#)

a company name in uppercase characters. This can lead to consistency problems if some applications enforce the logic to uppercase the company name and others do not.

Consider moving your logic into a single place all of your applications can call. Even better, if you want to ensure the data being inserted into your table conforms to certain business rules, like our capitalization rule, move the logic to capitalize the company name into a before insert trigger. This ensures no matter how your data is inserted, whether it be through a known application or an ad hoc driven insert, that business logic will be implemented.

### *Db2 Web Query*

Additionally, if you have existing data you want to ensure the quality of, products like Db2 Web Query can be used to automatically cleanse and consolidate your data using the integrated Extract, Transformation and Load (ETL) DataMigrator feature.

## **Effective Data Security**

It's critically important your data is secured. Data breaches cost companies millions of dollars along with customer trust. Additionally, many countries have placed the rights to data privacy into law. Understanding and defining who can access what data is a critical aspect of data management.

### *IBM i Services*

Db2 for i provides many IBM i Security Services to help you understand and analyze your current data security. These services return information related to security such as authorization lists, function usage, object ownership and user profiles. IBM i Services can be easily queried using SQL. Querying a service like OBJECT\_PRIVILEGES can help you identify problems like tables with excessive \*PUBLIC authority. This is a very fast and easy way to understand who has access to your data.

### *Security auditing*

It's vitally important you know not only who in your organization has access to your data but how that access is changing over time. The security audit journal allows you to log and analyze security related events that occur on your system. You can analyze changes to object authorities, changes to object ownership, authorization failures and many other security events. Auditing security events helps you evaluate your security plan and ensure the plan is keeping pace with the changes occurring on your system.

### *Authority collection*

Ensuring a user has enough authority—but not too much—can be difficult. You need to make sure you give a user authority to the objects they need to do their daily tasks, without granting them excessive authority. Authority



**Good data management is an ongoing practice allowing you to obtain the maximum value from your data. Db2 for i has many excellent built-in tools that you already own that are easy to implement to help you along your data management journey.**

collection is a useful tool that can be used to accomplish this.

Authority collection captures data during runtime authority checking and logs that data. This data can then be queried to determine the lowest level of authority needed to complete a task and excessive authority can be removed and appropriate authority granted.

### *Row and column access control*

Sometimes you need more granular rules to enforce access to your data. Db2 for i allows you to enforce this granular authority using Row and Column Access Control (RCAC), letting you place access controls around the data itself. These access controls are at the row level, column

level or both. Data can be masked at a column level and be subset on a row level based on the rules defined. All users of the tables are subjected to the rules, no matter who they are or how they access the data, allowing for separation of duties.

## **Effective Data Management: An Ongoing Practice**

There are many small steps you can take to improve your data management. Good data management is an ongoing practice allowing you to obtain the maximum value from your data. Db2 for i has many excellent built-in tools that you already own that are easy to implement to help you along your data management journey.

---

**SARAH MACKENZIE** is a software engineer at IBM.

# Why Having a Backup and Recovery Plan Improves IBM i Vitality

JENNIFER GOFORTH GREGORY

**D**ebbie Saugen has heard all the excuses for why a company isn't prioritizing backup and recovery for their IBM i system. "Everyone says they aren't going to have a disaster, their business isn't like our other customers', or they don't live in a disaster-prone area," says Saugen, owner of Debbie Saugen Consulting. "But I've worked with hundreds of customers in real-life disasters over the years. It's not just natural disasters—it's hardware failures, it's user errors and now it's a problem with security. If you have a breach, you may need to recover our system in order to recover from that security breach. Every company needs a backup and recovery plan and solution for IBM i."

Saugen finds that most IBM i clients currently back up the IBM i locally to physical tapes. Recently, she met with a client that backed up their IBM i system and data every day to the same physical tape, which she says is an

especially poor strategy. By moving to the cloud, clients have their IBM i data off-site and have multiple copies of the data, which can greatly improve the quality and speed of recovery.

Once you decide to move to the cloud, Saugen recommends focusing first on recovery and then backup. Because the entire focus of recovery is getting your business back online and fully operating, it's essential to understand exactly how quickly you need to recover your IBM i system.

Saugen also recommends testing your recovery of IBM i data at least twice a year. Many companies only think about backup and recovery when they're facing a disaster. By taking the time to plan now for the inevitable disaster for some time in the future, you can improve the vitality of both of your IBM i system and your business.

[Learn more HA/DR strategies in the full article](#)



# IBM Power System Data Migration Considerations and Strategies

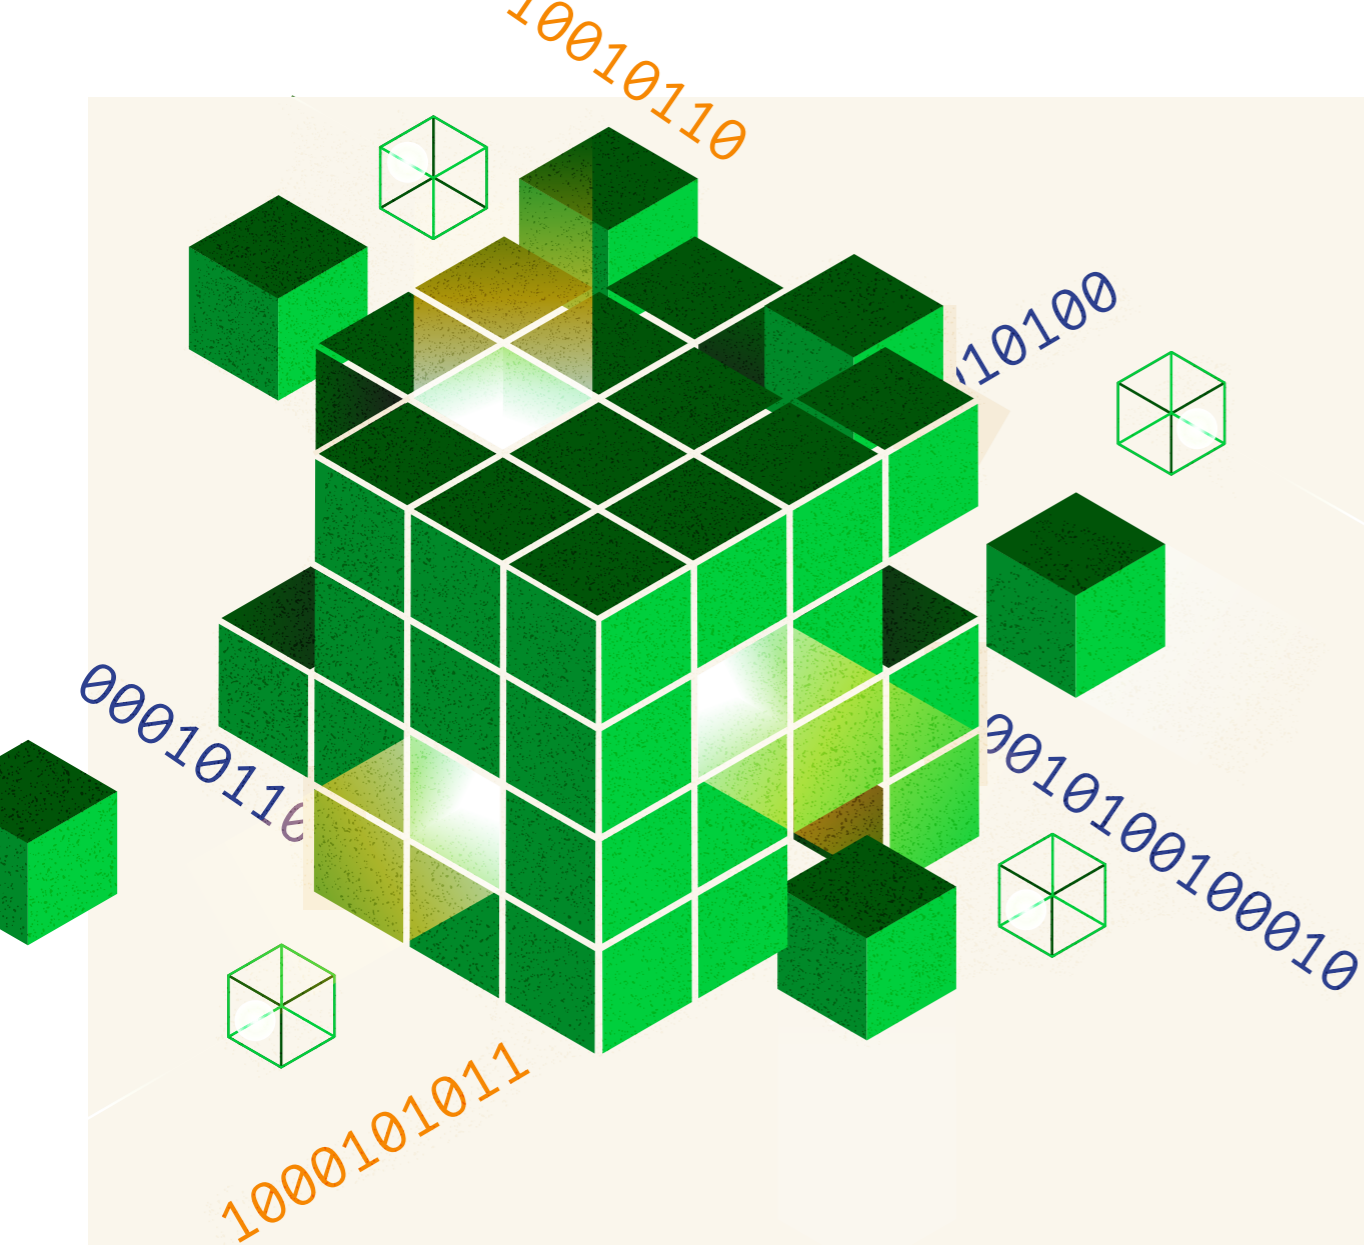
Richie Palma of iTech Solutions Group on IBM Power Systems migration considerations, and what migration strategies to use based on your business needs

BY RICHIE PALMA

**A**s businesses look to adopt the latest and greatest IBM Power Systems hardware into their data center landscapes and available outage tolerances are shrinking, the one-size-fits-all approach to migrations is no longer practical. Demands for 24-7-365 uptime are not only for large organizations anymore, but also small to mid-sized ones. These growing demands require a new mindset and new tools to support migration. Fortunately, for us folks running operations on IBM Power Systems, there are some amazing options to support migrations regardless of your downtime tolerance.

## Backup Planning

During the planning process, you need to take inventory of all the existing physical and virtual environments you will be migrating and any new environments that you need as part of the project. Review the amount of disk you have and plan for future growth for sizing. Once the environments that are being migrated are understood, it's important to size up an IBM Power System with adequate compute resources to support the future workloads.



Pete Massiello, president of iTech Solutions Group, notes his philosophy on upgrades: "iTech Solutions Group does quite a lot of IBM Power Systems upgrades and migrations every year. I can tell you from experience that all successful migrations encompass one major similarity: they all were executed based on a proven and detailed plan of attack—a roadmap of steps and procedures. It works every time. The steps may change based upon your migration methodology and available outage window, but it comes down to building a plan and executing it."

Regardless of your migration strategy, having a solid backup before any upgrade or migration is important. This is true whether you migrate your environment with replication, LPM or SAN to SAN replication. A solid backup to fall back on in the event something happens is key here. Even when pressing on with migration before getting a full system save on the source, it's wise to understand tape compatibility if you have to recover historical data. The general rule of thumb for LTO tape is that you can read tapes from two generations prior and can write to tapes from the prior generation.

## Understanding System Footprints and OS Compatibility

It's also important to understand the physical footprint of both your existing IBM Power Systems stack and your future IBM Power Stack. In many migration scenarios, the hardware will remain in the same data center, so ensuring adequate rack space to run both the existing hardware and new hardware simultaneously is important. The existing source machine should remain as stable and uninterrupted as possible while racking and cabling the new hardware.

OS compatibility is another key consideration at this point. This is especially true if you haven't kept up on



**Regardless of your migration strategy, having a solid backup before any upgrade or migration is very important. This is true whether you migrate your environment with replication, LPM, or SAN to SAN replication.**

regular PTFs and OS upgrades. There are scenarios where your OS is sitting so far back that you require an OS upgrade or PTFs on the source environment to get it to a level supported on the target machine. If you've stayed relatively current on your hardware and OS, this is typically not a problem. For example, let's look at an IBM i shop running IBM i 7.2 on an IBM POWER8 server. If the customer's 7.2 environment is below TR8, they will need to apply PTFs to get beyond TR8 to migrate. While there are ways to migrate data on top of a new OS version, anything you can do to keep the OS consistent between source and target significantly reduces risk and the overall complexity of the migration.

Third-party application vendors also need to be reviewed for any licensing changes that need to be made, OS compatibility if you're jumping your OS during the process, and keys if required for the new machine. Depending on the licensing logic of the third-party vendor, your migration strategy, the number of

environments being migrated, and the timeline for migration, it may be necessary for you to obtain an additional short-term license or temp keys for migration.

## Common Migration Strategies

Since I've touched on some of the critical areas in the planning phase of the process, let's dive into some of the more common migration strategies. It's important to understand the balance between risk, complexity and end-user outage time while examining the following migration options.

### *Full system save and restore*

If your business can afford some scheduled downtime, there is nothing like the reliability and stability of a full system save and restore. If you require 24-7 uptime or can only afford an outage window of a couple of hours, then



this isn't for you. This is for the shop that runs operations Monday through Friday and can afford a reasonably prolonged outage for save and restore. Backup media compatibility plays a significant role in this migration strategy. Backup and restore has low risk and low complexity but a larger outage window.

### *Live replication to a third target environment*

If downtime is a concern of yours, then chances are you already have a replication solution to support faster recovery in the event of a disaster. In most cases, this same tech can be utilized to support live replication to a third target environment, which would be the migration target for the source partition. In almost all cases, you still



## Best-in-Class SMB Solutions: Start Here

The TechChannel 2021 Solutions Directory is the SMB users' guide to the latest in innovative, best-in-class Power Systems, IBM i, AIX, Linux on POWER and storage solutions from ISVs and partners worldwide. Search by OS, company or by 22 product categories for the latest information and expert perspectives available on hardware, accessories, software products, education and services.

[Start your research now](#)

**TechChannel**

seed this second high availability (HA) target with a point in time backup, e.g., physical tape.

After restoring the migration target, HA replication would need to be established and configured/modified. After replication has started the business should run audits on the replicated data and tune the HA replication to ensure everything is replicated to the future target environment. IBM i shops that are confident in their HA configuration, test it regularly and closely monitor it will significantly reduce the time spent auditing the future migration target for completion. When the business is ready to migrate, simply execute a role swap and move operations to the new migration target. Utilizing replication technology for migrations does involve a moderate amount of risk due to all the moving parts, and has a fairly high level of complexity if confidence is not already established in the replication solution. Still, the outage window typically is measured in minutes, not hours.

## LPM


Another option for migrating users is Live Partition Mobility (LPM). It has some very specific prerequisites in order to use it, but is a solid setup for migration if you can swing it. The first major requirement is external SAN storage, which is much more widely adopted now in the IBM Power Systems Midrange space than 5-7 years ago due to advanced functionality like encryption, Flashcopy and remote replication.

LPM also requires PowerVM Enterprise Edition and that you run your IBM i partition on top of VIOS. POWER9 servers come with PowerVM Enterprise built in on every machine, so if you're moving to POWER9 then you automatically get to check that box. Plus, you can get 60 days of PowerVM Enterprise for your existing POWER7 or POWER8 when you purchase your new POWER9 at no charge. The VIOS requirement may be a barrier to entry, but could be considered as businesses build out a highly scalable virtualized infrastructure stack. An LPM setup would typically consist of two separate IBM Power Systems servers connected to a single SAN storage unit.

All storage exists on the SAN, and everything in memory is kept in sync through Active Memory Mirroring. Your migration would be handled with no outage to users by letting the new target server take over as the live partition. For clients with downtime concerns, this is a great way to have no interruption as you migrate your IBM i workload to POWER9. Utilizing LPM for your migration is pretty cut and dry once set up and comes with limited risk. Once you have all of the infrastructure prerequisites in place, LPM requires only a moderate level of complexity to get set up, and it has no outage for users.

## *Separating storage from hardware*

This last step regards external storage. When you separate the storage from the hardware, it opens some tremendous new avenues for simplifying upgrades



**It's important to understand that data migrations are not a one-size-fits-all process. It comes down to your needs and your configuration. The investment you make in planning on the front side will determine your success during execution.**

and reducing the outage window. For IBM i shops, once you have all of your storage on SAN, migrations are as simple as moving a couple of fiber cables and IPL-ing the server. This will require a very small outage but it's straightforward and carries a low risk.

I should also mention how easy it is to migrate from internal storage to external storage for IBM i shops. Because IBM i manages its storage so well, you can use it to your advantage when moving from internal disk to external SAN storage. All there is to do is connect the SAN and provision LUNs to IBM i. Then, tell IBM i that you want it to drain all the storage from the internal drives and move it to the external drives. This can be done on the fly while operations continue to run. Once all of the data is

migrated to external storage, all that remains on internal storage is the load source. This is the only part that requires an outage. Once moved, you're entirely running on external IBM SAN storage and have all of the advanced functionality of IBM's robust storage portfolio.

## Migration Isn't One Size Fits All

It's important to understand that data migrations are not a one-size-fits-all process. It comes down to your needs and your configuration. The investment you make in planning on the front side will determine your success during execution.

---

**RICHIE PALMA** is an account executive and tech solutions consultant with iTech Solutions Group. He is also an IBM Champion for Power Systems.



# Accelerating Clinical Trials With Big Data Analytics

KEELIA ESTRADA MOELLER

Over the last two years, reflects Raj Indupuri, CEO, eClinical Solutions, companies in the life-sciences industry invested more into digital transformation initiatives, with an emphasis on modernizing technology infrastructure. “That means more investments into data infrastructures, whether it’s to handle big data, deploy advanced analytics and machine learning models. It’s leading into larger investments into data ecosystems, along with modernizing analytic capabilities, which is much needed to tap into the data that’s already been collected and is available.”

However, one less fortunate trend in the space is: “The amount of data chaos is increasing significantly,” Indupuri notes. “So, we must find ways to tap into this data, transform this data into valuable insights.”

This brings us to the importance of data management in the healthcare industry. Data management is a very broadly used term, but it takes on a particularly important

meaning in our space: “Data management has a very critical and important function as it relates to clinical development or clinical trials—because you have all of these streams of data to be handled,” Indupuri says.

That’s where big data, advanced analytics, AI and machine learning come in. These strategies can help reduce overall clinical trial cycle times needed to bring emerging research to fruition.

Clinical trials are becoming increasingly complex, and the amount of data collected and used to conduct the trials only continues to increase. In addition to this growing data for new research, there is already a plethora of data from past clinical trials that can be harnessed. Big data, advanced analytics, AI and machine learning can help tap into this data, transforming it into valuable insights—accelerating research, and bringing medicine to patients faster.

[Learn more about data management in healthcare](#)

This e-book was published by



901 N. 3rd St., Suite 195, Minneapolis, MN 55401 // (612) 339-7571

# staff list

---

**Senior Content Strategist:** Claire O'Leary

---

**Senior Editor:** Keelia Estrada Moeller

---

**Art Directors:** Mike Kooiman & Kelsey Hanscom

---

**Project Manager:** Noelle Heaslip

---

**Audience Development Director:** Linda Holm

---

**Publishing Director:** Mari Adamson-Bray  
(612) 336-9241 // mbray@techchannel.com

---

**Account Executive:** Darryl Rowell  
(612) 313-1781 // drowell@techchannel.com

© Copyright 2021 by MSPC, a division of MSP Communications. This e-book could contain technical inaccuracies or typographical errors. Also, illustrations shown herein may show prototype equipment. Your system configuration may vary slightly. This e-book may contain small programs that are furnished by MSPC as simple examples to provide an illustration. These examples have not been tested under all conditions. MSPC, therefore, cannot guarantee or imply reliability, serviceability or function of these programs. All programs contained herein are provided to you "AS IS." IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT AND FITNESS FOR A PARTICULAR PURPOSE ARE EXPRESSLY DISCLAIMED.

All customer examples cited represent the results achieved by some customers. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

The articles in this e-book represent the views of the authors and are not necessarily those of MSPC or TechChannel.

# TechChannel

TechChannel.com is home to a variety of content to help you get started on your data management journey.

[Learn more about data management](#)