



# TechChannel

## Streamline Your Modernization Strategies With Open Source and SQL With Db2 for i

- Scott Forstie on how establishing a virtual layer comprised of SQL views fits into database modernization strategies
- Jesse Gorzinski on open-source modernization trends and advantages

- 
- 3** Achieving Database Modernization Using SQL With Db2 for i
- 
- 9** Using Power10's Superfast AI With Event Streaming
- 
- 10** Open Source Reaches the Mainstream
- 
- 15** AIX 7.3 Leverages Power10 Processor to Support Customers' Hybrid Multicloud and AI Endeavors

# Mainstream Modernization Strategies



Modernization strategies come in many forms—each with a unique blend of cost, risk and value. But some strategies are better than others, depending on your organizational needs.

However, some organizations still have hesitations about adopting the most effective modernization strategies—such as establishing a virtual layer of SQL views within physical data models or leveraging open-source technologies.

In this e-book, Scott Forstie directly addresses some of the most common “blockers” to database modernization, and explains how establishing a virtual layer comprised of SQL views fits into database modernization strategies. Jesse Gorzinski also weighs in on common open-source misconceptions, and highlights why open source has become a mainstream modernization strategy.

**Keelia Estrada Moeller, Senior Editor**






# Achieving Database Modernization Using SQL With Db2 for i

Scott Forstie on how establishing a virtual layer comprised of SQL views fits into database modernization strategies

BY SCOTT FORSTIE



**D**atabase modernization comes in many forms, each with a unique blend of cost, risk and most importantly, value. This article focuses on a very specific idea—establishing a virtual layer of SQL views within the physical data model. By the end, you'll see that this can be done with next to zero risk and next to zero cost, leaving the decision point with you as to the value.

## Common Database Modernization Blockers

Hold the phone. This sounds too good to be true. We don't have a database engineer on our team. We don't have deep SQL skills. Lastly, we don't have time to gain the requisite skills to tackle database modernization!

Let's address some of the most common blockers to database modernization:

**Blocker:** Our physical data model doesn't even use SQL!

**Counter:** SQL views can be built over DDS created database physical files, SQL tables and other SQL views.

**Blocker:** SQL views are dangerous.


**Counter:** SQL views are performance neutral. SQL views take minimal storage, they don't grow in size and they're safe to use.

**Blocker:** We don't have staffing. Management won't buy in!

**Counter:** Gain the awareness, trust and support of management by delivering unfunded modernization to help management understand that a balanced approach is possible, quantifiable and valuable.

**Blocker:** We don't have time!

**Counter:** This idea will take minutes to complete. I think you can squeeze that in. Further, a small amount of time spent on this topic will gradually yield value and return on your investment.





## Redirecting Programs and Users to an SQL View

For many IBM i clients, the physical data model is comprised solely of database physical and keyed logical files. By forcing the users and applications to directly consume the lowest level of the physical data model, there is a significant barrier to change. When a non-SQL language is used to consume a database file, the format of the file is hardbound to the consuming program. Later, if the business needs or wants to add a column to the file, a corresponding recreate of all the consuming programs must be orchestrated.

A closer-to-modern approach for the data model would be to redirect programs and users to an SQL view, even if the view is simply a reprojection of the same columns found in the underlying table. Why? Flexibility. When the applications are



## 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](#)

refocused on a virtual layer of SQL views, the door opens to adding columns to the underlying physical file without the historic concern of coordination, negotiation or otherwise doing the heavy lifting necessary to ripple the changes throughout the data center.

Instead of using the remainder of this article to wax poetic on theoretical matters, I offer you a working solution. Before I take you through the details, let me make one point clear: There are many independent software vendor (ISV) products that can help you with database modernization. I encourage you to take a close look at these options, as they are powerful and should be considered.

## A Working Example

SQL on Db2 for i can do so many amazing things. In the following example, I will show you how SQL can discover database physical files, create the SQL statement text to establish SQL views over those files, and finally, execute the SQL statements. This working example has been published on the Db2 for i SQL Tutor, a site focused on sharing SQL solutions to IBM i challenges.

As seen in Figure 1, the welcome page of the SQL Tutor has different avenues into the examples, and there is even an aggregated list of the **IBM iSee video blogs**, where my colleague Tim Rowe and I visually show how-to topics on the IBM i, frequently with a companion SQL script. From



**Figure 1.** The Db2 for i SQL Tutor

Photo by David Bowman

the SQL Tutor, if you drill into the Database Engineering Topics, scroll down a bit and you'll see the SQL examples and IBM iSee videos as depicted in Figure 2. If you click on the Virtually done.sql link, you should see the SQL source code and sample invocations of the build\_view\_statement() SQL scalar function and the build\_views\_over\_physicals() SQL procedure. To use this example as is, execute the CREATE FUNCTION and CREATE PROCEDURE SQL statements one time on your IBM i.

The build\_view\_statement() function is passed in the schema name and table name of your existing database physical file. The name of the SQL view that you would



like to have created is the third parameter. The example uses a naming scheme where the character “V” is tacked onto the end of the table name. You could adopt whatever naming convention you’d like to utilize. I prefer to have the beginning portion of the view name use the same character sequence as the table name, so that it’s easier to remember the view name and so that the objects appear next to each other in tools that list database objects. If you have tables that use 10-character names, the view will end up with a system generated file name. The users will use the nice, long name and if you want to adjust the system name, do so via Access Client Solutions (ACS).

The fourth and final input parameter is an optional system name for the view, which only needs to be considered if you’re using long names for your physical files, which is doubtful.

Figure 3 shows how easy it is to call the build\_view\_statement() function for a specific file, confirming that the construction of the CREATE VIEW statement is working.

The build\_views\_over\_physicals() procedure drives the show. It utilizes a recently added Db2 for i catalog (QSYS2/SYSFILES) to discover database physical files within the user supplied schema name.

<a href="#">Row level auditing.sql</a> I was asked how to incorporate row level auditing detail into tables. While Temporal tables with Generated columns is a powerful combination, the following example demonstrates a different approach.	Triggers Hidden columns Default values
<a href="#">SQL DDL with nc.sql</a> SQL DML includes the WITH NC clause to avoid having the data change participate in the transaction. SQL DDL does not include the WITH NC clause, but the savvy SQL user can leverage an AUTONOMOUS procedure to achieve the same behavior.	Transaction Avoidance
<a href="#">Virtually done.sql</a> Does your physical data model include a virtual layer? If no, this gist is for you...	Modernizing with Views

iSee Video Tutorials	Topic
<a href="#">Convert your DDS to DDL</a> Many of us have DDS today that still is used to describe our data. We have asked, 'how to I convert that to DDL?' , 'Is there some magic tool?' . In this session we are not going to show you any magic, but we are going to look at some methods that will empower you to be able to take those key files created with DDS and easily convert them to a table with DDL. This is an important step in being able to take advantage for the latest in database technology.	Modernization
<a href="#">Extract your DDL and Save it</a> We all have great procedures for saving the data on our systems. But if something were to happen to your database table, you might have the data, but do you have the DDL to be able to easily rebuild everything ? In this session we will look at how to easily extract the DDL such that you have that key description information for the future.	Reverse engineering SQL objects

**Figure 2.** Database engineering topics include working SQL and how-to videos

```
--
-- Test the SQL View builder function
--
values coolstuff.build_view_statement(
  P_TABLE_SCHEMA    => 'TOYSTORE',
  P_TABLE_NAME      => 'SALES',
  P_VIEW_NAME       => 'SALESV',
  P_VIEW_SYSTEM_NAME => default);

00001
create view TOYSTORE.SALESV ( SALES_DATE, SALES_PERSON FOR COLUMN SALES00001...
```

**Figure 3.** The build\_view\_statement() scalar function constructs a CREATE VIEW SQL statement

**For each file, the function does the following:**

1. Invoke the `coolstuff.build_view_statement()` function to have the CREATE VIEW statement text generated
2. Execute the CREATE VIEW statement
3. Transfer the ownership of the view to match the owner of the physical file
4. Grant the view the same authorizations as the physical file
5. Log the details of the operations in a Declared Global Temporary File (DGTF) named `session.views_created`. SQL DGTF's exist within the current job's version of QTEMP. If you want the results of the procedure call to be saved, use an INSERT with Subselect, querying `session.views_created`.

Once all database physical files, in the given schema, have been processed, the overall results of the function are returned as a result set, as shown in Figure 4.

That's all there is to it! See how easy it is to establish a virtual layer, comprised of SQL views?

**My advice is this:** Approach database modernization using modern tools, with an eye on cost, value and risk.

```
-- =====
-- To invoke, pass in the name of the schema:
-- =====
call coolstuff.build_views_over_physicals(p_table_schema => 'TOYSTORE');
```

TABLE_SCHEMA	TABLE_NAME	VIEW_SCHEMA	VIEW_NAME	VIEW_SQLCODE	VIEW_ERROR_TEXT
TOYSTORE	ACT	TOYSTORE	ACTV	0	<NULL>
TOYSTORE	CL_SCHED	TOYSTORE	CL_SCHEDV	0	<NULL>
TOYSTORE	DEPARTMENT	TOYSTORE	DEPARTMENTV	0	<NULL>
TOYSTORE	EMP_PHOTO	TOYSTORE	EMP_PHOTOV	0	<NULL>
TOYSTORE	EMP_RESUME	TOYSTORE	EMP_RESUMEV	0	<NULL>
TOYSTORE	EMPLOYEE	TOYSTORE	EMPLOYEEV	0	<NULL>
TOYSTORE	EMPLOYEE_DEETZ	TOYSTORE	EMPLOYEE_DEETZV	0	<NULL>
TOYSTORE	EMPPROJECT	TOYSTORE	EMPPROJECTV	0	<NULL>
TOYSTORE	EXTRAONE	TOYSTORE	EXTRAONEV	0	<NULL>
TOYSTORE	EXTRATWO	TOYSTORE	EXTRATWOV	0	<NULL>
TOYSTORE	IN_TRAY	TOYSTORE	IN_TRAYV	0	<NULL>
TOYSTORE	NAMES	TOYSTORE	NAMESV	0	<NULL>
TOYSTORE	NEWSTUFF	TOYSTORE	NEWSTUFFV	0	<NULL>
TOYSTORE	ORG	TOYSTORE	ORGV	0	<NULL>
TOYSTORE	PROJECT	TOYSTORE	PROJECTV	0	<NULL>
TOYSTORE	SALES	TOYSTORE	SALESV	0	<NULL>
TOYSTORE	SAVESTUFF	TOYSTORE	SAVESTUFFV	0	<NULL>
TOYSTORE	STAFF	TOYSTORE	STAFFV	0	<NULL>

**Figure 4.** The `build_views_over_physicals()` table function in action

In some cases, SQL with Db2 for i is the tool to use to modernize your IBM i.

**SCOTT FORSTIE** is a senior technical staff member within IBM i development and spends most of his time working as the Db2 for i Architect. He has worked on IBM OS development since joining IBM in 1989. Scott is a frequently published author, speaker at industry events around the globe, and can be found @Forstie\_IBMi and forstie@us.ibm.com. Scott is an avid runner and enjoys cooking.



# Using Power10's Superfast AI With Event Streaming

BY JESSE GORZINSKI

**T**he latest generation of IBM's Power processors is no incremental update. It comes with built-in acceleration for artificial intelligence (AI)! This is done by integrating matrix math assist (MMA) units right into the chip.

If you haven't already, now is the time to explore AI. IBM i is a natural fit for this purpose. In order to be effective, AI algorithms need "ground truth" data. Simply put, the more "ground truth" is available, the better the intelligence. And guess what houses tons of data: Db2 for i! Some companies have decades of historical data, and some process mind-numbing amounts of transactions every minute. It's a treasure trove for cognitive computing.

There are a number of ways for an IBM i company to exploit this revolutionary processor technology.

In some scenarios, you may wish to analyze transactional data in real time. A logical approach for this is to decouple the transaction processing from the analytics. An event streaming model can be used, whereby transactions are published in real time over some medium for another component to examine as needed.

This event streaming approach, like others, can leverage real-time information to provide the best insights. It's unique in that it takes an offline or out-of-band method. This could be a suitable end goal for a variety of AI applications. In this case, the event-streaming model offers complete flexibility. Ultimately, there are three main components to this approach: a data source, a streaming platform and the AI.

[Learn more about Power10 in the full article](#)

# Open Source Reaches the Mainstream

In a few short years, IBM i clients have embraced open-source technology. Jesse Gorzinski weighs in on open-source trends and advantages here.

BY NEIL TARDY

One way to gauge the value of open-source solutions on the IBM i platform is to consider what actual IBM i clients are accomplishing with the technology. But to truly comprehend the impact of open source, it's necessary to take a step back—perhaps even a literal step back to look around your home.

Open source underpins technology that allows you to check your bank account balance on your phone or start your car from your kitchen. Beyond that, many prominent companies rely on open source in their internal operations—not just tech giants like IBM and Microsoft, not just GitLab (which went public at \$77 a share in October), but businesses you'd never expect.

The reality is that open source is fully mainstream now, and that prominence and prevalence makes it all the more compelling to IBM i clients looking to modernize applications.

“In the past five years, open source has become the forefront classification of technology—for IBM i and the industry as a whole,” says IBM's Jesse Gorzinski, business architect for open source on IBM i. “We have smart cars and smart refrigerators and all these things. I can start my Hyundai from my phone because my Hyundai has a bunch of code that talks to the internet and talks to message brokers to get the signal through. That's powered by open source. Open source is just how we do this stuff now.”



## Fighting FUD

It can be argued that open source has existed since the earliest days of computing. In the 1950s, the SHARE user group worked with IBM to publish applications and utilities as source code for use on IBM mainframes. Certainly if you consider the Zend Server or support for the Node.js JavaScript runtime or Python programming language, IBM i clients have benefited from open source for years. However, only recently has open source really taken off on the platform—early 2018 is a logical milestone.

This could explain why common misconceptions about open source persist within the IBM i world. For instance, is open-source software secure?

“There’s this large misconception that open source is somehow inherently not secure, because the bad guys can see the code,” Gorzinski says. “But who do you do your personal banking with? Chase? Citigroup? Because I can pull up their web portals and show you where they’re using open source to secure your information. These are just a couple of the prominent stakeholders who trust open-source security.”

He goes on to point out initiatives like the Internet Bug Bounty and the Open Source Security Foundation (OpenSSF)—cross-industry collaboratives that strive



to improve the security of open-source software by identifying vulnerabilities and establishing and promoting best practices.

Of course none of this is to suggest that the technology is invulnerable and security isn't a concern when using

open-source software and implementing open-source solutions. The point is there's no inherent reason to fear the "openness" of open source.

"There is fear, uncertainty and doubt—or FUD—around modernizing applications and changing them to do



## 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**



something new,” Gorzinski adds. “And it’s perfectly reasonable to stop and think, ‘Am I opening up an attack vector for malicious players?’ You should ask that question. But the good news is that open source brings industry standard security solutions backed by a large user community and freely available information. You’ll find all kinds of resources at places like Stack Overflow and more reputable sources. So you can start implementing security with a high degree of confidence that you’re doing it right.”

IBM i installations should also be confident in the support structure backing open-source development. In that prior sentence, the word “support” takes on dual meaning. First, there’s actual software support, which IBM recommends and provides for commercial open-source packages it **makes available for download**. Then again, there’s the open-source support emanating from the world at large.

“A lot of the Fortune 500 tech companies invest in open source security. It’s a wide community,” Gorzinski says. “For instance, Google ‘Airbnb open source.’ Airbnb is actually one of the most ethical, contributive companies in the open-source space outside of the big tech firms. People will go to a JavaScript conference, see an Airbnb booth and be like, ‘What? Are they trying to sell us cheap apartments to rent in New York?’ No, they’re actually there to talk about all the great things they’re doing





**“All over the industry people worry about finding skilled people who understand the core business. But open source helps because open-source technology stacks can be used to augment today's business applications in a non-disruptive manner. ”**

—Jesse Gorzinski, business architect for open source on IBM i, IBM

with open source because that's what they're running on their stack. They understand that supporting the open-source community is good business.”

The IT skills shortage is a final factor that causes some to hesitate when considering open source. Sure you can post an ad on a job site and hear from two dozen PHP programmers within a day, but how can a PHP programmer possibly comprehend, much less manage, core business applications running RPG code written back in the heyday of hair bands?

“That's probably the thing you hear most: It's hard to find people with the right skills,” Gorzinski says. “It's especially true with IBM i, but this concern is not unique to IBM i clientele. All over the industry people worry about finding skilled people who understand the core business. But open source helps because open-source technology stacks can be used to augment today's business applications in a non-disruptive manner.

You don't necessarily have to rewrite anything to build a new website, expose APIs, or implement DevOps. You just need to find folks with the skills that you need for the technology you're about to embrace.”

## Modernization Tools

For those about to embark on a new project, or simply get started with open-source development, Gorzinski offers these recommendations:

**Use current programming languages:** Python, Ruby, PHP and other common languages afford users access to millions of modules in community repositories. The Node.js npm repository alone is said to contain more than a million open-source packages. “That means if you're looking to do [Internet of things] (IoT) or microservices or whatever, the heavy lifting has likely already been done for you,” Gorzinski says.

**Check out these tools:** Node-RED is a browser-based low-code programming tool for integrating IoT devices with applications. Even by techie standards, Apache Camel is tough to sum up. But it allows users to integrate disparate systems using the same API. This capability is obviously valuable in any modernization scenario.

Of course, numerous IBM i customers are already doing impressive things with open source. To name just two, there's Wijnen Van Maele and Cras Woodgroup. More examples can be found in the IT Infrastructure section of ibm.com.

Gorzinski describes the recent growth of open source on the platform as transformative. "Certainly adoption rates continue to rise, but with IBM i in particular, we've seen an incredible transformation," he says. "IBM i shops are doing cutting edge stuff with [artificial intelligence] (AI), IoT, microservices and quantum computing. Even more importantly, they're running it in production. That's the MO, the modus operandi, now."

# AIX 7.3 Leverages Power10 Processor to Support Customers' Hybrid Multicloud and AI Endeavors

BY SCOTT MCKINNEY

The new AIX 7.3 release brings customers new ways to integrate their mission critical infrastructure with new and enhanced capabilities. IBM designed AIX 7.3 in response to feedback from customers, particularly around extracting business value from data and supporting their hybrid multicloud strategies. Its new features take advantage of the unique capabilities of the Power10 processor.

The innovations in AIX 7.3 focus on further improving AIX's core strengths customers depend on and exploits new Power10 hardware features, while evolving AIX to better take advantage of new and emerging technologies.



“As we continue to advance the OS, we have an extremely strong commitment to investment protection,” says Carl Burnett, Distinguished Engineer, AIX, IBM. “The investments our customers have made in AIX remain there in a compatible way, so they can get additional value out of the hardware without having to change code in their applications.”

**The new features center around:**

- Streamlined business insights (enterprise artificial intelligence (AI))

---

- Enhanced support for continuous computing

---

- Increased security from core to cloud

---

- Integration into existing environments

---

- Support for hybrid multicloud

**The combination of AIX 7.3 and Power10 helps clients support two broad AI use cases:**

- 1.** In-transaction workloads: You can perform real-time inferencing on an arriving transaction quickly to leverage predictive capabilities versus having to call another AI platform. This includes use cases like risk and fraud detection, or loan and claim pre-approvals—all on the same Power system, close to the data.
- 2.** Operational analytics: AIX 7.3 combined with MMA allows for more sophisticated AI processing on fresh data as it arrives, without having to move it to a different platform and losing the value of insights from newly arriving data. Two examples of this are dynamic price optimization and anomaly detection in a batch of sales orders.

In both cases, AIX 7.3 and Power10 help to meet SLAs around the speed of transactions and response time window with maximum security where you have to make a detection or decision.

[Learn more about AIX 7.3 in the full article](#)

This e-book was published by



953 Westgate Drive, Suite 107, St. Paul, MN 55114 // (612) 339-7571

staff list

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

**Senior Editor:** Keelia Estrada Moeller

**Art Director:** 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

**Sr. Strategic Account Manager:** Lisa Kilwein  
(480) 428-9780 // lkilwein@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 security journey.

[Learn more about modernization](#)