

IBM i 7.2 – The great beyond

An overview of the new IBM i 7.2 operating system release

Tim Rowe

May 08, 2014

There is a wealth of new capabilities available with the IBM i 7.2 release. In this article, you will learn about some of the database, security, manageability, and mobile device integration enhancements that keep your IBM i solutions well-positioned for today and into the future.

Introduction

The world of IBM i is taking a major step forward as we look to the next decade. It has been a while since we delivered a new operating system version for IBM i. Actually, it has been four years since IBM i 7.1 was launched. That in its self is a significant shift from how the IBM i world has worked for years. For the longest time, IBM had been on a 2-year cadence. Similar to clockwork, every 2 years there was a new level of the IBM i operating system. A different approach was taken for the IBM i 7.1 release. Despite the 4-year gap between releases, IBM has continued to deliver new capabilities with a Technology Refresh model. IBM has delivered refreshes twice a year since the GA of the 7.1 release. If you were to look at the functionality delivered with those eight Technology Refreshes and added that all up, yep you guessed it, you would find that we delivered an entire release worth of content! So why change? Why do we need a new release?

Great questions. While we have delivered a great number of updates and enhancements, there are just some things that we just can't deliver in a PTF. For some new function, we need the formal process of recompiling everything along with the extended operating system testing that comes with a new release. For other things, we need that release boundary to make changes and updates to reposition the IBM i operating system for the future.

What did we focus on for this new release? When we sit down to figure out what makes sense for us to deliver in an operating system release, here is the list of guiding principles that help us along this path.

- **Information into insights**

The integrated database has been at the core, IBM i has since its inception. This guiding principle remains true today with all of the new, powerful features that have been added to IBM® DB2® for i.

- **Security without compromise**

Security has and continues to be the hall mark of the IBM i operating system. Seems silly to have a system where you can't protect your data. We continue to focus on insuring that the IBM i operating system has the required controls as the industry continues to emphasize the protection of sensitive data.

- **Solutions for a modern, mobile world**

The world demands that your applications and systems be mobile accessible; with updates to the web and application server infrastructures and a brand new mobile-based system access solution, IBM i is well-positioned for the future.

- **Resiliency without downtime**

The uptime for your IBM i continues to be legendary. There are countless stories of systems just running untouched for years. But you need to ensure that you have a disaster recovery solution ready when disaster strikes.

- **Virtualization without limits**

May be the only thing bigger than *mobile* is *cloud computing*. We continue to invest in virtualization technology to not only help you get the best return from your system investment, but also provide IBM i with the necessary infrastructure for the cloud.

- **Strong integrated value**

Having a great database that is secure with great features is certainly important, but you need something to run that software on. The new IBM POWER® technology leads the industry with processor technology, I/O capabilities, and tight integration with core operating system features.

Information into insights – DB2 for i

The database for IBM i continues to demand our attention. Over the past several technology refreshes, a great deal of new enhancements have been included in the database, well, IBM i 7.2 takes this to the next level. The actual list of new features and enhancements for DB2 on i 7.2 is quite lengthy. I could easily devote this entire article to cover all of the DB2 enhancements, so I've included just couple of the key highlights.

Row and column access control (RCAC)

This is a subject that I am sure you will be hearing a great deal about in the coming days. As you watch the news, there continues to be reports of companies having their data breeched. Many industries are taking a much closer look at who has and who should have access to the actual data. Even worse, some industries are being regulated and being told to control data access. Sure, the IBM i platform is known as being a highly secure system, but some of these rules for governance are intended to help you 'protect you from your self'. I am sure each and every one of you has security on your machines. You may have users, programmers, security officers to control who has access to an object on a system. But just because you have access to an object such as a DB2 table, do you really need access to every row and column value within that table? Sure, your security officer needs to be able to grant authority to other users, but should they also have access to data such as your social security number or other personal data? Yeah, I know you trust them and all, but it really only takes one data breach to cause significant damage to your business. Did you realize that because your security officer has the *ALLOBJ special authority, the officer

can both examine and change any of your data? With RCAC, you can reduce the security officer's access to your data.

Okay, so maybe you're better than that. You have created views over the data to control data access for various users or different groups and roles. Depending on the complexities of your business, the number of views can multiply like rabbits!!! (let me guess, you think rabbits are a good thing, so cute, [well think again!](#)) Just like rabbits running around in your backyard, they all look the same and there are lots of them. How do you gain control over the management of all these views! There is now a better way. This is where RCAC comes into the picture. Instead of trying to control data security at the application layer, data access can be controlled at the database layer across all interfaces. Leveraging RCAC gives you the ability to implement object-level controls without even giving access to all of the data. Let's dig in to how the RCAC support works.

Masks and permissions

A column mask allows you to logically control the data values that are returned for a given column. So, what does that mean? Consider the following example: You have a column for social security number in your database. When that column is queried, most likely, every user sees actual the social security number value. Using a mask allows you to ask questions and determine if the user is authorized to see the actual value. Some users need access to the actual value while only a masked value should be returned to other users! Let's review a simple column mask example.

```
CREATE MASK SSN_MASK ON EMPLOYEE
FOR COLUMN SSN RETURN
CASE
  WHEN(VERIFY_GROUP_FOR_USER(SESSION_USER, 'PAYROLL') = 1)
    THEN SSN
  WHEN(VERIFY_GROUP_FOR_USER(SESSION_USER, 'MANAGER') = 1)
    THEN 'XXX-XX-' CONCAT SUBSTR(SSN, 8, 4)
  ELSE NULL
END
ENABLE;
ALTER TABLE EMPLOYEE
ACTIVATE COLUMN ACCESS CONTROL;
```

In this example, users who belong to the PAYROLL group profile are returned the actual social security number value. In contrast, users who belong to the MANAGER group profile can see only a masked version of the number with the first five digits masked with Xs. For all other users, a null value is returned. Again, the key point here is that this security processing is not being done at the application level. It is performed at the database level. Meaning, the security constraints are now part of the operating system and enforced against all interfaces that try to access the employee table. While masks can be used to protect a column containing sensitive data, row permission can be established to limit the rows accessible to a user, group profile, or whatever logic you'd like to use to limit access to rows in table. By establishing permissions, you can start to eliminate some of those rabbits. The RCAC support is a key tool in really helping you to define and implement granular data access controls to ensure that users get access to only the minimal amount of data required.

SQL improvements

This section alone is worth a book, in addition to continual improvements in the performance of the SQL Query Engine (SQE) and the improved indexing support. SQE is now also used to process native query access interfaces including the Open Query File (OPNQUERYF) command and Query/400 requests. What does this mean? These older style queries will now have the potential to run a little faster and be analyzed with the newer DB2 performance tools. Over the past years, we have been pushing you to move to a modern query and reporting tool such as IBM DB2 Web Query, this new improvement is not an excuse to remain using the old tools - it is still best practice to migrate these older queries to SQL-based interfaces such as DB2 Web Query so that you have easy access to the functional advantages offered by SQL.

DB2 for i also includes support for new built-in global variables. What are these objects? They provide a convenient method for you to learn about your runtime environment from within your SQL scripts or trigger programs. Instead of having to call cryptic system APIs, you can let the database keep track of the information and simply interrogate the built-in global variables. Here are some of the more interesting built-in global variables:

- CLIENT_IPADDR
- CLIENT_HOST
- CLIENT_PORT
- ROUTINE_SCHEMA
- ROUTINE_SPECIFIC_NAME
- ROUTINE_TYPE

Timestamp precision

When a row is inserted into a table, you often are required to keep track of the time the transaction occurred. Seems logical, well today you have six positions of precision for your timestamp. For many years, that was more than enough precision. Today's servers are much faster, so transactions can appear to have happened at the same time with only six digits of precision! To address this issue, DB2 for i now supports an additional six digits of precision. You can now track 12 digits of precision. In addition, you can go the other direction and drop the sub-second precision points to reduce the storage requirements for your table. Today, that may mean you have six digits of precision being wasted for each row in your database table. Multiply the storage allocated for those six digits times the number of rows, and that can be a lot of wasted storage for a large table. Now, you can specify a precision of 0 to remove that un-needed precision.

Well, I am just having too much fun with all these very important new updates to DB2 for i (and I am not even a database guy!!!), but it is time to move on to some other things.

Solutions for a modern mobile world

The web and application infrastructure space has a large number of updates and improvements. We have made focused changes to each of the components so that the infrastructure is well-positioned for the future.

IBM HTTP Server for i

To help keep our system current, we have updated the Apache HTTP server up to the 2.4 level. This is an important update for the future. We have incorporated the latest updates for web serving with this new server. This update includes the latest support for security and other features proved by the open source community.

Integrated web application server

The runtime engine for the integrated web application server has also been updated. We are now using the IBM WebSphere® Application Server - Liberty Core profile as the basis of this integrated server. You can quickly and easily create your own servers without having to install WebSphere Application Server! As your instance needs grow, you have a guaranteed growth path to the full WebSphere Application Server profile. This integrated solution is a multi-platform web container - a true replacement for the ASF Tomcat server that I know many of you use.

Integrated web services server

Our integrated web services run time has also been enhanced to run on this new and improved integrated runtime engine. In addition, the web services engine is based on the Liberty Java API for XML-Based Web Services (JAX-WS) runtime environment. The really good part about these updates is that web service developers will see little change when deploying web services other than faster startup times, improved performance, and greater stability. These updates are critical to ensure that the IBM i platform has a secure and compliant web presence into the future.


Java and PASE

The world of Java has been updated with this latest (at the time of this publication) IBM i operating system release. The PASE environment is an embedded IBM AIX® kernel built into the IBM i operating system. This kernel provides the base run time for the IBM Java engine. We have updated this kernel with a newer level of AIX. These changes can help the Java engine to fully use the IBM Power® hardware. From a Java perspective, there is one significant change; Java 5 is no longer supported. I know, hard to believe but Java 5 was released in 2004! It's time to move to Java 6 and Java 7.

Mobile

We have a very exciting new mobile solution for IBM i that we are delivering. Ever needed to manage certain aspects of your IBM i system from your mobile device? Want to access a 5250 screen from anywhere? Run an SQL script from your mobile device? Now you can. We have updated the IBM i Access for Web (5772-XH2) product to now be mobile-enabled! This mobile enablement allows you to check on the status of a job, view spool files, run SQL scripts, start a 5250 emulator along with many other tasks. Figure 1 demonstrates how a mobile device can access the list active of jobs.

Figure 1. Mobile solutions for accessing IBM i




This new mobile solution is going to be delivered as a technical preview for IBM i 7.2 GA. The fun thing about this new offering is that it also supports servers running the IBM i 6.1 and IBM i 7.1 releases. The 7.2 version of the XH2 product can be ordered and loaded on the previous IBM i operating systems releases without any issues. You can find additional details on the [IBM i Access product page](#).

Navigator

The IBM Navigator for i web interface has continued to progress and a great deal of new features and functions have been added:

- **Monitors**
From the web-based Navigator for i, you can now create monitors for many system performance attributes and create monitors to watch for any message on your system. When a monitor hits a threshold, a set of response actions can be defined to automatically run.
- **PTF management**
You can load and apply PTFs from the web-based interface, as shown in Figure 2. You can do this on your local system or you can send PTF and apply to a group of systems. Compare your local system to any number of other systems in your network and then automatically load any missing PTFs.

Figure 2. Web Navigator for i PTF management




- Batch modeler

This is really cool stuff! First, you start a collection for an existing batch job workload. After you have this collection, then using the modeler, you can change a wide variety of parameters such as run priorities, memory, processor, and disk to then model how these changes affects your batch runtime window. Graphical charts such as the ones in Figure 3 provide a side-by-side comparison. It is easy to run the modeler many times to understand the performance of different configurations.

Figure 3. Batch modeler

Measured vs Modeled Resource Utilization Overview




- IBM Performance Data Investigator (PDI)

The PDI support is now properly placed in the left navigation area and is enhanced with new IBM i 7.2 metrics. Within PDI, there are a very large number of graphs and performance metrics which sometimes can make it difficult to find the chart or metric you are looking for. That problem is no longer an issue with the new fast PDI search interface. Enter the metric (or graph you want, and PDI will find and present the artifacts that contain the specified item. Click the link, and immediately you are taken to that page!

- Performance and usability

Chrome is now a supported browser that can be used. A new 'My Favorites' list has also been added to provide fast-path navigation to important pages and tables. The browser interface also contains a fast-path integrated file system (IFS) control. Instead of clicking and navigating through folders, you can just enter the file path and go directly there. Lastly, Navigator for i includes a fast search interface that allows you to find interfaces using the CL command names! For example, just enter the search string, WRKACTJOB, in the fast search control as shown in Figure 4 and you can go directly to that interface with no extra clicks!

Figure 4. Navigator for i - fast search



Samba

Any of you want to Samba? Besides sounding like some dance from Brazilian heritage, what are we talking about? Samba is an open source SMB-based networking protocol for providing fast, stable, and secure file access. File serving, something we have had on IBM i for a long time is provided with IBM i NetServer. IBM i NetServer has many great features when it comes to file serving, although in some instances performance has been an issue. Samba is a server that uses TCP/IP on IBM i to interact with Microsoft® Windows® clients or servers as if it is a Windows file and print server. Samba is not intended to be a full replacement for file serving on IBM i, but rather give customers an additional option. For example, Samba does not support Kerberos, automatic CCSID conversions, or integration with IBM i auditing exit programs. For those features, IBM i

NetServer is the best choice. But, if you require just basic file serving with performance, then Samba might be the choice for you. Consider the following key features:

- A fast, modern, lightweight Common Internet File System (CIFS) file server
- An FTP style CIFS client (smbclient)
- NTLM and NTLMv2 authentication with session security enhancements
- Enhanced error code reporting
- Better protocol compatibility with current Windows clients
- SMB 2.0 protocol support

Resiliency without downtime

The latest updates to the IBM PowerHA® Express Edition provides IBM HyperSwap® support. You now have the ability to *instantly* switch from production IBM System Storage DS8000® instance to a remote DS8000 system. This switch can be triggered manually for something like planned maintenance or even better, and it can happen automatically in the case of a DS8000 system failure. This enhancement contributes toward the goal of 'continuous availability'. In addition, there are key updates to the base PowerHA solution, including new updates to the SYSBAS support. You can now replicate object authority and ownership with the Administrative domain. This is just some of the updates that better ensures your business stays running.

Security without compromise

Obviously, the RCAC support we talked about is a huge improvement that certainly falls under the umbrella of security. Because it is built right into the database, which is an integrated part of the IBM i operating system, there is no compromise. The RCAC data security rules are guaranteed to be enforced across all system interfaces. In addition to RCAC, a number of other security updates have been included. With the IBM POWER8™ processor technology, cryptographic performance acceleration is available. This is a special support that is built into the processor itself and doesn't require any additional products or hardware. The performance acceleration is geared for certain cryptographic algorithms such as AES and SHA-2. This acceleration should also improve the performance of SSL, VPN, tape encryption and other software applications that leverage the cryptographic services APIs.

Single sign-on (SSO) has been enhanced to support several additional interfaces. You can now configure SSO for both the FTP and Telnet interfaces. With the new FTP support, SSO can be enabled between Kerberos-enabled FTP clients and the IBM i FTP server or between the IBM i FTP client and some other FTP server. Likewise, you can configure any Kerberos-enabled Telnet client and the IBM i Telnet server. These are supported through Kerberos authentication and Enterprise Identity Mapping functionality.

The security audit support has been enhanced to help better track changes on your system. For instance, many of the security audit records now include both the before and after data. There is also a new value for the QPWDRULES system value. Specifying the new value of *ALLCRTCHG ensures that the password composition rules are enforced on both the CRTUSRPRF (Create User Profile) and CHGUSRPRF (Change User Profiles) commands. This eliminates a loop hole where system password rules could be compromised.

Strong integrated value

One of the big reasons for doing a new release is that it is easier for the operating system to support new hardware. IBM continues to invest in IBM POWER® technology and IBM i clients maximize their investment when the operating system is able to use the hardware advancements. With the IBM i 7.2 release, you now have support for the new POWER8 processor chip. This new processor technology provides a pretty amazing platform for running your applications. This new chip is loaded with multiple caches to keep your data closer to the processor – this helps the processor spend more time running than waiting. In a way, it is just like when you are our driving your vehicle, the car is meant to drive, not sit idle. You want to have all the traffic signals to all turn green before you get there. The POWER8 processor is designed to accomplish a similar task, get everything in place so that the 'lights are always green'. Additionally, the POWER8 processors contain a significant bump in the number of cores on the chip and the number of threads that can be supported by each. This should further strengthen the ability of the IBM i operating system to scale as more workloads are added to your system.

I cannot cover all of the hardware and operating system enhancements in this this article, but here are a few of the additional highlights that you might want to investigate further:

- BRMS enhancements
 - Dashboard of all BRMS systems
 - Central site monitoring
 - Faster install
 - DUPMEDBRM concurrent duplications in batch
 - ASYNCBRING to improve IFS backups
- Better tracking of temporary storage usage
- Limits to growth updates
 - 2 TB memory pools
 - Double the job table limits
- LDAP updated to ITDS 6.3
- DNS BIND 9.7.4-P1
- New SMTP server, removes SSD limitation
- C++ compiler include C++ 10x support
- Many networking enhancements
- Color printer image support – including JPEG, TIFF, GIF
- Support for WAN over LAN

As you can now see, there is a wealth of new enhancements available with the IBM i 7.2 release. Several of these updates have the potential to be a game changer for many of you. So it is time to start planning your next system upgrade – speaking of upgrades, the upgrade path is easy with no object conversion requirements. Enjoy the brand new IBM i release that can keep your solutions well-positioned for today and in the future.

Resources

- [IBM Knowledge Center – IBM i 7.2](#)

- [IBM i Technology Refresh Wiki](#)
- [IBM i Access products](#)

© Copyright IBM Corporation 2014

(www.ibm.com/legal/copytrade.shtml)

[Trademarks](#)

(www.ibm.com/developerworks/ibm/trademarks/)