Securing IBM i: A Dual Responsibility

Carol Woodbury, CRISC, CISSP

President and CTO, DXR Security

DXR SECURITY

Securing IBM i: A Dual Responsibility

We've long thought of IBM i as being 'secure-able' but where does the responsibility lie for making that happen? I assert that it's a dual responsibility between IBM and organizations using IBM i.

IBM's Responsibility

Obviously, if the operating system itself is full of security issues, as a community we're going to have a hard time securing our data. In fact, it might be impossible to secure our data. So what does IBM do to hold up their part of the responsibility of securing IBM i?

The thought of security starts where it should: with the design of the operating system itself and with every new feature that is implemented. In fact, the entire IBM Corporation has a stated mission of making sure products are designed with security and privacy in mind. Not only are designs scrutinized for security and privacy issues, but testing occurs during several phases of the development process to look for deficiencies in these areas. But even with these efforts, we still see security-related patches (PTFs, in IBM i terms) being released. How is this possible? Several reasons exist.

Technology Changes

Unfortunately, as technology evolves, the reality of what was once a secure architecture may be unveiled as insufficient. Take, for instance, the early days of AS/400. At that time, it was quite easy to secure your data. You simply had to restrict your users to a menu, make sure their session was signed off if they got out of the menu, and configure the profile as a limited user so users couldn't enter commands from a command line. Since the only way users could access the system was through a 5250 emulator, data was secure. Enter TCP/IP, PCs, and other technology, and that security scheme is FAR from sufficient for securing data today.

People

While the men and women that develop IBM i are some of the best, they are, after all, human. And as we all know, humans make mistakes. Intentions and testing aside, unfortunately, some bugs slip out the door.

Open Source

Because of the desire to stay current with technology and to meet the demands of the IBM i user community to make it available, Open Source has become widely used by both the operating system and business partners alike. As I'm sure you're aware, the Open Source Community is not without its share of security issues. Therefore, if an Open Source product or technology has a security issue and it's available on IBM i, IBM i may also be affected.

Responding to Vulnerabilities

When I talk with the IBM i community, I almost get the sense that they're offended at the very thought that IBM i could have a security vulnerability. That we would think IBM i is somehow immune from having vulnerabilities is just not realistic. Because of that, we must focus on how IBM responds to these vulnerabilities.

What is the process that IBM goes through to identify whether something constitutes a vulnerability and therefore needs to be fixed? Let's first look at how vulnerabilities get reported. A customer or business partner of IBM typically submits a vulnerability or what they believe to be a vulnerability through the normal support process. But contrary to popular belief, the hacker community does have eyes on IBM i and those individuals usually submit vulnerabilities via https://hackerone.com/ibm. (Some hackers notify vendors when they discover a vulnerability, giving the vendor a chance to provide a fix prior to the hacker publicizing and exploiting the issue.) Regardless of how IBM is notified of vulnerabilities, the IBM i team takes the input, evaluates it and, if necessary, issues a fix and assigns a CVE. (CVE stands for Common Vulnerabilities and Exposures.) A CVE provides a unique identifier for a security vulnerability. IBM participates in the well-known, industry-accepted CVE Program to ensure that the vulnerabilities identified within their products are communicated via known channels, documented in industry-standard terms and follow industry standards when classifying the vulnerability as Critical, High, Medium or Low.

- ➔ I cannot emphasize enough that, if you are alerted to or otherwise discover a CVE affecting IBM i that's classified as Critical or High, you MUST take immediate action and not simply ensure the PTF is applied during your next PTF cycle. Many organizations wait 90 or 180 days to apply PTFs after they've been issued or routinely apply PTFs once a quarter. Critical and High fixes should not be treated like a 'typical' fix! Leaving them unapplied leaves your systems vulnerable. Every week I read about an organization that was breached because a known vulnerability has been left unpatched. Do you risk your organization by leaving non-IBM i Critical or High Risk CVEs unpatched? Then why would you leave IBM i unpatched? The *maximum* time you should leave a Critical or High Risk issue unpatched is 30 days. But the sooner the better.

You may also see IBM issue a Security Bulletin when a security vulnerability has been discovered in an Open Source product. All CVEs for Open Source products are reviewed by the IBM i team. If an Open Source product used within IBM i has a CVE issued for it, an analysis is performed to determine if the vulnerability also applies to IBM i. If it does, an IBM i-specific Security Bulletin is created. If you are unsure as to whether a particular CVE affects IBM i, you can search the IBM Security Bulletin website to determine the answer and take action if necessary.

In addition to providing fixes, IBM often makes security adjustments in new releases. For example, in IBM i 7.5 the *PUBLIC authority setting of numerous objects was set to a more restrictive value to come in line with the security best practice of not granting more authority than required. And adjustments were made for some commands to require a special authority to better align their use to specific roles rather than allow any profile to use them. Last but not least, IBM has added a myriad of security features and enhancements over the years both through major release cycles as well as Technology Refreshes (TRs) that allow us to secure IBM i and its data.

Your Responsibility

As you can see, IBM has a very robust process for identifying, evaluating, and fixing vulnerabilities as well as providing us with integrated security and features for securing our systems. All this is well and good, but if we don't use those features or upgrade our systems to a supported release and stay up to date with the latest TRs and especially PTFs, the IBM i Community isn't holding up its end of the responsibility for securing IBM i, and it goes without saying that systems are left vulnerable. The fact that organizations choose to ignore IBM i in this way is stunning to me. IBM i most often holds the 'crown jewels' of an organization's data. Why organizations don't value this data and ensure that it's secured appropriate to its worth is baffling. Or IBM i may drive one or more critical business processes – perhaps running a manufacturing line, an online order entry system, a warehouse, or another critical piece of the organization's technology. So critical in fact, that if it's not available, the organization will, at best, lose a significant portion of their revenue or at worst, be forced to shut down entirely and hope they can recover quickly. Again, the fact that IBM i's security is not maintained according to the role the system plays in the organization, is stunning to me.

While I am disappointed over how many organizations fail to take advantage of the security features IBM has provided for securing IBM i, I believe the larger issue today is that of not keeping systems current. I see three major areas of concern when systems are not kept current.

Continued Use of Old Technology

I am amazed (and not in a good way) at the number of organizations that continue to run the old Client Access for Windows product. The excuse I hear is, "It still works." Seriously? You're ok with putting your organization at risk by having an unsupported product on everyone's desktop? This product has been out of service for so long (since April 30, 2019, to be exact) that it has not received the security scrutiny of current products. So who knows if exploitable vulnerabilities exist? Are you really going to stick with this technology and hope some zero day vulnerability isn't discovered and exploited? (A zero-day attack is exploiting a vulnerability that was previously unknown to the vendor and, therefore, no fix (patch or PTF) is available when the initial attack is launched.)

Also, how are you going to defend your position to stay on such an old product if Microsoft breaks Client Access during a Patch Tuesday cycle? IBM's not going to bail you out of this situation. So rather than architecting a staged roll-out of Access Client Solutions (ACS), I predict you're going to find yourself in an all-hands-on-deck scramble to get all desktops converted so users can resume their daily tasks… all because you justified staying on old technology because, "It works."

Heritage Navigator for i is another technology that should no longer be in use. The log4j industry-wide scramble and subsequent analysis by IBM determined that it should no longer be in use and that the New (and I must say, GREATLY improved) Navigator for i should now be your browser interface into IBM i. Unfortunately, I continue to see Heritage Nav used, thus leaving those systems vulnerable to log4j exploitation. Which reminds me to point out that just because a vulnerability has gone out of the news cycle doesn't mean it doesn't continue to be exploited.

The continued use of old technology is not limited to just Client Access and Heritage Navigator for i. Anyone that lived through the log4j saga understands the pain of finding and patching all vulnerable

implementations on all technologies across their entire organization (including IBM i.) What are you going to do to protect your system and your organization should another issue like log4j arise that affects an unsupported version of IBM i? The thought should make organizations running unsupported versions of the operating system and vendor products shudder. I'll remind you that the hacker community DOES have an interest in IBM i. Do you really want to be running ancient and/or known-to-be-vulnerable technologies and hope nothing happens?

The Vulnerability of Unpatched Systems

To stay abreast of the current security landscape, I subscribe to five different security-related newsletters and bulletins. Every day at least one of them describes a new zero-day vulnerability. More disturbing are the follow-on articles describing how organizations are being breached because they've failed to apply the patch once the vendor released the fix.

What does this have to do with IBM i? Everything! We MUST keep our systems patched – including IBM i – if we wish to stave off attackers and/or limit the damage they can do if your organization is breached. This is not a plot for a movie or a best-selling novel. I'm talking about the reality of today's world whether you wish to acknowledge it or not. Wise organizations take the posture that it's not a matter of 'if' they will be breached but 'when'. They have an incident response plan in place to respond to the breach and regularly practice it so they aren't caught flat-footed when an incident occurs. Their patch strategy is included as part of their incident response plan. Part of that plan includes a technology inventory so when they are alerted to a zero-day exploit, they can take the systems offline or isolate them and apply the patch when one becomes available. In other words, they have an out-of-band process defined for applying patches for Critical and High Risk vulnerabilities – including those that apply to IBM i.

IBM provides several methods for organizations to stay current with PTFs. First, if you haven't already, I encourage you to sign up for alerts from IBM. You can sign up for security bulletins (alerts regarding security vulnerabilities) as well as notifications for other types of events, such as the release of a HIPER PTF. (Note you can also sign up for alerts for other IBM products – not just IBM i.) You'll need an IBM ID, but that's the only requirement. If you're not signed up, here are the instructions from IBM.

I also encourage you to check at least monthly, to determine if your system is current with group PTFs. The easiest way to be aware of your PTF levels is to use the following IBM i Service which does a 'phone home' to IBM to find the most recent PTFs and compares that list to what's installed.

```
SELECT *
    FROM SYSTOOLS.GROUP_PTF_CURRENCY
    ORDER BY PTF_GROUP_LEVEL_AVAILABLE - PTF_GROUP_LEVEL_INSTALLED DESC;
```

Unfortunately, there is an occasional defective PTF, so another service has been created that looks at defective PTFs and shows whether the replacement PTF is installed.

```
SELECT *
    FROM SYSTOOLS.DEFECTIVE_PTF_CURRENCY;
```

You can even determine your firmware levels

```
SELECT *
    FROM SYSTOOLS.FIRMWARE_CURRENCY;
```

I understand that some organizations have a difficult time taking an outage. But over the years, IBM has improved the PTF process and fewer PTFs require an outage to get applied. If you haven't reviewed your patch strategy recently, I highly recommend you do so and, as part of that effort, determine if you are using the most recent processes for obtaining and applying PTFs as well as develop an out-of-band process to address fixes for Critical and High Risk issues.

Finally, don't forget about any of the Open Source packages that you may have installed yourself. IBM will obviously not be providing fixes for those. You must keep those up-to-date yourself. Do not forget this important step for keeping your system patched!

Compliance Requirements

Finally, not only do organizations that fail to keep their systems current lose out on the plethora of security features added each release (IBM i 7.5 was particularly bountiful with new security features and updates), but those systems are often not able to comply with the requirements of regulations such as the Payment Card Industry's (PCI) Data Security Standard (DSS) or laws such as the EU's General Data Protection Regulation (GDPR) or meet the requirements of government entities such as Singapore's Monetary Authority of Singapore (MAS) or individual U.S. State requirements. For example, systems running older versions of IBM i cannot meet the encryption requirements of the PCI DSS when making secure connections because they don't support TLS 1.2 and only run the deprecated (aka vulnerable) versions of this technology. In addition to simply being a good business practice, many laws and regulations require that an organization only run supported applications and use current technology.

Summary

I usually write about IBM i Security features and describe how to use them, but I felt compelled to raise awareness over the risks of not keeping systems current. I wish we lived in a world where no security vulnerabilities exist, and the threat of a cybersecurity incident was just a theory that couldn't possibility happen to us. Unfortunately, that's not the case, and we must take the topic of securing IBM i seriously. I trust that IBM will continue to hold up its end of the equation for ensuring we can properly secure IBM i. My question is … will you?

Carol Woodbury is President and CTO of DXR Security. Carol has specialized in IBM i Security starting with her tenure as iSeries Team Leader and Chief Engineering Manager for Security for IBM in Rochester, MN. Since leaving IBM in 2000 she has co-founded two businesses focused on helping organizations running IBM i to implement the security features provided by IBM. Carol is a world-renowned speaker and writer on IBM i Security including two books: IBM i Security Administration and Compliance, Third Edition and the complementary, Mastering IBM i Security: A Modern Step-by-Step Guide. Carol holds her CISSP and CRISC security certifications.