2019 State of IBM i Security Study

Data breaches and lax cybersecurity make headlines with alarming frequency, and the past year was no different:

- Facebook was still dealing with the Cambridge Analytica scandal when it revealed a data breach in September of 2018
- Google faced multiple GDPR complaints from privacy rights groups and a fine of 50 million Euros
- Panera Bread, Quora, Marriott, and many others exposed data belonging to hundreds of millions of people

While the headlines are concerning, thousands of other breaches didn't make news, but did affect those organizations—and their IT teams. Facebook can withstand a cyberattack. A small or mid-sized business might not.

So, it's no surprise 70 percent of IT pros working on the highly securable IBM i OS are concerned about cybersecurity.

The 2019 State of IBM i Security Study reveals concrete, impartial data about how IBM i systems can be protected—but often aren't.

EXECUTIVE SUMMARY

For the 16th year, this study provides compelling insight into security weaknesses affecting many IBM i systems—systems that are often used for business-critical data, payment card data, and personally identifiable information (PII).

The 2019 State of IBM i Security Study analyzed 244 servers and partitions, drawing participants from finance, retail, manufacturing, and many other industries.

This is not a recurring study of the same systems each year, but general trends are apparent. Overall, cybersecurity is becoming a higher priority for participating organizations, and in recent years businesses have made gradual improvements with basic system security and password controls.

Despite some improvements, the study results show many organizations are still in the early stages of implementing IBM i security controls.

DATA FROM SEVEN CRITICAL AREAS OF IBM i SECURITY, SUMMARIZED BELOW, REVEALS THE EXTENT OF THE RISK:

Setting the Tone: Basic System Security Levels

70 percent of the systems studied follow best practices for overall system security as recommended by IBM and industry experts.

Users with Superpowers

Overwhelmingly, the IBM i servers in this sample have too many profiles with powerful authorities. In the hands of careless or disgruntled employees, this could result in data loss, theft, or damage. Auditors check for the abuse of special authorities as part of any standard IBM i audit. Even auditors who are not very familiar with the IBM i environment are aware of this issue from other platforms.

Faulty Passwords and Profile Security

In IBM i, user profiles with a default password have a password that's the same as the user name. Twenty-six percent of systems studied have more than 100 user profiles with default passwords. One system has a total of 12,000 user profiles with default passwords, almost one-third of those maintained in an enabled state.

Data Access Goes *Public

Virtually every system user has access to data far beyond their demonstrated need. Auditors typically look to ensure that the company has adequate separation of duties and appropriate controls in place to enforce the separation of duties.

Through the Networks' Back Doors

Network access control and auditing is nonexistent in most IBM i shops, so both authorized and unauthorized access occurs without accountability or traceability. IBM's exit point technology provides the ability to control and monitor network data access. However, the study indicates that the adoption rate of exit programs has not kept pace with the adoption rate of network data access utilities.

Straying from the Audit Trail

On most of the systems studied, security violations could occur undetected. 85 percent of organizations log security-related event to a secure repository, but most lack an efficient strategy for monitoring and interpreting security data.

Susceptible to Viruses and Malware

11 percent of servers are now scanning files on open. This means the other 89 percent remain at risk of having internal objects impacted or of spreading a hosted infection to another server in their network, but it is an improvement over previous years.

TABLE OF CONTENTS

4. Introduction

4. Why This Study Is Important for You

5. The Power Systems Landscape

5. Our Methods

7. Setting the Tone: Basic System Security Levels

8. Key System Values for Restoring Objects

9. Users with Superpowers

10. Faulty Passwords and Profile Security

10. Prime Targets: Inactive Profiles

11. The Open Secret of Default Passwords

12. Setting a Standard for Minimum Password Length

13. Capitalizing on Other Password Settings

14. Forgotten Passwords and Other Invalid Sign-On Attempts

16. Data Access Goes *Public

19. Through the Networks' Back Doors

21. Users Crossing the Command Line

21. Straying from the Audit Trail

23. Susceptible to Viruses and Malware

25. Conclusion

27. About the Author

27. About HelpSystems

Introduction

Cyberthreats become increasingly sophisticated every year, raising the importance of proper security controls. The 2019 State of IBM i Security Study proves once more that many organizations running the IBM i operating system rely on system settings that leave data vulnerable.

BUT IN RECENT YEARS, HELPSYSTEMS HAS OBSERVED AN ENCOURAGING TREND: ORGANIZATIONS LARGE AND SMALL INCREASINGLY PRIORITIZE IBM i SECURITY.

But in recent years, HelpSystems has observed an encouraging trend: organizations large and small increasingly prioritize IBM i security. As IBM i pros began to understand that this OS is highly securable—not inherently secure by default—more of them invested time and resources in IBM i security education.

A deeper understanding of the risks and the security controls built into the platform is currently driving a wave of interest in prioritizing cybersecurity issues and getting help implementing controls.

No matter where your organization falls on the IBM i security continuum, the annual State of IBM i Security Study strives to help you understand common IBM i security exposures and how to correct them quickly and effectively.

Why This Study Is Important for You

Over the past 16 years, the State of IBM i Security Study has provided invaluable security insight from thousands of participants worldwide. The results from the 2019 study lead us to conclude that if you have IBM i systems in your data center, your organization might also suffer from internal control deficiencies.

Your IBM i server likely runs mission-critical business applications—and has for years. By now, the staff that set up server security may no longer be with your organization.

To complicate things, the nature of many IBM i security controls causes confusion over who is responsible for the security configuration—IBM, the customer, or the application providers. As such, many systems operate with default settings due to lack of ownership.

You know an IBM i audit is long overdue, but you’re too busy grappling with:

- Knowledge gaps
- Overextended staff
- Lean IT budgets

Because Windows and UNIX platforms tend to require more resources to secure them, it’s much easier to let IBM i security projects take a back seat.

Consequently, the administration of IBM i security controls has lapsed and guards are down even as threats to your system grow.

Here’s the good news: the weaknesses identified through our scans and documented in this study are caused by poor or missing configurations that can—and should—be corrected.

This study shows you the most common and dangerous IBM i security exposures; outlines best practices for improvement; and explains how these relate to compliance legislation, industry regulations, and IT guidelines and standards.

The Power Systems Landscape

The IBM i community is a large and loyal one, with IBM estimating 120,000 customers around the world use IBM i. Over 70 percent of organizations on IBM i run more than half their business applications on this platform, according to the 2019 IBM i Marketplace Survey Results.

Companies in retail, financial, manufacturing, and distribution industries typically purchased their Power Systems server as part of an integrated business system. Today, approximately 16,000 banks run their core banking and financial applications on an IBM i server.

No matter what industry they operate in, organizations store a wide variety of mission-critical information on IBM i, including:

- Financial data
- Personally identifiable information (PII) for employees and customers
- Payroll data
- Inventory levels
- Pricing information
- Customer lists
- Intellectual property
- Manufacturing processes
- Business strategies

Many organizations are subject to government regulations and industry standards, like Sarbanes-Oxley, HIPAA for the U.S. healthcare industry, PCI DSS for organizations that handle credit cards, and their equivalents around the world. Staying compliant provides a compelling reason for prioritizing a secure configuration.

Compliance and security are becoming even more important, as new requirements are added and more regulations are passed. PCI DSS began requiring penetration testing in 2015; multi-factor authentication became mandatory in February of 2018. The EU's General Data Protection Requirement (GDPR) took effect in May 2018, while the state of New York's cybersecurity law for financial institutions has been enforceable since March 2017.

Across the U.S., more states are moving to add to their existing cybersecurity regulations. California is the most notable example. In 2018, the state passed the California Consumer Privacy Act, which is set to take effect in 2020. As it stands now, this law includes many provisions similar to GDPR.

Security mandates don't cover all types of data, and not all organizations are affected by security mandates. But consider the consequences of leaking information that gives your organization a competitive advantage, such as pricing information or inventory levels. Businesses in highly regulated industries aren't the only ones that should be concerned with IBM i security.

Our Methods

To conduct this study, the HelpSystems security experts audit IBM i systems using our Security Scan. This software is free and runs directly from any network-attached PC without modifying IBM i systems settings, interrogating Power Systems running IBM i (System i, iSeries, AS/400) across **seven critical audit areas:**

- Server-level security controls
- Profile and password settings
- Administrative capabilities
- Network-initiated commands and data access
- Public accessibility to corporate data
- System event auditing
- Virus scanning

After the analysis is complete, the anonymous security statistics are returned directly to one of our servers. The software does not collect any application-specific data; therefore, no information is available regarding the purpose of the server. Participation in the study is optional.

This year's study includes 244 IBM i servers and partitions that were audited between January and December of 2018. The sample includes systems of varying sizes, the most common being 41A, which comprised 36 percent of the reviewed servers.

Organizations can voluntarily submit demographic information to the study. The organizations that supplied this information self-classified their industries as:

- Finance
- Insurance
- Retail
- Healthcare
- Manufacturing
- Technology
- Government
- Other (free format)

As in previous years, participating organizations spanned a broad range of sizes, but the sample is not random. Security officers or other staff at these companies were concerned enough about IBM i security to request a scan. This tends to result in a sample that are in the early stages of hardening their IBM i security.

Finally, this is not a recurring study of the same systems each year. Direct year-over-year comparisons cannot be made, but some general trends are apparent.

The average system scanned for the 2019 study has 1,228 users and 555 libraries. These numbers are a bit higher than the median because there were several large servers in the data sample (Table 1).

Table 1: Average System Size

System Size	Average	Median
# of Users	1,228	477
# of Libraries	555	378

The majority of scanned servers were running on supported versions of the OS; however, 27 percent were on V7.1, which IBM stopped supporting in April 2018. That's a significant change from last year, when 60 percent of systems studied were on V7.1. The number of systems running on V7.3 increased from just three percent last year to 45 percent this year. Systems running on V7.2 held steady at nearly 30 percent (Figure 1).

Figure 1: Installed Versions of IBM i

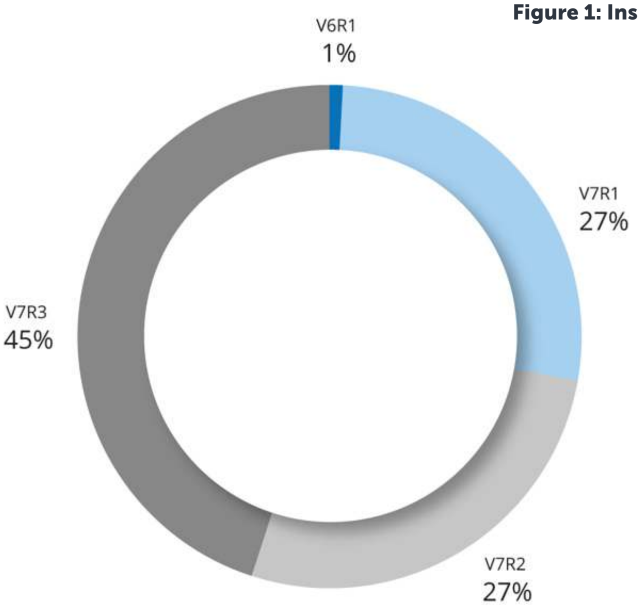

Setting the Tone: Basic System Security Levels

IBM i security best practices start with the configuration of numerous system values, which regulate how easy or difficult it is for an outsider to use or abuse your system. Poorly configured or unmonitored system values are an unacceptable security risk.

QSECURITY Level

What is it and what's the risk?

The system security level (QSECURTY) sets the overall tone, although it is often undermined by other settings. IBM recommends and ships security level 40 as the minimum, due to documented vulnerability found in level 30 and below. It should be noted that, despite the change to the default setting, a server migration will typically reload this to the same value as found on the previous generation of the server.

Power Systems servers can be configured at one of five different security levels:

- **Level 10** — No Security. No password required. User IDs are created for any user who attempts to sign on. IBM no longer supports level 10.
- **Level 20** — Password Security. Every user must have a valid ID and password. Every user with a valid ID and password assumes root-level authority (*ALLOBJ) by default.
- **Level 30** — Resource Security. Object-level authority is enforced as users do not assume root-level authority by default. A moderately knowledgeable programmer or operator can bypass resource-level security and assume root-level authority.
- **Level 40** — Operating System Security. Level 30 protection plus additional operating system integrity. It is possible for an extremely knowledgeable programmer with access to your system to elevate his or her level of authority, possibly as high as root-level authority.
- **Level 50** — Enhanced Operating System Security. Level 40 protection plus enhanced operating system integrity. A properly secured system at security level 50 is the best defense. However, even at level 50, other system configuration issues must be addressed.

What's the data?

Figure 2 shows the distribution of security settings on the systems included in the 2019 dataset. Out of the 244 systems studied, 24 percent were running system security level 30 and four percent were running at security level 20. Overall, nearly 30 percent fell short of IBM's recommended minimum level (Figure 2A). Many running on a sub-par security level are doing so without deliberate intent after having migrated their system values from an older server and are now recognizing the need to take corrective action.

Figure 2: System Security Level

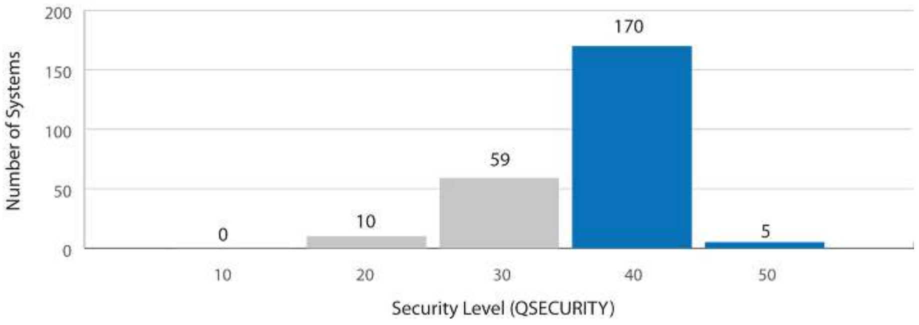

Figure 2A: Meeting the Recommended Minimum Level

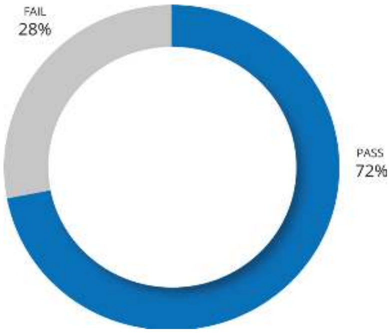

What's the solution?

Bringing your system up to QSECURITY level 40 or higher is a critical step toward protecting your system. People are now largely aware of the inherent risks of running below the minimum recommended level but remain cautious about changing settings that may prove problematic for the business applications. Outsourcing this task to IBM i security professionals like the team at HelpSystems is a way to eliminate quickly all the guesswork from the process. Discussions with HelpSystems security experts have assured many that changes to the security level can be made in a carefully architected manner to virtually eliminate the risk of an unplanned outage.

Key System Values for Restoring Objects

What is it and what's the risk?

Several other system values related to object restoration often remain at their shipped levels, reflecting a typical IBM i configuration of "load and go."

The system values in question are designed to work together as a filter that prevents restoration of malicious or tampered objects. But IBM i's default values fail to provide this protection, which may leave the system vulnerable.

What's the data?

The system values below work consecutively to determine if an object should be restored, or if it is to be converted during the restore:

- **Verify Object on Restore** (QVFYOBJRST)—More than 70 percent of servers are running below the recommended level of 3.

  This value, preset at level 1, controls whether a signature will be validated when a digitally signed object is restored.

- **Force Conversion on Restore** (QFRCCVNRST)—Nearly 93 percent of servers are running below the recommended level of 3.

  This value, preset at level 1, controls whether some types of objects are converted during a restore.

- **Allow Object Restore** (QALWOBJRST)—Only four servers had altered this system value from its default *ALL setting.

  This value controls whether programs with certain security attributes, such as system-state and authority adoption, can be restored.

DEFAULT SETTINGS RARELY PROVIDE THE DEGREE OF SECURITY ORGANIZATIONS NEED, AND SYSTEM VALUES FOR OBJECT RESTORATION ARE A PRIME EXAMPLE.

What's the solution?

Default settings rarely provide the degree of security organizations need, and system values for object restoration are a prime example. A proactive approach to system values starts with defining and implementing a security policy that incorporates the most secure settings your environment will tolerate. (Seek professional expertise if you are unsure of the impact of certain settings.) The free open source IBM i Security Standard from HelpSystems can help you get started defining your own policy.

Defining your security policy with Powertech Policy Minder for IBM i and reporting on exceptions will ensure that your system settings match your policy.

Users with Superpowers

What is it and what's the risk?

IT professionals require special authorities to manage servers. In addition to changing system configuration, these authorities may permit the ability to view or change financial applications, customer credit card data, and confidential employee files.

In careless, misguided, or malicious hands, these special authorities can cause serious damage. Recent data shows that data breaches caused by insiders are increasing in both frequency and cost. This type of breach includes users who take revenge on a previous employer by exposing corporate data along with users who unwittingly click on phishing emails. The insiders in these two cases have very different intentions, but the result is the same: private data exposed to unauthorized parties.

IBM i special authorities are administrative privileges and always pose a security risk, so auditors require you to limit the users who have these special authorities and carefully monitor and audit their use.

What's the data?

There are eight types of special authority in IBM i. Figure 3 shows the average number of user profiles for each special authority.

Figure 3: Powerful Users (Special Authorities)

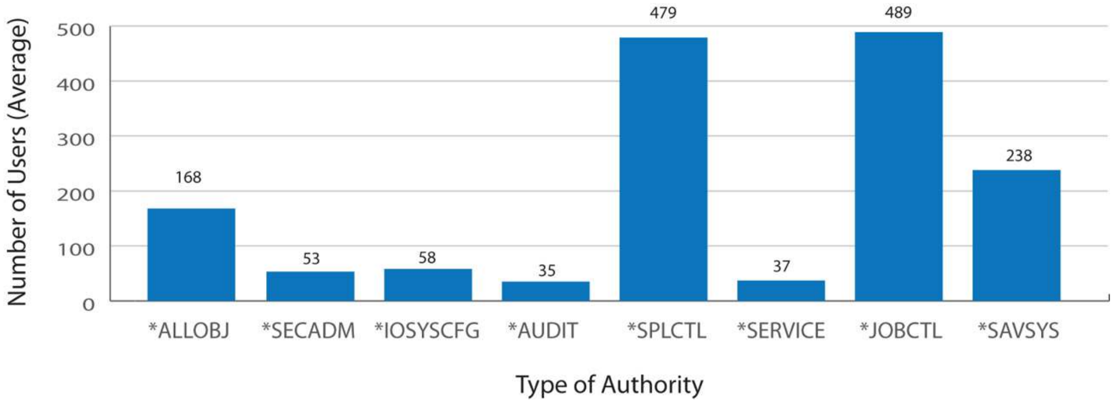

Of the special authorities, *ALLOBJ is the one providing users with the unrestricted ability to view, change, and delete every file and program on the system. As shown in Figure 3, this authority is granted to users in unacceptably high numbers.

Only seven of the systems reviewed had 10 or fewer users with *ALLOBJ authority. The most frequently granted special authorities were Job Control (*JOBCTL) and Spool Control (*SPLCTL), both of which were granted to nearly 40 percent of users. Job Control provides the capability to change the priority of jobs and printing, or even terminate subsystems in some cases. Spool Control enables users to fully access any spooled file in any output queue, regardless of imposed spool restrictions.

Many organizations have begun to embrace a role-based access control (RBAC) model in an attempt to standardize the user configuration. This is typically implemented on IBM i using a mechanism known as Group Profiles. In this study, all but three of the servers had one or more group profiles defined and 53 percent had 10 or more of them. Of all the group profiles, 95 percent were passing special authorities and excessive private permissions down to their members, an inheritance that is sometimes overlooked during a review of user permissions.

What's the solution?

IBM does not publish any documentation for the functions available with each of the special authorities, which leads to resistance by IT to remove authorities for fear of "breaking" existing operations.

While it is difficult to create a hard and fast rule for all environments, IBM i security experts agree that the number of users with special authorities, and private or public permissions, should be kept to the barest minimum. Regulations often permit assignment of administrative privileges only on an "as-needed" basis. This is known as the principle of least privilege.

In general, it's a good security practice to keep the number of users with special authority to fewer than 10, or less than three percent of the user community. Reducing this number is best done after consulting with an IBM i security expert, who can advise on ways to determine if authorities are necessary and suggest possible alternatives in marginal cases.

Here are some best practices for powerful users:

- Document and enforce separation of duties for powerful users.
- Avoid having any all-powerful users, all the time.
- Monitor, log, and report on the use of powerful authorities.
- Be prepared to justify the use of powerful authorities to auditors and managers.

To make the work of monitoring and documenting user privileges easier, a solution like Powertech Authority Broker for IBM i can automatically monitor, control, and audit users who need to access higher levels of authority. Powertech Command Security for IBM i is an effective way to prevent unauthorized users from executing a monitored command.

IBM i SECURITY EXPERTS AGREE THAT THE NUMBER OF USERS WITH SPECIAL AUTHORITIES SHOULD BE KEPT TO THE BAREST MINIMUM.

Faulty Passwords and Profile Security

User and password security issues are critical because they represent the most obvious—and most easily exploited—method to compromise your system.

Without proper user and password security measures in place, efforts to secure other areas of an IBM i network are largely ineffective. How can you be sure that the user signed on is the same user that the profile and password were assigned to?

Prime Targets: Inactive Profiles

What is it and what's the risk?

Inactive profiles are user profiles that have not been used in the last 30 days or more. They create a security exposure because these accounts are not actively maintained by their users, which make them prime targets for hijacking.

Many of these inactive profiles belong to former employees or contractors—people who might carry a grudge or who might find their former employer's data useful in their new roles at competitors.

The threat persists even if ex-employees never attempt to utilize these profiles. Other users within the organization might know, for example, that the former IT director's profile is still on the system. And whether an inactive profile is exploited by a former employee, a malicious insider, or a hacker, unusual use of the profile won't be detected and reported by the profile owner.

What's the data?

Figure 4 shows an average of 438 profiles (36 percent of the total) have not signed on in the past 30 days or more. Of these, 295 of them remain enabled and ready to be used.

Figure 4: Inactive Profiles

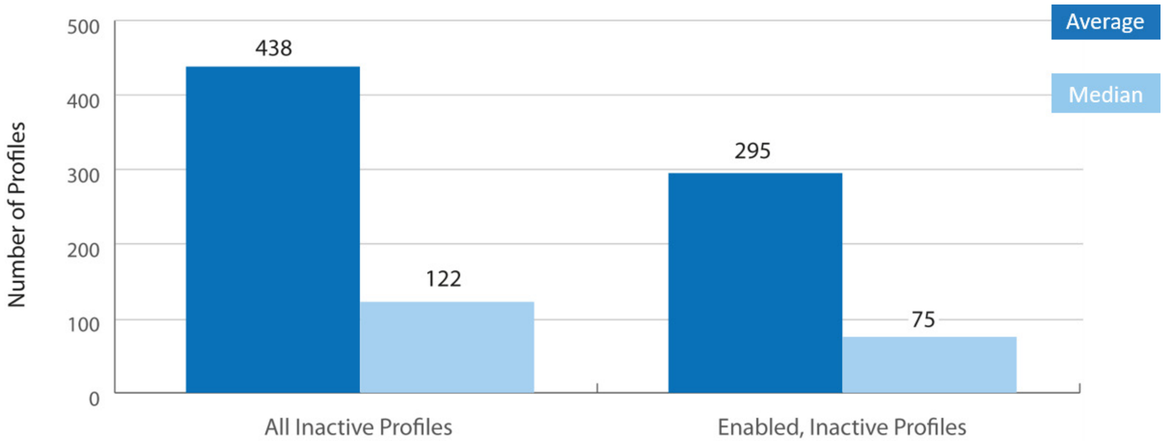

What's the solution?

Develop a process for inactive profiles. Start by defining how long a profile must be inactive before you take action (perhaps 60 days), disable the inactive profiles, and remove all special authorities and group profile assignments. Wait another 30 days to make sure the profile really is inactive before removing it from the system, or until the name of the user is no longer required for reconciling with the audit trail.

This process can be performed manually or automated using IBM's built-in security tools.

The Open Secret of Default Passwords

What is it and what's the risk?

On IBM i, profiles that have a default password have a password that's the same as the user name. Because this is the default when new user profiles are created, it is a particularly high-risk factor for IBM i servers.

Many companies have policies to name their user accounts or profiles based on a standard format, such as first name initial followed by surname (for example, jsmith or tjones). A hacker can guess profile names like jsmith and try default passwords. It's even easier for an employee who understands the internal convention for user profile names to guess account names and try default passwords, especially if they are aware of accounts that have been created but not yet used.

Regulatory and legislative standards typically mandate that users must utilize unique credentials known only to the user, ensuring that any actions can be tied to that specific individual. Organizations might struggle to prosecute illegal or unauthorized activity if it became evident that the credentials couldn't unequivocally identify the culprit. This prevalence of default passwords means guessing a password becomes an incredibly simple task and this ultimately translates into a compliance failing.

What's the data?

In this study, nearly 15 percent of user profiles have default passwords (Figure 5). Half of the systems studied have more than 30 user profiles with default passwords. 25 percent are even worse off, with more than 100 users with default passwords. One system has a total of 12,642 user profiles with default passwords and nearly four thousand were in an enabled state.

Figure 5: Default Passwords

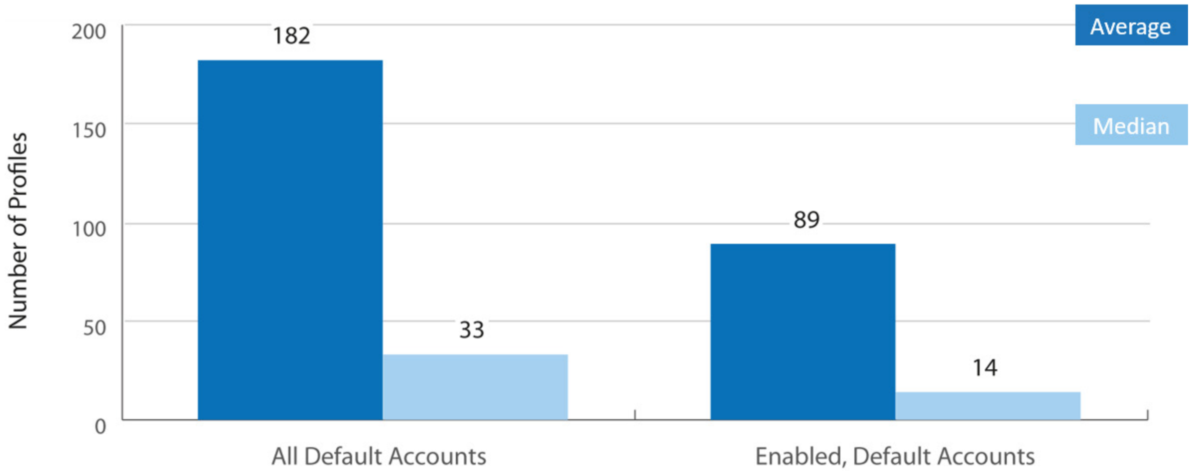

What's the solution?

Establish and enforce password policies that make it difficult to compromise a user's account. As of V7.2, IBM i supports the banning of default passwords via the QPWDRULES system value, although consideration must be given to applications or vendor software that creates profiles during installation.

Reporting tools like Powertech Compliance Monitor for IBM i make it easy to generate audit reports on a regular basis that compare IBM i user and password information against policy.

Setting a Standard for Minimum Password Length

What is it and what's the risk?

IBM i provides the capability to require a minimum length for passwords. Shorter passwords may be easier to remember, but they're also easier for others to guess. Although short passwords can be strengthened by using random characters, the odds of correctly guessing a four-character password are greater than a six-character password.

That's why password recommendations have evolved. Previous editions of the State of IBM i Security Study advised using six-character passwords, but NIST updated its recommendations at the end of 2017 and HelpSystems is following suit. Password best practices now call for eight-character passwords.

What's the data?

Figure 6 shows the setting for the minimum password value on the systems reviewed. According to our results, 41 percent meet or surpass the best practices standard of eight characters or more.

Regulatory compliance mandates, such as the Payment Card Industry Data Security Standard (PCI DSS), recognize the benefit of an even longer password; however, nearly 60 percent of servers in this study fail to satisfy PCI's requirement of seven-character passwords. Shockingly, 22 percent of systems permit users to select a password that is less than five characters long and 16 servers permitted the use of single character passwords.

IMAGINE WHAT COULD HAPPEN IF YOUR USERS WITH SIMPLE PASSWORDS HAVE SPECIAL AUTHORITIES OR ACCESS TO SENSITIVE DATA.

Figure 6: Minimum Password Length

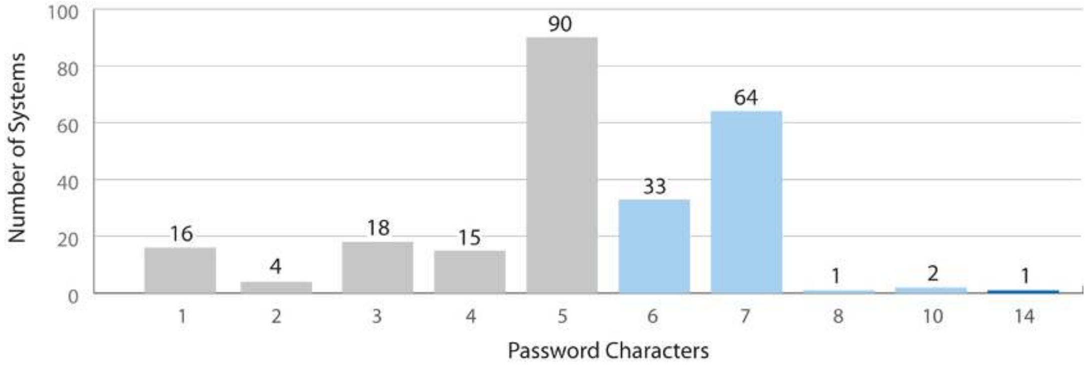

What's the solution?

Create a password policy that requires users to use eight or more characters in their passwords—and at least seven if your organization is required to comply with PCI DSS.

NEARLY 30 PERCENT OF SYSTEMS DO NOT REQUIRE PASSWORDS TO DIFFER FROM THE PREVIOUS PASSWORD.

Capitalizing on Other Password Settings

What is it and what's the risk?

IBM i allows systems administrators to define password policy at a granular level. Password settings include length, character restrictions, digit requirement, expiration time, and how soon a password can be reused.

These settings help make passwords harder to guess and increase the protection of your system, since simple, easy-to-guess passwords like "123456" and "password" that remain disturbingly common. Imagine what could happen if your users with simple passwords have special authorities or access to sensitive data.

What's the data?

The latest data shows that IBM i administrators aren't utilizing all the password controls available to them:

- 54 percent of systems don't require a digit in passwords, facilitating users' common desire to use simple (weak) dictionary words as passwords.
- 98 percent of systems do not impose any restrictions on characters. Simply restricting vowels would add extra security by preventing users from choosing simple, easily guessable words for their passwords.
- 43 percent of systems do not set an expiration time for passwords—users are never forced to change their password. This can also be controlled at a user level, but that is typically reserved for the exception profiles.
- Nearly 30 percent of systems do not require passwords to differ from the previous password. Only 34 percent require at least 10 unique passwords and just 13 percent require users to adhere to the maximum setting of 32 unique passwords.

IBM i V6.1 introduced QPWDRULES, a system value that provides a way of designating numerous password policy settings in a single repository. All systems included in this study have access to this system value, and 12 percent are now benefiting from the flexibility it offers.

Another value added in V6.1, QPWDCHGBLK, restricts the frequency with which users can voluntarily request a password change. This prevents users from repeatedly changing passwords to return to their favorite. The most common—and default—setting was *NONE, found on 90 percent of servers studied. Less than 10 percent of systems in the study specified a value of 1–24 hours, which is the range of settings HelpSystems recommends.

While good password controls are important, a password expiration policy is also valuable. Traditionally, the password expiration best practice is to set the expiration interval at a maximum of 90 days. According to systems in our study, the average password expiration interval is 95 days, and of the systems setting a password expiration interval, the most common value is 90.

What's the solution?

IBM i includes settings that give system administrators the power to require stronger passwords, but the latest password recommendations from NIST call for a re-evaluation of what "stronger" actually means.

Requiring your users to use a password of eight characters or more can make the password harder for hackers to crack and easier for users to remember. IBM i can even support passwords up to 128 characters, which are more accurately called passphrases. New passwords should be checked against a dictionary of banned words and the number of attempts allowed to enter it should be limited.

With these controls in place, it may not be necessary to set an interval for password expiration. Work with your auditors to determine the best policy for your system.

Multi-factor authentication is another way to protect your systems from unauthorized access. Users must verify their identities by supplying a possession (a one-time password or a YubiKey, for example) or an inherent characteristic (a fingerprint or retina scan) in addition to their password. Requiring users to authenticate in multiple ways reduces the risk associated with a weak or compromised password. A solution like Powertech Multi-Factor Authentication makes it easy to implement multi-factor authentication for IBM i.

Another option is eliminating passwords entirely by implementing single sign-on (SSO) based on technology that is included in the IBM i operating system. HelpSystems' knowledgeable services team can assist with deploying the configuration needed to leverage this popular elimination of passwords.

Forgotten Passwords and Other Invalid Sign-On Attempts

What is it and what's the risk?

Passwords are forgotten, mistyped, or simply mixed up with other passwords. Invalid sign-on attempts happen to everyone from time to time, and IBM i users are no different. Help desk personnel charged with resetting these passwords often work with the same users over and over. How do you track which users have multiple invalid sign-on attempts? What if your powerful profiles are targeted?

A single invalid attempt, or even a handful of unsuccessful tries, may not be cause for concern. But what if your system had one user profile with hundreds or even thousands of invalid sign-on attempts?

Larger numbers could indicate an intrusion attempt, while three, five, or even ten attempts are probably the sign of a frustrated user.

It's also possible that attempts numbering in the thousands or even hundreds of thousands are the sign of a broken application, perhaps one lacking a built-in mechanism to recognize when its attempts to connect to the server are being denied. But that assumption should never be made without investigation. And a broken application is still an application that is not fulfilling the business purpose for which it was written.
The level of risk increases significantly if the offending profile is determined to be, for example, QSECOFR, and is not disabled automatically, or if the security team has no way to be notified of failed access attempts in a timely manner.

What's the data?

Virtually every system had at least one profile that had experienced an invalid sign-on attempt, which is not surprising. 51 percent of systems had a profile that had experienced more than 100 denied attempts. 25 percent had more than 1,000 invalid sign-on attempts against a single profile. One system in our study had 1.3 million attempts against a single profile.

It's worth noting that the count of the number of attempts continues even if the profile is disabled. There is nothing an administrator can do to prevent the attempts as long as the user can obtain a connection to the server. For that reason, the most important element is timely detection and notification.

Figure 7 shows the action taken when the maximum number of allowed sign-on attempts is exceeded. In 90 percent of cases, the profile is disabled and this is always recommended. When using explicitly named devices (as opposed to virtual device names) the recommendation is expanded to include disablement of the device description. It is not recommended to disable virtual devices, as the system typically creates a new device when the user reconnects. The device setting does not apply to all connections, such as ODBC and REXEC services.

The other 10 percent of servers disable the device, but leave the profile enabled. This creates risk if the user re-establishes a connection, or perhaps connects to a service that does not require a workstation device

Figure 7: Default Action for Exceeding Invalid Sign-On Attempts

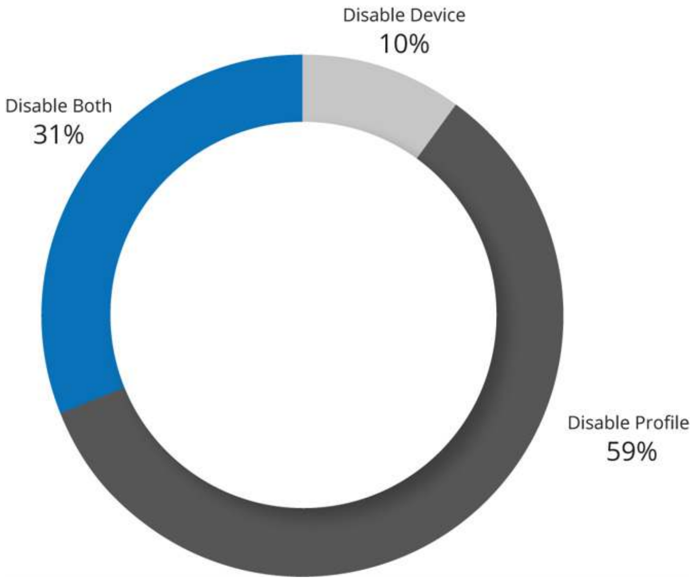

What's the solution?

Timely notification and investigation of an unusually large number of denied access attempts is a critical first phase in the detection of a possible compromise.

Too often, data breaches hit the headlines accompanied by a startling revelation over just how long the breach was permitted to occur. An organization cannot stop a breach if they don't know it's happening, and invalid sign-on attempts are one of the most obvious indicators.

To protect your system, make sure profiles are disabled by default after the maximum allowed sign-on attempts is exceeded.

A tool for self-service password resets can help the users who have truly forgotten their passwords. Powertech Password Self Help for IBM i is one option that makes it easy for IBM i users reset a password and it sends instant alerts to designated personnel when unsuccessful resets occur. This allows administrators to take corrective action before any damage is done.

A multi-factor authentication solution can also protect your systems by requiring another credential in addition to a password. Powertech Multi-Factor Authentication for IBM i is one such option that can help protect servers inside and outside the firewall, as required by PCI DSS.

Data Access Goes *Public

What is it and what's the risk?

On non-IBM i servers, users who are not granted permission to an object or task typically have no authority. With IBM i, this is not the case. Every object has a default permission that is applicable to non-named users, known collectively as *PUBLIC.

Unless the user is granted a specific authority—granting or denying access—then the user will operate with the object's default permission. This isn't a problem until we discover that this default is initially set by IBM and is sufficient to allow a user to invoke a program and to read, change, and even delete data from a file.

In other words, unless proactive steps are taken to restrict *PUBLIC access rights, users who have not been granted a specific authority to an object or task can read, change, and delete data.

This situation creates a risk of unauthorized program changes and database alterations—a red flag for auditors, who recommend that users should not be authorized to read or change production databases or source code without a proven business requirement.

What's the data?

This study uses the *PUBLIC access rights to libraries as a simple measurement indicating how accessible IBM i data would be to the average end user.

Figure 8 details the level of access that *PUBLIC has to libraries on the systems in our study. If *PUBLIC has at least *USE authority to a library, anyone who can log in to the system can get a catalog of all objects in that library, and use or access any object in the library. Assuming that the user or *PUBLIC also has the necessary authority to the specific object, they may even be able to delete objects from the library.

Based on a common lack of security on individual objects, users with PC access can often download (read) and manipulate (change and delete) data stored in the library. The FTP GET function or ODBC operations in tools like Microsoft Excel may allow even a novice end user to access your data.

***CHANGE** authority to a library allows the user to place new objects in the library and to change some of the library characteristics.

***ALL** access allows anyone on the system to manage, rename, specify security for, or even delete a library (if they have delete authority to the objects in the library).

Figure 8: *PUBLIC Authority to Data

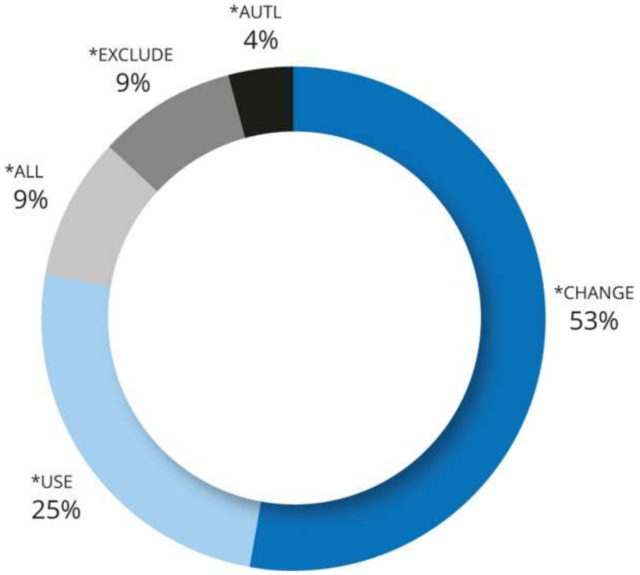

Our findings demonstrate that IBM i shops still have far too many libraries accessible to the average user. The statistics for DB2 libraries indicate a lack of adequate control over the data, which often includes critical corporate financial information.

The method used to determine what authority *PUBLIC will have to newly created files and programs typically comes from the library's Default Create Authority (CRTAUT) parameter.

Figure 9 indicates that 19 percent of libraries reviewed had Default Create Authority set to *USE, *CHANGE, or *ALL. However, 79 percent of libraries deferred the setting to the QCRTAUT system value (*SYSVAL).

Figure 9A extends the library level assignment of *SYSVAL and reflects that the system value typically remains at the shipped default of *CHANGE. In fact, just nine percent of servers are configured to default to the deny-by-default requirement of common regulatory standards such as PCI.

This means that when new files and programs are created on these systems, the average user automatically has change rights to the vast majority of those new objects. On these systems, non-named users have the authority to read, add, change, and delete data from the file. These same users can copy data from, or upload data to, the file, and even change some of the object characteristics of the file.

An added twist to this concept occurs when a user profile is created with permissions granted to the general user population (*PUBLIC). When *PUBLIC permissions exceed the strongly recommended setting of *EXCLUDE, this is known as an "unsecured profile." It is possible for an alternate user to run a job that leverages the privileges of the unsecured profile. This activity will not be logged by the operating system as a security violation as it is deemed permissible at all security levels. 183 systems have at least one unsecured profile and 45 systems have 10 or more profiles that are publicly accessible.

Figure 9: Default Create Authority by Library

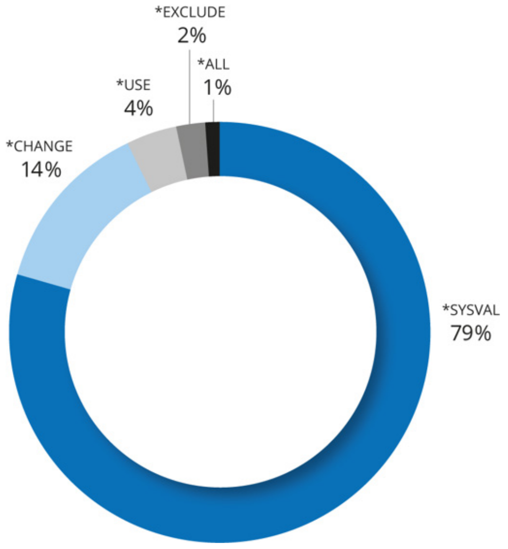

Figure 9a: Default Create Authority by System

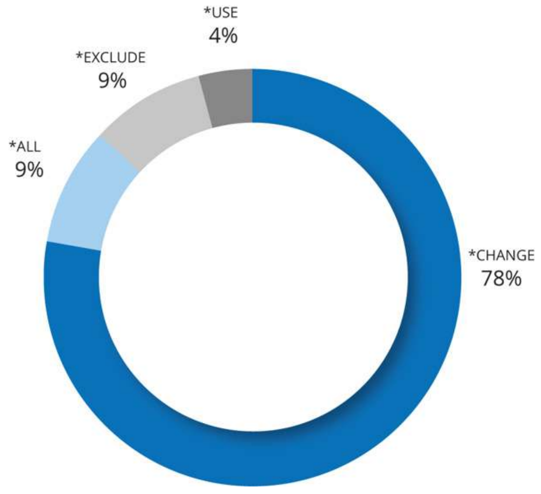

A powerful IBM i programming technique known as "adopted" authority enables a program to perform functions that calling users may not be authorized to run on their own. The QUSEADPAUT system value defines which users can create programs with the user adopted authority (USEADPAUT(*YES)) runtime attribute. Only nine systems impose a restriction on the setting of this value.

What's the solution?

With virtually every system user having access to data far beyond their demonstrated need, administrators need better processes to control access to IBM i data.

In the 2019 IBM i Marketplace Survey, IBM i pros reported that one of their top security challenges is the struggle between cybersecurity controls and business efficiency. It's true that securing your data has the potential to make it harder for end users to access the information they need. But a data breach can bring day-to-day operations to a standstill.

There's a clear need to prioritize cybersecurity and implement security tools that provide users with secure, frictionless access to the data they need. Tools like Authority Broker, Powertech MFA, and Password Self Help do just that.

And don't overlook the security capabilities of the IBM i OS. Where possible, secure data using resource-level security to protect individual application and data objects.

When it is not possible or practical to protect data with resource-level security, or when an additional layer is desired, use exit program technology to regulate access to the data by network services. Powertech Exit Point Manager for IBM i is an industry leading off-the-shelf IBM i exit program solution.

Monitor changes to your database information. Powertech Database Monitor for IBM i creates before-and-after snapshots of database changes and can even require users to sign for changes, so you can meet compliance requirements.

Investigate how well your third-party software suppliers use operating system resource level security. Seek assistance from the vendor in protecting application objects.

Finally, ensure that application libraries are secured from general users on the system. (Although it requires some planning, consider setting the System Value and Library values for Default Create Authority to the most restrictive setting [*EXCLUDE].)

Through the Networks' Back Doors

What is it and what's the risk?

Over the years, IBM has extended the power of IBM i by adding tools that allow data to be accessed from other platforms, especially PCs. Well-known services such as FTP, ODBC, JDBC, and DDM are active and ready to send data across the network as soon as the machine is powered on. Any user with a profile on the system and authority to the objects can access critical corporate data on your Power Systems server.

This is possible even when administrators do not purposely install data access tools on users' PCs. End users can download free tools from the internet or even use tools included with other software loaded on their PCs to access sensitive data. For example, Windows comes with FTP client software that easily sends or retrieves data from an IBM i server.

Worse yet, the results from the data access section indicate that object-level authority is poorly implemented on most systems. The combination of open access rights to data, overly powerful users, and convenient tools to access the data from a PC is a perfect storm of IBM i security exposures.

Beyond data access, some TCP services permit the execution of server commands. The easily-accessed FTP service enables certain commands like Delete Library (DLTLIB) to be run by all users—even those without command line permission on their profile. This is still a shock to many system administrators and unknown to many managers and auditors.

What's the data?

The statistics in Figure 10 show that REXEC, a TCP/IP application that allows users to submit commands to a remote system, is often not automatically started. FTP, however, is almost always active and listening. This means most users are only a few keystrokes away from using it to send data across the network.

Figure 10: REXEC and FTP Autostart

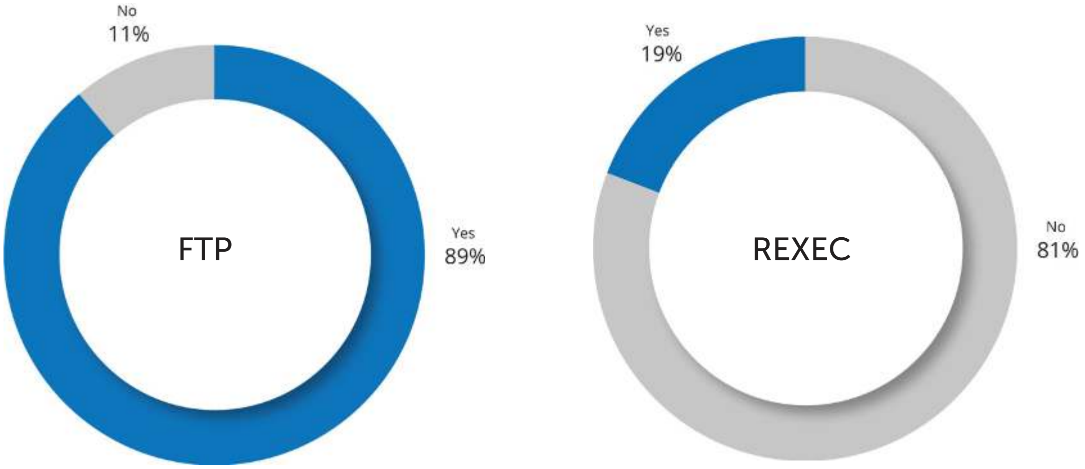

To reduce this serious exposure, IBM provides interfaces known as exit points that allow administrators to secure their systems. An exit program attached to an exit point can monitor and restrict network access to the system.

An exit program should have two main functions: to audit access requests and to provide access control that augments IBM i object-level security. The study assumes that all designated exit programs satisfy both of these minimum requirements.

HelpSystems reviewed 27 different network exit point interfaces on each system to check whether an exit program was registered. More than 70 percent of the systems have no exit programs in place to allow them to log and control network access (Figure 11).

Even on the systems with exit programs, coverage is often incomplete. Of the systems with programs in place, 10 percent have only one or two registered exit programs, while just eight percent have programs registered to all of the network access exit points.

Having very few exit programs registered is highly indicative of self-policing via a client authoring their own exit programs. While it suggests that the concept of registering exit programs is grasped, coverage is often limited to just those servers that are considered most notorious or risk-laden.

In this study, 8 percent of servers have more than five exit programs but fewer than the 27 indicative of the presence of a commercial-grade exit program solution. The most common exit point covered was FTP Server Request, followed by FTP Client Request. ODBC/JDBC data access functions remain largely unmonitored, further exposing the server to the risk of transparent data loss.

MORE THAN 70% OF THE SYSTEMS HAVE NO EXIT PROGRAMS IN PLACE TO ALLOW THEM TO LOG AND CONTROL NETWORK ACCESS.

Figure 11: One or More Exit Programs in Place

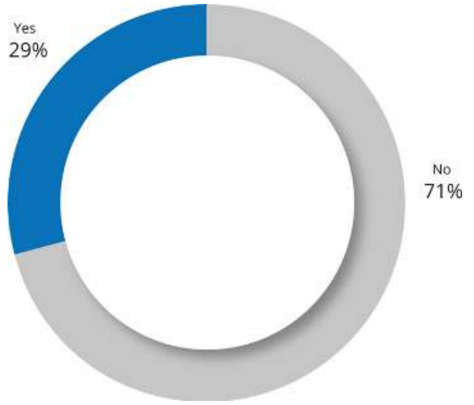

What's the solution?

This data contrasts with the 2019 IBM i Marketplace Survey, where just over one-third of survey takers reported having exit point security in place. Although this study and the survey have different participants, IBM i users who think their exit points are secure should do their due diligence and verify the appropriate controls are in place. One quick way to identify unsecured exit points is to request a free Security Scan from HelpSystems.

Without exit programs in place, IBM i does not provide any audit trail of user activity originating through common network access tools such as FTP and ODBC. Even companies that have installed exit program solutions to protect their data frequently neglect some of the critical access points.

It appears that many companies in the IBM i community are dangerously unaware of the wide-open network access problem. The lack of monitoring and control of network access is a serious deficiency in many shops.

IBM i shops can write their own exit programs or use packaged software to accomplish this task. The advantage of using a commercial exit program solution like Powertech Exit Point Manager for IBM i to monitor and control users' access through network interfaces is that you get broader coverage that protects all critical access points.

Users Crossing the Command Line

What is it and what's the risk?

The traditional way to control access to sensitive data and powerful commands was to limit command line access for end users. And in the past, this method was effective.

In addition to configuring the user profile with limited capabilities, application menus controlled how users accessed data and when they had access to a command line. However, as IBM opens new interfaces that provide access to data and the opportunity to run remote commands, this approach isn't as sound as it used to be.

What's the data?

According to our 2019 results, 72 percent of users have command line access through traditional menu-based interfaces. Of those users, we found that 40 percent of the profiles were enabled.

Several network interfaces do not acknowledge the command line limitations configured in a user profile and must be controlled in other ways. This means that users can run commands remotely, even when system administrators have purposely taken precautions to restrict them from using a command line.

THIS MEANS THAT USERS CAN RUN COMMANDS REMOTELY, EVEN WHEN SYSTEM ADMINISTRATORS HAVE PURPOSELY TAKEN PRECAUTIONS TO RESTRICT THEM FROM USING A COMMAND LINE.

What's the solution?

Based on the broad *PUBLIC authority demonstrated in the Data Access Goes *Public section, anyone on these systems can access data, commands, and programs without the operating system keeping a record.

Start addressing this problem by reviewing network data access transactions for inappropriate or dangerous activity. Be sure to establish clear guidelines for file download and file sharing permissions. Remove default DB2 access in tools like Microsoft Excel and IBM i Client Access.

Straying from the Audit Trail

What is it and what's the risk?

One of the significant security features of IBM i is its ability to log important security-related events into a tamper-proof repository—the Security Audit Journal. This feature allows organizations to determine the source of critical security events, such as:

- Who deleted this file?
- Who gave this user *ALLOBJ authority?

This information can make the difference between responding promptly to a security event and discovering a breach *after* significant damage has occurred.

The challenge is that the volume of data contained in the Security Audit Journal is so large and the contents so cryptic that most IT staff have trouble monitoring the logged activity with the tools available in the operating system.

Making sense of security data is the crucial second half of the auditing equation. Without a way to turn the data into meaningful, actionable information, organizations run the risk of missing important warning signs.

What's the data?

16 percent of the systems reviewed do not have an audit journal repository, indicating a very low level of scrutiny.

In a related statistic, almost 20 percent of systems are operating with the QAUDCTL system value setting at its shipped value of *NONE (Figure 12). This is the master on/off switch for auditing and globally blocks any system or object level events from being logged, regardless of the existence of the system audit journal.

Figure 12: IBM Security Audit Journal in Place

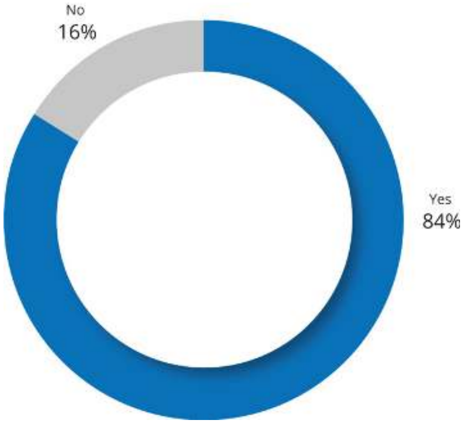

The absence of the IBM Security Audit Journal indicates a very low level of scrutiny for the system in question. In addition, the presence of the journal on a system where the QAUDCTOL system value is set as *NONE suggests that administrators fired up the auditing function but subsequently turned it back off or perhaps were unaware of the necessity for additional configuration.

There is also inconsistency in the types of events that are audited. Some configurations suggest that auditing has been activated by high availability (HA) applications that must replicate events to backup systems. Audit event types such as *AUTFAIL (Authority Failures) are not required in an HA infrastructure and often identify that customers are not using the auditing facility for security purposes.

When organizations have activated the Security Audit Journal, it's unclear how much insight the extensive data is providing them. A few software vendors provide auditing tools that report on and review the system data written to the Security Audit Journal. But only 24 percent of the systems in this study have a recognizable tool installed.

What's the solution?

On most of the systems surveyed, security violations could occur undetected. Companies that use the Security Audit Journal are in a far better position than those that don't because, at any time, they can use an automated tool to sift through and interpret the audit journal entries.

Given the voluminous amounts of raw data collected in the IBM Security Audit Journal, it's not realistic to expect system administrators to manually review the logs regularly. The job of filtering and analyzing massive amounts of complex raw data requires software tools.

ONE BUSINESS SCANNED IBM i FOR VIRUSES FOR THE FIRST TIME, AND WAS SHOCKED TO FIND NEARLY 250,000 FILES INFECTED BY THE CRYPTOWALL VIRUS.

At the same time, many organizations are overwhelmed by the amount of reporting required to demonstrate compliance with regulations such as Sarbanes-Oxley (SOX) and the Payment Card Industry Data Security Standard (PCI DSS), yet it appears that very few of them take advantage of the tools that are available to automate and simplify reporting tasks.

Utilizing a software auditing tool like Powertech Compliance Monitor for IBM i reduces the costs associated with compliance reporting and increases the likelihood that this work will get done. Implementing Powertech SIEM Agent for IBM i will feed IBM i security data into your Enterprise Security Solutions that support Security Information and Event Management (SIEM) or Syslog formats, allowing you to identify and analyze security events quickly.

Susceptible to Viruses and Malware

What is it and what's the risk?

One of the more controversial IBM i security topics is the risk posed by viruses and other malicious programs. While the traditional IBM i library and object infrastructure is considered to be highly virus-resistant, it is acknowledged that other files structures within the Integrated File System (IFS) are susceptible to hosting virally infected files, which can then be propagated throughout the network.

Recognizing this reality, IBM created system values and registry exit points to support native virus scanning a number of years ago.

One business scanned IBM i for viruses for the first time, and was shocked to find nearly 250,000 infected files. If anyone doubted the need for virus protection, this example proves the risk is real.

Scanning IBM i for viruses and malware is becoming increasingly popular as administrators start to recognize that IBM i contains file systems that are not immune to infection and, under certain circumstances, native applications and even IBM i itself can be impacted. This will likely become more commonplace in the future, considering the many high-profile malware attacks in recent years.

What's the data?

We record the configuration of primary scan-related system values (QSCANFSCTL and QSCANFS), as well as the QIBM_QP0L_SCAN_OPEN exit point (scan on stream file open) and QIBM_QP0L_SCAN_CLOSE (scan on stream file close).

The exit point QIBM_QP0L_SCAN_OPEN allows an exit program (antivirus engine) to be registered that will intercept all file open attempts and scan the file before it can be opened. This ensures you will not spread an infection outside the IBM i environment.

When the servers were reviewed for antivirus controls, 11 percent were scanning on file open, which is a noticeable increase over prior years. This means the other 89 percent are at risk of having internal objects impacted or of spreading an infection to another server in their network (Figure 13).

Figure 13: Scanning on IFS File Open

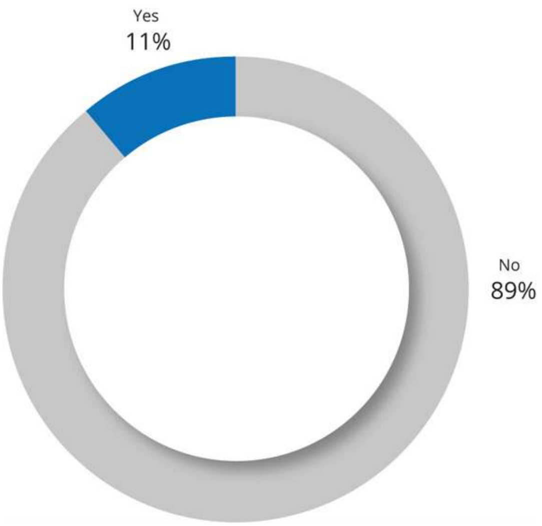

All of the systems assessed for virus scanning use the OS default value for system value QSCANFS (*ROOTOPNUD). This setting allows for stream files in the root(/), QOpenSys, and user-defined file systems to be scanned for virus threats if a scanning application were installed.

Of the systems scanning for malware, just three are optimized for performance by configuring system value QSCANFSCTL to include value *FSVRONLY. If the values include *FSVRONLY, only access through the file server will be scanned for malicious threats and file access from native jobs will not trigger a scan or negatively impact performance.

What's the solution?

Although IBM i itself may not be infected by PC- or Unix-based malware, the objects stored in the IFS can be. In addition, any active virus will operate with the authority of the user that activated it and therefore can impact native objects through actions such as object renaming or deletion.

If you are using IBM i as a file server, you must take action to ensure infected objects do not have the ability to execute on a Windows server. A proper antivirus defense would provide detection, removal, and prevent the spreading of infection beyond the current environment.

Register an exit program to exit point QIBM_QP0L_SCAN_OPEN to intercept file open attempts from the network and scan files before they are opened. This prevents viruses from spreading outside the IBM i environment.

In addition, utilizing an exit program, such as the one provided with Powertech Exit Point Manager, registered to the QIBM_QPWFS_FILE_SERV exit point can help limit actions of remote viruses operating on other servers on the network.

Once again, the data in this study is vastly different from the 2019 IBM i Marketplace Survey, where 32 percent of survey taker claimed their IBM i servers were protected with antivirus solutions. This suggests that some IBM i professionals might not realize the importance of using a solution that's built for IBM i's unique features. Using a native IBM i virus scanner is more secure, faster, and more reliable than a PC-based scanner.

CONCLUSION

IBM i has developed a reputation as one of the most securable platforms available. One of IBM i's great advantages is that sophisticated tools for securing, monitoring, and logging are built into the OS. But experts agree that IBM i security is only as effective as the policies, procedures, and configurations put in place to manage it.

This study highlighted a number of common security exposures and configuration management practices that must be addressed to protect the data on IBM i systems.

While all organizations could improve the IT controls on their IBM i servers, deciding which controls to improve first can be a challenge. No system became vulnerable overnight, nor is it possible to fix every security problem in a single day. **What's important is starting somewhere and making continued progress toward a stronger security profile**.

While every system faces unique challenges, in general there are four top priorities for IBM i security. If you're unsure how to proceed, start with this list:

- System Security: Check the QSECURITY level and make sure it's 40 or higher
- Security Auditing: Enable QAUDJRN and find a tool to help interpret it
- Network Access: Register the most common exit points like FTP and ODBC first
- Reduce unneccessary user privileges

What most experts recommend is starting with an assessment of vulnerabilities to understand where your system security stands today and how it could be improved. Security professionals with IBM i expertise and user-friendly software solutions are available to make this project faster and easier. HelpSystems offers a range of options, from a very thorough Risk Assessment to a quick, no-charge Security Scan.

Once you have all the information, you can begin formulating a plan that addresses your organization's security vulnerabilities. And from there, security will become business as usual—not a moment of panic after a failed audit or a data breach.

HELPSYSTEMS IS HERE TO HELP WITH IBM i SECURITY

Check how secure your IBM i is with a Security Scan from HelpSystems. Security Scan is free, fast, and reveals your system's security gaps. Our Security Advisers can then help you formulate a plan to remedy your security vulnerabilities.

APPENDIX

As the leading expert in IBM i security, HelpSystems has developed an extensive line of powerful solutions designed to address shortcomings in the operating system, provide advanced functionality in access control and auditing, and ease the cost and burden of maintaining regulatory compliance.

Table 2 outlines the available security modules and their purpose.

Table 2: Comprehensive Suite of Security Solutions

SECURITY SOFTWARE

Powertech Compliance Monitor for IBM i	Custom auditing and reporting
Powertech Exit Point Manager for IBM i	Access control by exit programs
Powertech Authority Broker for IBM i	Management of privileged users
Powertech SIEM Agent for IBM i	Real-time security reporting
Powertech Database Monitor for IBM i	Real-time database monitoring
Powertech Command Security for IBM i	Command monitoring and control
Powertech Identity Manager for IBM i	Centralized user profile management
Powertech Antivirus for IBM i	Advanced native virus detection
Powertech Password Self Help for IBM i	Self-service password reset
Powertech SecurID Agent for IBM i	Two-factor authentication for IBM i
Powertech Policy Minder for IBM i	Automated security policy enforcement
Powertech Risk Assessor for IBM i	Comprehensive security assessment
Security Scan	Free IBM i snapshot
Powertech Encryption for IBM i	Encryption and key management
GoAnywhere	Managed file transfer
Powertech Multi-Factor Authentication	Multi-factor authentication software for IBM i
Powertech Event Manager	Real-time cybersecurity insight and response

SECURITY SERVICES

Risk Assessment	Detailed vulnerability assessment
Penetration Testing	Unbiased test of vulnerabilities
Architecture	Security action plan
Remediation	Architecture implementation
Managed Security Services	Monthly monitoring and reports
Single Sign On Managed Service	Risk-free SSO implementation

ABOUT THE AUTHOR

Robin Tatam is a midrange industry veteran with three decades of IBM i experience. He is an IBM Champion, an award-wining speaker/subject matter expert in security for COMMON, and a member of their Speaker Excellence Hall of Fame. Robin is certified with ISACA as a Certified Information Security Manager (CISM) and the co-author of IBM's Redbook publication on IBM i data encryption.

ROBIN TATAM

Director of Security Technologies

About HelpSystems

HelpSystems is the leading expert in automated security solutions for IBM Power Systems servers, helping users manage today's compliance regulations and data privacy threats. Our security solutions and services save your valuable IT resources, giving you ongoing protection and peace of mind.

Because Power Systems servers often host sensitive corporate data, organizations need to practice proactive compliance security. As an IBM Advanced Business Partner with an expansive worldwide customer base, HelpSystems understands corporate vulnerability and the risks associated with data privacy and access control. HelpSystems security solutions and services are the corporate standard for IBM i security at many major international financial institutions.

HelpSystems has demonstrated a proven commitment to the security and compliance market and leads the industry in raising awareness of IBM i security issues and solutions, leveraging the experience of the world's foremost IBM i security experts, Robin Tatam and Carol Woodbury.

- HelpSystems is a member of the PCI Security Standards Council, a global open standards body providing guidance to the Payment Card Industry Data Security Standard. HelpSystems works with the council to evolve the PCI DSS and other payment and data protection standards.
- HelpSystems is a member the IBM i Independent Software Vendor (ISV) council.
- HelpSystems publishes an open source IBM i Security Standard as a part of its mission to promote awareness of common security challenges and ensure the integrity and confidentiality of IBM i data.

helpsystems

www.helpsystems.com

About HelpSystems
HelpSystems is a leading provider of systems & network management, business intelligence, and security & compliance software. We help businesses reduce data center costs by improving operational control and delivery of IT services.