OUTPUT MANAGEMENT

Cyber Security and Output Management

CONTENTS

EXECUTIVE SUMMARY

This white paper presents a comprehensive overview of the current state of vulnerabilities in print and output management, then outlines strategies for businesses to enhance their security posture regarding this oft-overlooked aspect of computing. The paper examines the evolving nature of threats and discusses the steps businesses must take to mitigate the risk of data loss. It underscores the importance of educating employees on best practices and investing in technology solutions that provide a layered security approach. Finally, the report outlines the regulatory and certification environment governing information security, offering guidelines to ensure compliance with these requirements. Overall, the paper provides valuable insights for business leaders looking to safeguard their operations against the ever-increasing threat of data loss in output management.

Dynamic Shift in Security

Companies, governments, and individuals are increasingly plagued by a wide range of security attacks including Phishing, MITM attacks, viruses, and ransomware. The historical way of guarding against virus attacks was to harden a shell around central resources and attempt to prevent penetration attacks. These so-called castle-and-moat security systems are commonly viewed as insufficient for the modern threat environment. They indeed present a tempting target to individuals, groups, or even state-sponsored attackers since a hardened shell implies that there is something important behind the wall.

The consistent failure of the castle-and-moat approach has caused a shift in thinking about cyber security. If the castle itself cannot be safeguarded, then there must be some better way to protect the corporate information. A de-perimeterization of resources has come to the forefront of security thinking.

This philosophy and architecture is now known as Zero Trust Architecture (ZTA).

The core concept is that if the castle cannot be guarded, then each individual component and transaction must be safeguarded.

The simple definition is to “never trust, always verify” each system, user, and transaction in all business workflows.

Alongside this development, a need for authorization and authentication above the traditional LDAP models was required in order to accommodate the growth of common open internet resources like Facebook and Google. This development has generated the OAuth and OpenID Connect standards. These methods first authorize the user, and then provide a method for the ongoing authentication (or perhaps denial) of that user’s data. As such, the user is authorized by some source (like Microsoft Entra, Ping ID, OKTA or others) and once authorized, is provided a “token” to include with all continued transactional work. That way, each transaction can be examined for the token, and trust is then verified at that granular level. Along with high-level TLS encryption methods, this level of security may suffice, but an organization really does need to evaluate the need for other complex options like firewalls, Intrusion Detection Systems, and Virtual Private Networks (VPNs). LRS® software allows for all company workflows to proceed via the open internet with or without other complex options.

Many areas of IT struggle to keep pace with this trend. In 2019, LRS started a journey to bring output management and the enablement of print and scan throughout this de-perimeterized world. LRS has the unique advantage of having workstation-resident software that can capture output before it is sent from the local spool as well as agents that run in multi-function printers (MFPs) and server technology that can run on-premise or in any cloud space.

This philosophy and architecture is now known as Zero Trust Architecture (ZTA). The core concept is that if the castle cannot be guarded, then each individual component and transaction must be safeguarded.

This allows LRS to meet and exceed all of the authentication and encryption needs of a pure zero-trust philosophy.

Vulnerabilities in Output Management

Print spooling, devices, and workstation methods create a number of vulnerabilities that are both known and predictable based on experience in other areas. Denial of Service attacks, PrintJack exploits, default passwords, unsecured network connections, outdated firmware, unsecured print jobs, remote access, unsecured document storage, poor administrative security, and lack of encryption (both at rest in devices and spools and in motion to the printer) pose significant risk to an organization's well-being. A few of these are discussed here.

Because print processes are integrated into the operating systems in many areas, and because they require system-level authority to do their job (and also contain communication protocol needs like Server Message Block, aka SMB), they are a tempting target for virus / ransomware distribution and infection.

Vulnerabilities in Ad Hoc Office Printing

Ad hoc printing is often the only area that companies consider in a discussion of output management and printing. Ad hoc printing is best described as users sending everyday documents, emails, and reports to printers. Some of this may be necessary, such as a letter or a check, but often ad hoc printing is more discretionary in nature. In some cases, it may be a waste of company resources, and should be controlled using enforceable software policies.

During the last three decades, the technical community predicted a reduction and elimination of ad hoc print. Pressures brought on by the pandemic and worker preferences and demands have pushed the workforce into a far more mobile stance than was predicted in 2019. Yet the printed page remains a strong force in the office.

Print processes are often controlled, and indeed dictated, by operating systems like Windows, macOS, UNIX / LINUX, and large platforms such as z/OS. Because print processes are integrated into the operating systems in many areas, and because they require system-level authority to do their job (and also contain communication protocol needs like Server Message Block, aka SMB), they are a tempting target for virus/ransomware distribution and infection. One main reason is that these processes have both the authority to distribute malware and the built in mechanism to spread it.

The prime example of this is PrintNightmare. PrintNightmare exposed a set of vulnerabilities in the Windows print mechanisms, and the attempt to patch these vulnerabilities made many companies lose their ability to print for a time. Because of the underlying, low security mechanisms provided by the Windows print mechanisms, this virus created a near-panic situation in businesses around the world. Shutting down all print for security reasons brought many organizations to a standstill.

Electronic vulnerabilities are not the only area that merit consideration in the ad hoc printing environment. In many organizations, stacks of paper found next to printers or in wastebaskets often contain vital information. For instance, one large healthcare organization was required to pay a $49 million dollar settlement for, among other things, Protected Health Information (PHI) found in the trash. It is clear that ad hoc printing needs careful controls to help prevent this kind of thing. Policy printing is one area to be considered. Another is so-called “pull printing.” Pull printing is the practice of requiring employees to physically go to a printer and authenticate themselves using a card, user credentials, or other mechanisms to complete the printing process. That way, the user is standing by the device as the paper comes out. This does not eliminate the risk but does provide an audit trail as to when a print occurred and who picked it up.

The pull-print process may introduce vulnerabilities in the electronic realm. For many pull-print products, jobs are stored on the operating system spool until release is initiated. As discussed, these system-level processes are vulnerable. The space where jobs are stored is unencrypted. Moreover, it is unlikely that the underlying drivers are capable of encryption to the device.

Because OS printing is integral to the process, it is critical that jobs, once rendered, be captured, encrypted, and retained either in a locally encrypted file or in an encrypted central spool. It is also important in both cases that these captured jobs can be sent via TLS-encrypted mechanisms to the end device. With pull printing in particular, the OS spool introduces a physical vulnerability. If a job is printing, and the printer runs out of paper or otherwise breaks, what happens if paper is not added until later or the printer repair is delayed? The OS spool will just happily send the rest of the data. Instead, a managed spool (either central or local, yet removed from the OS spooling system) can sense a printer outage and stop the job from printing while the device is unsupervised.

The devices themselves can indeed be vulnerable. Lexmark, HP, Brother, Xerox, and others have published many firmware updates to harden their devices, but organizations can be lax on applying these updates. It is important to remember that printers are not simple devices today, but rather complex processors with strong CPUs, ample memory, and persistent storage. They often use public domain platforms like Android and as such, employ JAVA routines in their core. Hackers are clever and flexible and can use these platforms as a base of attack. It is critical to keep these devices updated, use current SNMP network protocols, and change default passwords to discourage access to these powerful print devices.

Vulnerabilities in Business-Critical Printing

The systems that create and deliver business-critical documents offer very tempting targets for outside attackers. The reason is simple: business-critical output contains some of the most important information assets in the organization, including customer identities, revenue or billing amounts, private health and/or financial data, etc. Such valuable information can be directly exploited for financial gain or used to damage the reputation of a market competitor. Depending on the industry, information from business critical documents can also be used for certain types of espionage.

Business-Critical printing is output that drives workflows. Some typical examples are:

- Manufacturing and distribution
  - Box labels and enclosures
  - RFIDs
  - Safety and regulatory documents
- Healthcare
  - Patient information documents
  - Prescriptions
  - Armbands
  - Lab / blood labels
- Banking
  - Checks
  - Batch output

There are many other examples. Many of these are driven by automated workflows, EMR systems, ERP / WHM systems, and other business-critical workflows. Even though many processes are dependent on these for day-to-day work, they escape notice as companies think about print.

These areas share some of the same vulnerabilities discussed above but are compounded by the complex systems in which they run. Often these systems employ OS-style printing or have their own antiquated methods for driving these devices.

It is unlikely that malware penetration will find a path in these older methods, but using simple internet attacks can halt these processes and workflows at the simple print level with catastrophic results costing companies thousands of dollars per minute for failed print.

To avoid the inherent problems with these systems, it is important to remove / intercept output data from these systems at the earliest opportunity and use an output management system to ensure that workflows do not fail. Once again, providing an encrypted method of transferring data to the central spool, encrypting the data at rest, and then driving output devices with TLS-encrypted data streams allows for both safer transport and reducing the risk of disruption.

Mitigation and Defense

More than five years ago, LRS made a concerted effort to address the challenges of business-critical output. As a result, we have incorporated best practices into our products and processes to help organizations reduce or eliminate threats from these many vulnerabilities. This section will highlight the security measures incorporated into LRS's holistic output management solution set.

The LRS Approach

The first steps taken by LRS involved scrutinizing our internal processes. Many new vulnerabilities originate with the use of external routines. LRS has a mix of programming environments that suit the specific product, portability and platform needs including assembler, 'C', and C++. LRS also uses .NET, OpenSSL, jQuery and similar tools to speed development and keep pace with modern standards. With each patch level of our code, we ensure that we are using the most current stable version of these products. Microsoft Corporation and large open-source projects watch security closely, and they close gaps in their code when, or more often, before they are discovered.

Internally, LRS uses third-party code-scanning products to double-check any coding practices that can introduce vulnerabilities. This occurs several times in any release cycle, and our development and security staff carefully address any areas that are uncovered by this process.

Penetration Testing, also known as Pen Testing, is done both internally, by third parties, and by our customer base. Because our products are often exposed to the open internet and to internal resources like printers, mobile devices and workstations, the pen testing must be robust and extensive. This testing is conducted many times each year. As issues are found, they are researched and remediated. No environment can guarantee fully secure software, but LRS does create secure software through our dedication to secure design, testing, and by following industry best practices.

Transactional Security

Several years ago, LRS noted that normal print traffic was woefully behind the times in terms of security. Outside of the LRS environment, many vulnerabilities regularly occur, including:

- Improper use of operating system-level spooling
- Failure to encrypt data at rest
- Failure to encrypt data in motion both to spooling entities and to devices
- Failure to identify the correct user ID to assign to the job
- Failure to provide transactional authentication to print jobs
- Failure to recognize and capture output from systems outside of the local workstation
- Failure to provide access control over print
- Failure to integrate with existing Identity Provider systems

These may seem par for the course. In many senses they are indeed the norm. But output and print processes need to be hardened before they are exploited.

The following pages list some methods that LRS uses to shore up vulnerable output practices such as those listed above.

The LRS Personal Print Manager (PPM) solution intercepts rendered data from the operating system. Removing data from the spooling system mitigates the many vulnerabilities that have been exposed in these OS-level processes.

Operating System-Level Spooling

Printing at the workstation level necessarily involves operating system (OS) spooling. That is the one and only way that applications running at the workstation level can take the basic graphic and text information displayed on the user's screen and change the data into a format that a printer can understand. As such, these systems are necessary to do our work.

These spooling subsystems, however, do not need to be the end-all solution. The LRS Personal Print Manager (PPM) solution involves a robust set of processes that intercept rendered data from the operating system. As will be further discussed, this data, once captured, is treated with the same high-level care as any other transaction on the system. Removing data from the spooling system mitigates the many vulnerabilities that have been exposed in these processes.

Encrypting Data at Rest

OS-level spooling simply writes files to the local file system. These files may include rendered data (PCL, PostScript, PDF, etc.) or published and known graphic formats like GDI+. This data contains all the information that was called for at print time including confidential employee data, sensitive competitive information, patient health information and almost any of the numerous forms of data that companies seek to protect.

PPM captures the data rather than consigning it to the OS methods of file writing. Instead, PPM encrypts the data locally before storing it in its own local repository. From there, the data may be further sent to a central spool, to the printer, or left in the encrypted repository waiting to be retrieved by the user at the printer using Pull Print mechanisms. This eliminates the vulnerability of unencrypted data sitting in the spool.

Pull Printing functionality is offered by many companies. This is the practice of storing the rendered job and waiting for a user to physically go to the device and authenticate before the job is printed. Authentication is often done via a card tap, though many other mechanisms are available from a variety of vendors including LRS. Almost all other vendors, however, store the rendered job in the local file system as described above – unencrypted and vulnerable.

Encrypting Data in Motion

Any rendered jobs that are in the spool are then forwarded on to devices (usually by using print-driver technology) to the printer. These data streams are the same as were spooled in most cases and are well-known data formats that are easy to read and exploit.

This seems surprising. Encrypted print protocols like IPPS have long been available, and many modern printers support their use. Because the PPM software captures the data and drives output devices directly, it can and does send the data using the IPPS protocol. This data is TLS encrypted using certificates and methods commonly used on the Internet today. This approach protects the data in transit and allows for end-to-end encryption of data.

Failure to Identify the Correct User ID to Assign to the Print Job

In order for output to be audited and jobs to be properly pulled, the correct user ID must be assigned to the job. At first blush, this seems simple; just use the user ID from the OS.

Times have changed. Many organizations have adopted a Bring Your Own Device (BYOD) philosophy. When I, as a user, provide my own device, chances are good that I built my own login, or worse, used the default for that device. For business application access, this poses no issues as most apps themselves have authentication front ends that use some central user ID mechanism like AD, LDAP, SAML or OIDC sources. Sometimes, these systems have a backend print mechanism that uses the credentials provided. More often, output is forwarded to the browser and is handled by the local system as described earlier. If that browser is on a BYOD device, this local print will carry whatever ID that the user created – “Admin” being a common one – and will not have any association with the corporate credentials.

This means that Audits will provide information of dubious value – there may be many employees with the ID of Admin – and pull print becomes impossible. A card tap at a device authenticates with corporate credentials, but the print job is stored without them.

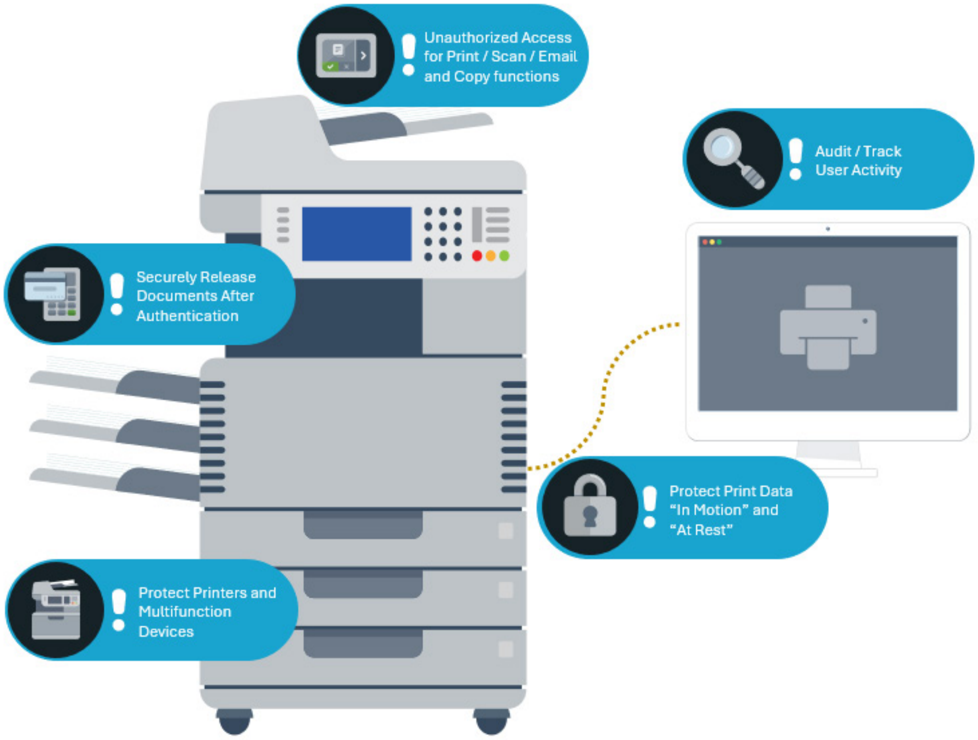
Unauthorized Access for Print / Scan / Email and Copy functions
Audit / Track User Activity
Securely Release Documents After Authentication
Protect Print Data “In Motion” and “At Rest”
Protect Printers and Multifunction Devices

PPM resolves this issue. PPM works with the LRS/Gateway component to protect print. When a user starts PPM or initiates print, PPM will ensure that the user is authenticated before proceeding.

Failure to Provide Transactional Authentication to Print Job

When PPM is first used either by accessing the local application or during the generation of print, it starts a process of authentication. PPM connects with the LRS/Gateway which itself acts as an OIDC provider. The LRS/Gateway, in turn, provides access to whatever authentication method the company uses (AD, LDAP, OIDC) such that those well-known systems actually do the real authentication of the user. Once authenticated, the LRS/Gateway, using standard OIDC methods, provides an Access token back to PPM to include in all further transactions. Any print data that goes to the central VPSX® spool will be required to carry that token to ensure that this is the actual user, and that the user is properly identified. Also, in line with this approach, any jobs that PPM sends (or releases) to the system locally will be assured of having the correct user ID rather than the value stored in the OS.

This OIDC token method also allows for transactions to be authenticated without the constant nuisance of repeated logins. Tokens are encrypted and stored locally for use. Expired tokens can be refreshed within the bounds of the company’s configurations and needs. This robust process allows for full encryption, authorization, and authentication of each and every print job in the workstation arena.

Failure to Recognize and Capture Output From Systems Outside of the Local Workstation

Back-end print from systems such as SAP, Epic, Oracle Millennium, E-Business suite, and countless other applications is among the most vulnerable data of all. Unfortunately, data kept in these systems is often the most important information asset in any organization. Competitive data, sensitive employee and patient data, and corporate / government secrets are stored here. Print from these platforms is almost always unencrypted both in motion and at rest. LRS has worked with many of these systems to harden this vulnerability using the same methods – intercept the rendered data, encrypt it, and get it to the device. This aspect of security is, without fail, ignored by most other vendors.

Failure to Provide Access Control Over Print

Once authenticated on the network, it's important to be able to group like users together by department, division, location, etc. LRS provides security access panels to let print administrators group print access lists, enabling users to display, update, and configure various print options. Whether it's mobile printing, pull printing, PPM direct printing, or scanning, LRS can customize information access to fit the customer requirements.

Failure to Integrate with Existing Identity Provider Systems

As mentioned above, The LRS/Gateway component provides access to whatever authentication method the company uses (e.g., AD, LDAP, OIDC) such that those well-known systems actually do the real authentication of the user. LRS offers this capability so user access management can remain with the customer and, through the use of AD groupings, can actually integrate with the LRS security group structure.

Conclusion

Many areas of vulnerability are uncovered throughout the IT landscape every day. Output management has only recently come under the microscope due to high-profile examples of exploitation. In contrast, LRS has focused on security measures including encryption for well over 20 years.

It is important to protect your governmental and corporate print data with the same care as other data produced and stored in your organization. LRS strives to make this easier for organizations of every size and every sector.

LRS® OUTPUT MANAGEMENT