
A WHITEPAPER ON SECURITY

Cloud Services Security in PaperCut MF

by PaperCut Software – March, 2020





Risky business	4
Location, location, location	5
Cost	5
Security	6
Scalability	7
Reliability	7
Accessibility	7
Layer upon layer	7
Physical	8
Administrative	8
Technological	9
PaperCut MF's Cloud Service	10
Document processing	10
Document path	11
Actions at the MFD	11
The PaperCut MF Application Server	11
The PaperCut Cloud Service	12
Encrypted storage	12
Delivery to the cloud storage provider	12
Data retention	13
User authorization	13
Actions at the MFD	13
The PaperCut MF Application Server	14
PaperCut Cloud Service	14
Unauthorized user	14
OAuth	14
Authorized user	15
Good to go	15
Home sweet home	15



ATFAQs (answers to frequently asked questions)	16
Business process	16
Software development	16
Vulnerability awareness	16
Process and procedure	17
Conclusion	17
Glossary	18
Authors	20
Contributors	20
PaperCut HQ	20



Data security in the cloud from the ground up

Securing important data is a centuries-old pursuit.

In 1586 an English nobleman named Anthony Babington began sending messages to the imprisoned Mary, Queen of Scots. The content of the correspondence centered around a plot to assassinate Queen Elizabeth and install Mary on the throne of England. Babington used a number of techniques to keep his messages as secret as possible since treason was rather frowned upon.



In today's cyber world we would refer to his approach as layered security. First, he encrypted the messages, then hid them in the corks of beer barrels and then had them transported through a network of trusted accomplices. Unfortunately for Babington, he used a weak cipher and his network was also compromised. Within a few months, he paid the highest price for his lack of data security.

At PaperCut we pay special attention to multiple aspects of data security, and it's designed into everything we build. You might also want to read the [Securing your Print System](#) whitepaper.

Risky business

Protecting a valuable asset always has some inherent risk. Usually, the goal of data security is to take the risk down to the lowest acceptable level while also maintaining business objectives and containing the cost. So we should always keep these three items in mind: risk, objectives, and cost.

We could create a data security plan that has absolutely zero risk by making the data inaccessible to anyone, but this would likely not meet business objectives or would come with an exorbitant cost.

Electronic data introduces some extra risks to consider. It used to be that to steal an asset, thieves had to physically steal it. Now, with electronic data, they can just make a copy, and in most cases the copy is as valuable as the original. And even worse, there's often little evidence that the copy has been taken so you don't even know you've been robbed.



In 1980 there was a string of thefts from commercial aircraft in the US. Quite by chance, the case was solved on May 14th in Atlanta, Georgia, when a cargo trunk popped open during loading. The crew discovered a stowaway hiding inside the trunk with some food and an oxygen tank. It turns out that he would crawl out of the trunk mid-flight, grab whatever valuables he could, and then take them back into his trunk with him. Compare this physical theft of valuables to the theft of electronic data—we'd call his actions a classic *Man in the Middle* attack and theft of data in transit. Luckily, stealing physical items made the airborne theft quickly evident, but with digital assets, the theft usually goes unnoticed for a period of time.

To properly identify your business objectives for data security can take quite some time. The objectives need to include legal requirements (for example, GDPR, HIPAA), operational efficiency for things like Business Continuity planning and supply chain, and basic items such as who requires access to the data and how quickly. We could cover some of those objectives by putting all the company data on a public web server, unencrypted. It might be highly efficient for business, but could be disastrous for reputation, confidentiality, and integrity. So your business objectives need to cover those, too.

Electronic assets like Protected Health Information (PHI), Personally Identifiable Information (PII), credit card data, intellectual property, and others must be secured in a way that minimizes the risk of loss, exposure or copying.

Location, location, location

A key factor in your decision to secure electronic data is where to locate the stored data. Typically, this means choosing locally hosted (that is, on-premise) or in the cloud. However, the decision is not as clear cut as it used to be, mainly due to significant advances in cloud security.

Let's break down the data location decision into five objectives.

Cost

Locally hosted data centers have significantly higher upfront and running costs than private cloud data centers (that is, Infrastructure as a Service¹). You'll need servers, disks, licenses, network gear, electrical power, cooling systems, and some amazing IT staff to manage, maintain, and update it all. A private cloud data center, on the other hand, has a very low cost of entry, but you will pay for just about every kilobyte of disk and cycle of a processor you use. You'll also bear most of the license, management, and software upgrade costs. Yes, this can also add up over time, but there does come a breakeven point (usually measured in years) where private cloud has a higher total cost of ownership. Other costs to consider and balance against the business objectives include scalability, legal compliance, and cash flow.

¹ [IBM, Learn IaaS](#)



Note: PaperCut MF Cloud Services are included as part of your license when you have active Maintenance and Support.

Security

This one used to be a slam dunk for locally hosted data centers, but that was “yonks” ago (Australian for a “quite a few years”). A locally-hosted data center gives you complete control over where and how to secure the data. For some organizations, this control is non-negotiable for certain data they handle. One reason is to comply with data protection regulations (for example, GDPR, HIPAA, FERPA). Compliance with regulations cannot always be ascertained for third-parties that have access to the data in the public cloud.

Locally hosted data only goes where you tell it to, and is only accessible to people and systems you authorize unless you store it on a disk or put it on a network. Because if anyone can access the data, they can steal it. Consider a few statistics on data breaches:

- ▶ 22% of breaches in 2017 use stolen credentials² (Verizon, 2018 Data Breach Investigations Report)
- ▶ 93% of malware entered via email¹
- ▶ Supply chain attacks were up 78% in 2018³
- ▶ 27% of breaches are a result of human error, 25% from system glitches, and 48% are malicious attacks (many of which were “insiders”)⁴

Most organizations that host their own data are no longer able to keep pace with the rise in quantity and sophistication of threats. Very few are prepared to defend against hardware chip vulnerabilities, zero-day exploits, multi-vector attacks, or polymorphic malware. We are at the point where most dedicated cloud providers offer greater security than self-hosted. However, there are still significant concerns for data security in the cloud, with the majority of them falling into the category of how the cloud is used. For example, millions of data records have been exposed in the last couple of years due to poorly configured Amazon S3 buckets⁵. Spoiler alert: if you don’t configure authentication and encryption, it turns out, others can see your data.

One more area where breaches are reported as “online access” happens in the supply chain. These breaches can happen to third parties that play a role in your business operations and that have access to your data. The majority of organizations do not have adequate third-party vetting procedures⁶. Your responsibility for data protection includes those with whom you share that data. A school that provides student information to a third party that prepares its meals is as vulnerable as the weakest link in the supply chain. A hospital that shares patient data with a third-party medical bill collector is vulnerable to and liable for PHI breaches. Before you trust your data protection responsibilities to a cloud provider, ensure they guarantee security of all third parties.

² Verizon, 2018 Data Breach Investigations Report

³ Symantec, 2019 Internet Security Threat Report

⁴ IBM, 2018 Cost of a Data Breach Study

⁵ [Bitdefender, Business Insights Blog](#)

⁶ [Help Net Security, Third Party Cyber Risk Management](#)



Scalability

Self-hosted infrastructure loses the battle to cloud unless we include the network in the reliability metric. Your network should have superior performance and reliability compared to the internet. Where cloud architectures excel is in their design as a massive pool of resources that can be allocated to consumers on an as-needed basis. If you have a spike in web traffic, transactions, reporting requests, etc., it's not a problem for a cloud service to scale up and down with demand. For scalability, cloud is a clear winner.

Reliability

Reliability for cloud services, excluding the internet connection to get there, is most often measured as an uptime percent. For every hour in the year, subtract all time required for hardware and software upgrades, repairs, patches, failures, migrations, etc., or any other reason that services were offline. Divide that number by 8,760 (the hours in a year) and you get your uptime percentage. On the low end, dedicated cloud services have 99.95% uptime. That's only four hours and twenty-three minutes per year of downtime. Cloud wins on reliability.

Accessibility

Data is only useful if you can access it when you need to, from the location you need to. Have you ever left home without your mobile phone? It's inaccessible and useless and you're a basket case. If your goal is to maximize accessibility to data (think 24x7, global) then cloud is a good choice. On the other hand, if you want to maintain a high level of control over accessibility due to confidentiality, liability, compliance, etc., then locally hosted could be the right choice.

Layer upon layer

Now that we've covered some of the differences between cloud and locally hosted data, let's dive a little deeper on the security topic. Data security isn't just one thing; it's layers. Security layers are often categorized as the:

- ▶ *physical layer*—things like fences, vaults and armed guards.
- ▶ *administrative layer*—covers training procedures, security audits and job rotation.
- ▶ *technological layer*—usually what we think of first and is made up of passwords, encryption and firewalls. Okay, there's a lot more to it than that, but you get the idea.

The concepts are the same as securing non-virtual assets, like precious stones, in that they are put in a safe, inside a locked building, with guards that patrol on regular intervals. Most of the time these multiple layers thwart would-be thieves.

From 2011 to 2017 there was a string of near-perfect jewelry store robberies in the southeast of the US, mostly in Florida⁷. The crew had sophisticated methods for getting around multiple layers of security. They rarely struck in the same police precinct more than once so patterns couldn't be determined. They carried electronic jamming equipment to confuse alarm systems. They picked

⁷ *The Atlantic*, December 2019, The Rise and Fall of an All-Star Crew of Jewel Thieves



jewelry stores in strip malls that were adjacent to shops that had little to no physical security (such as alarms, bars, or cameras). They broke into the insecure shop and used a simple drywall saw to create an opening to the jewelry store. Once inside the target store, their drill-man went to work “cracking supposedly impenetrable safes in record time, often drilling just one precise hole.” The crew specifically targeted stores with weak security layers.

PaperCut MF’s Scan to Cloud Storage and Document Processing services are designed for the cloud, with data security center stage. It’s hosted on Google Cloud, which uses multiple industry-leading security layers, and even layers within layers⁸. Any cloud provider you choose for your data to pass through or be stored in should employ comparable measures.

Here’s a brief summary of PaperCut MF’s Scan to Cloud Storage security layers that result from running on Google Cloud.

Physical

There are multiple options for physical layers to protect data centers, such as biometric identification, metal detection, cameras, vehicle barriers, and laser-based intrusion detection systems. All servers and network equipment are custom-designed by Google and include security chips to verify identity and integrity. And here’s a mind-blowing fact: a lot of “the cloud” is underground.

PaperCut’s cloud service is built on Google Cloud, and their data centers have an impressive list of compliance and certifications: ISO/IEC 27001, ISO/IEC 27017, ISO/IEC 27018, SOC 1, SOC 2, SOC 3, PCI DSS, CSA STAR, HITRUST CSF, FedRAMP, IRAP, UK's Cloud Security Principles, NIST 800-53, HIPAA, NHS Digital Commercial Third-Party Information Governance, and GDPR.

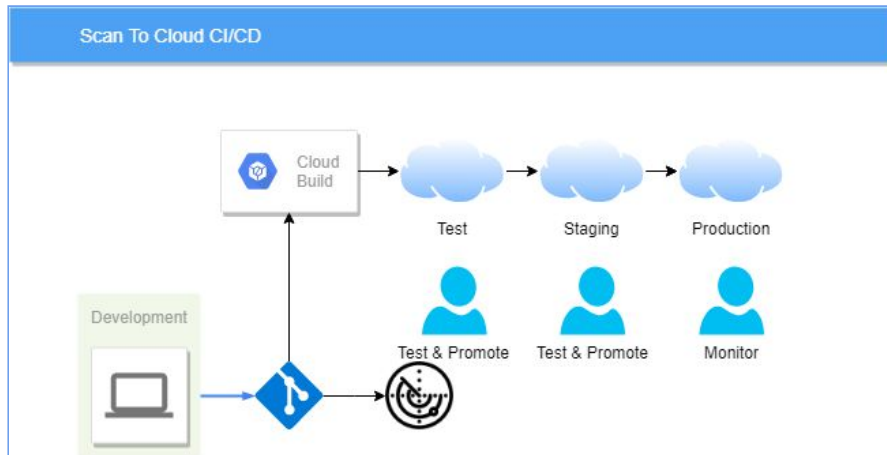
Administrative

When it comes to data breaches, people are the worst offenders. If we combine human error, insiders, and malicious attacks, people account for 75% of breaches⁹. That’s why Google and PaperCut allow a very small percentage of their employees to be in roles that have access to customer data. Our organizations also do thorough background checks, require regular security training, and put an audit trail on everything they do.

PaperCut also has well-defined software development practices to prevent security defects from ever seeing the light of day. Best practices for Continuous Integration and Continuous Delivery are strictly followed. This means PaperCut also has continuous testing (automated and human) and multiple stages between development and what a customer eventually uses. This helps catch mistakes before they make it into a release.

⁸ [Google Infrastructure Security Design Overview](#)

⁹ IBM, 2018 Cost of a Data Breach Study



PaperCut uses multiple stages in development and test for best practice CI/CD

That helps a great deal with PaperCut’s code, but what about all those third-party libraries and packages? PaperCut also tests those with code analysis, ad-hoc pen-testing, and active engagement with NIST, CVE, and other groups.

Technological

Here again it’s a multi-layered approach. The Google Cloud service itself is protected against fraudulent endpoints (that is, a bad actor trying to connect to the PaperCut API), network intrusion, data theft, and a host of other threats like malware (even at the hardware level) and DoS attacks. Check out the full white paper from Google¹⁰. These protection measures can stop virtually all attacks at the network level.


Other layers are added to protect customer data; things like encryption in transit and encryption at rest. “Google encrypts and authenticates all data in transit at one or more network layers when data moves outside physical boundaries not controlled by Google”¹¹. This means your data is never in clear text whenever it is moving between your site and PaperCut MF’s Scan to Cloud Storage service. In fact, it’s always using the highest level of TLS (Transport Layer Security) available.

Data is also encrypted at rest (even data needs a rest sometimes). This encryption isn’t just a single-pass encryption, there are more layers here, too. First, each data file is split into multiple blocks, then each block is encrypted using a different key, and then each encrypted block is stored on separate encrypted drives. Any data thief is facing the ultimate jigsaw puzzle where you can’t see the pieces or even know how many there are¹².

¹⁰ [Google Infrastructure Security Design Overview](#)

¹¹ [Google's Encryption in Transit](#)

¹² [Encryption at Rest in Google Cloud Platform](#)



Google's multi-layered file encryption

PaperCut MF's Cloud Service

Now let's take a look at the data flow and the security at each point for PaperCut MF's Scan to Cloud Storage service. We'll pull back the curtain for document processing, the document path, and user authorization.

If document processing is configured to use PaperCut MF Cloud Services, just about every scan will follow the document path. Even Scan to Email will send the scanned image to the cloud for document processing, if those options are selected in PaperCut MF's admin console.

Document Processing	Hosting Configuration
Document Processing is a collection of features to enhance and automate scanning. It includes OCR (Optical Character Recognition), Batch Splitting and Blank Page Removal (configured per Scan Action), and Despeckle and Deskew global settings.	<input checked="" type="radio"/> Use PaperCut MF Cloud Services for Document Processing (default) <input type="radio"/> Use Self-Hosted Document Processing (requires additional setup)

PaperCut MF's Document Processing is configurable for cloud or self-hosted

One very important bit of information to remember before we get into the weeds, is that PaperCut MF users can only scan to allowed destinations. The PaperCut MF administrator has complete control over scan destinations and document processing options. Using the PaperCut Cloud Service is entirely optional for each customer of PaperCut MF. No documents or metadata are sent to the PaperCut Cloud Service without explicit configuration by the PaperCut MF administrator.

An administrator must create a "Scan Action" to determine what each user sees when they login to the multifunction device (MFD). If the administrator has not created a Scan Action for Scan to Cloud Storage or Document Processing in the cloud, then the scanned images will not be sent to the cloud at all. For even greater control when setting up Scan Actions, access can be granted at the user or group level.

Document processing

The PaperCut Cloud Service has advanced document processing features, which include OCR, despeckling, deskewing, and others. Enabling these features requires the document to be unencrypted in the Cloud Service process for a very short amount of time. These features are optional and can be disabled by your PaperCut MF administrator.

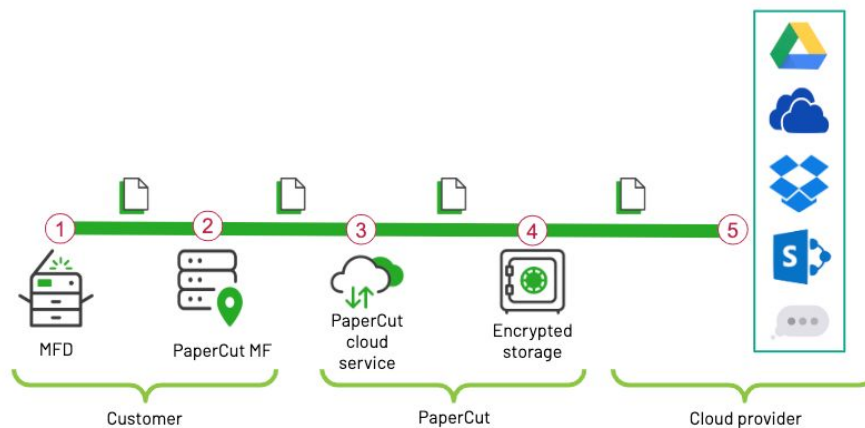
Your administrator can also set document processing in the PaperCut Cloud Service to be retained within a geographic region. Currently, there are three options: Australia, Germany, and the USA. Documents only leave the selected region upon delivery to the destination if the Scan Action specifies a cloud provider that is hosted elsewhere (for example, Evernote).



Document path

The life of the document for Scan to Cloud Storage begins at the MFD in PaperCut MF's Integrated Scanning feature. We utilize the in-built capabilities of the MFD to create an image of the document.

Note: All communication for the life of the scanned document uses TLS 1.2 unless otherwise noted.



The document path in PaperCut MF Scan to Cloud Storage

1. Actions at the MFD

- The user authenticates to PaperCut MF.
- User selects an available Scan Action with a cloud destination.
- User scans their document.
- The MFD creates the scanned image.
- The MFD sends the image document to the PaperCut MF server.

Note: Weak ciphers can be disabled in PaperCut MF so that an MFD is required to use more secure ciphers.

2. The PaperCut MF Application Server

- PaperCut MF receives the image document and metadata from the MFD.
- PaperCut MF creates a secure connection to the PaperCut Cloud Service. PaperCut MF uses only TLS 1.2 Perfect Forward Secrecy compliant cipher suites.
- The PaperCut MF server is authenticated and authorized to the Cloud Service. This is done with a unique combination of the PaperCut MF installation id and its license key.
- PaperCut MF sends the image document and related metadata (for example, the user's email address) to the PaperCut Cloud Service.



Note: Important considerations such as securing the PaperCut MF application server and file system, IT staff training, firewalls, etc. are the responsibility of the organization hosting PaperCut MF.

3. The PaperCut Cloud Service

- a. The PaperCut Cloud Service receives the document and metadata uploaded by the PaperCut MF application server. This is where all the magic happens.
- b. Depending on the options the PaperCut MF administrator has selected, the uploaded document is opened for processing. This includes OCR, despeckle, and deskew.
- c. The PaperCut Cloud Service writes the document to encrypted storage.
- d. The PaperCut Cloud Service is secured by and hosted on Google's Cloud Platform.
- e. The PaperCut Cloud Service runs entirely in the region selected by the customer (Australia, Germany, or the USA).

4. Encrypted storage

- Using Google's encryption at rest mentioned earlier, the document is encrypted and stored on encrypted drives.
- All documents are separated using a unique customer identifier.
- Only a very small number of PaperCut employees have access to documents in storage.
- All access is managed according to Google's IAM¹³ rules on an as-needed basis.
- Any access to data storage, including by operations personnel, generates an audit trail.
- If the PaperCut Cloud Service is not able to process or deliver the document, it will retry for a maximum of 24 hours.
- Once the document has been delivered, or the 24-hour timeout has been reached, it is securely deleted.

5. Delivery to the cloud storage provider

- It's probably only been a minute or two since the user pressed "Scan" on the MFD, and now the PaperCut Cloud Service sends the processed document to the selected cloud storage provider (for example, OneDrive or G Drive)
- The PaperCut Cloud Service always requests the minimum permission required to deliver documents (write-only), however not all cloud providers grant access to this granular level.
- After successful delivery to the cloud storage provider, the user's document is securely deleted from the PaperCut Cloud Service.
- If the delivery of the document is unsuccessful, it can reside in encrypted storage for up to 24 hours after which time it will be securely deleted.

¹³ [Cloud Identity and Access Management](#)



Data retention

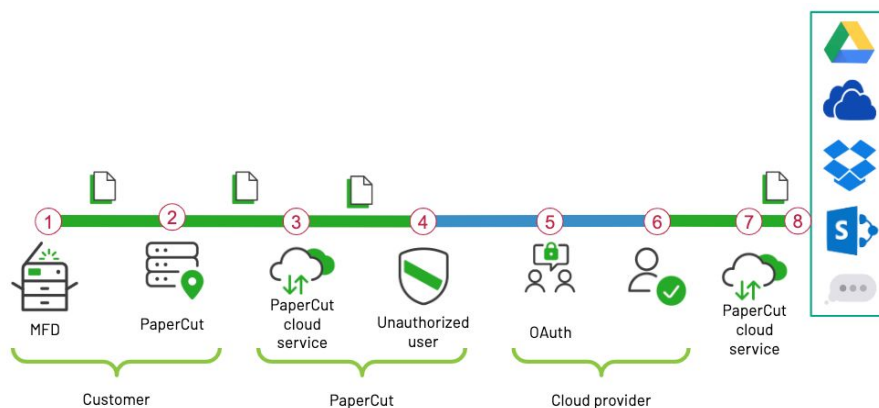
Neither PaperCut MF nor the PaperCut Cloud Service retain the scanned document or any of its contents. However the PaperCut Cloud Service retains the following information about each scanned document:

- ▶ User's email address
- ▶ User's full name or username
- ▶ User's locale
- ▶ Scanned document filename

All data retention is subject to governing laws and regulations providing for redaction (for example, GDPR).

User authorization

The first time a user scans to a cloud storage destination, there is a one-off step where they need to authorize PaperCut MF to upload documents to their chosen cloud storage provider. This authorization can be revoked by the user at any time.



The user authorization path in PaperCut MF Scan to Cloud Storage

1. Actions at the MFD

The actions at the MFD for user authorization are the same as the Document path mentioned above.

- a. The user authenticates to PaperCut MF.
- b. User selects an available Scan Action with a cloud destination.
- c. User scans their document.
- d. The MFD creates the scanned image.
- e. The MFD sends the image document to the PaperCut MF server.

Note: Weak ciphers can be disabled in PaperCut MF so that an MFD is required to use more secure ciphers.



2. The PaperCut MF Application Server

The same actions as those listed for the PaperCut MF application server in the Document path.

- a. PaperCut MF receives the image document and metadata from the MFD.
- b. PaperCut MF creates a secure connection to the PaperCut Cloud Service. PaperCut MF uses only TLS 1.2 Perfect Forward Secrecy compliant cipher suites.
- c. The PaperCut MF server is authenticated and authorized to the Cloud Service. This is done with a unique combination of the PaperCut MF installation id and its license key.
- d. PaperCut MF sends the image document and related metadata (for example, the user's email address) to the PaperCut Cloud Service.

Note: Important considerations such as: securing the PaperCut MF application server and file system, IT staff training, firewalls, etc. are the responsibility of the organization hosting PaperCut MF.

3. PaperCut Cloud Service

The same actions as those listed in the Document path.

- a. The PaperCut Cloud Service receives the document and metadata uploaded by the PaperCut MF application server. This is where all the magic happens.
- b. Depending on the options selected by the PaperCut MF administrator, the uploaded document is opened for processing. This includes OCR, despeckle and deskew.
- c. The PaperCut Cloud Service writes the document to encrypted storage.
- d. The PaperCut Cloud Service is secured by and hosted on Google's Cloud Platform.
- e. The PaperCut Cloud Service runs entirely in the region selected by the customer (Australia, Germany, or the USA).

4. Unauthorized user

- The PaperCut Cloud Service detects if the user is authorized for the selected cloud storage provider.
- If the user is not authorized, PaperCut Cloud Service sends an email to the user containing a secure link (TLS) to the cloud storage provider's authorization page.
- The PaperCut Cloud Service holds the user's document for up to 24 hours, during which time the user needs to complete the authorization.

5. OAuth

- Authorization is established between the user and the cloud storage provider using OAuth2¹⁴. All OAuth communication (including the initial email link) is always TLS 1.2.

¹⁴ <https://auth0.com/docs/protocols/oauth2>



- If the user authorization is successful, the cloud storage provider sends a user-specific token to the PaperCut Cloud Service on behalf of the user.
- Subsequent scans to the same cloud provider do not require re-authorization.
- If the authorization fails or is not completed within 24 hours, the PaperCut Cloud Service securely deletes¹⁵ the user's document.

6. Authorized user

- PaperCut Cloud Service stores the user's OAuth token (encrypted AES256-GCM), and never has access to the user's credentials.
- The token can only be used to deliver documents for a specific user to a specific cloud provider.
- The token never leaves the PaperCut Cloud Service.
- The token has an expiry date.
- PaperCut Cloud Service stores the token encryption keys using Google's Cloud Key Management Service.
- The keys never leave the PaperCut Cloud Service.
- A user can revoke the token at any time by visiting the cloud storage provider's site and removing authorization for PaperCut Cloud Service.

7. Good to go

- a. PaperCut Cloud Service sends the user's document to the cloud storage provider using the most secure TLS available.
- b. The user's document is securely deleted from the PaperCut Cloud Service.
 - If the delivery of the document to the cloud storage provider is unsuccessful, the document can reside in the PaperCut Cloud Service for up to 48 hours (24 hours for user authorization, and 24 hours for processing and delivery).
 - If retries to deliver the document fail, it will be securely deleted.

8. Home sweet home

It's probably only been a minute or two since the user scanned the document, and now it's safely home in the cloud storage provider.

¹⁵ <https://cloud.google.com/security/deletion/>



ATFAQs (answers to frequently asked questions)

A normal part of evaluating PaperCut MF and the Scan to Cloud Storage feature is asking lots of great questions about security. Here are answers to some of the most asked.

Business process

- The PaperCut Cloud Service (PCS) is HIPAA compliant. For cloud-based services, we have entered into a Business Associate Agreement (BAA) with Google that ensures any Personally Identifiable Information (PII) that transits across the cloud instance will be treated appropriately under the Health Insurance Portability and Accountability Act (HIPAA). If you believe you'll be using PCS with PII, consider engaging with PaperCut under a HIPAA BAA, which (like PaperCut's agreement with Google) puts some responsibility on PaperCut when it comes to handling any PII that passes through PCS.
- PaperCut MF (PMF) and PCS are General Data Protection Regulation (GDPR) compliant. In accordance with GDPR you can configure PMF and PCS to redact user details.
- PaperCut assigns an owner to be responsible for implementation and review of Business Continuity and Disaster Recovery (BC/DR) plans.
- All components of the BC/DR plans are reviewed at least annually and updated as needed.
- The BC/DR plan has been tested in the last year.

Software development

- PaperCut conducts ongoing training to validate employee understanding roles and responsibilities during a crisis.
- PMF and PCS follow a documented software development life cycle (SDLC).
- PMF and PCS are developed using secure coding techniques.
- Information security principles are designed into the PMF and PCS life cycle.
- PMF and PCS source code undergo static analysis and static application security testing prior to release.
- PMF and PCS are subjected to automated software testing prior to release.

Vulnerability awareness

- PMF and PCS are externally scanned for vulnerabilities on an ad-hoc basis.
- PMF and PCS participate in bug bounty programs.
- PaperCut supports direct submission of security issues¹⁶.
- PMF undergoes automated scans using OWASP Dependency-Check¹⁷.
- PaperCut reviews all security issue submissions and assigns a risk assessment based on DREAD¹⁸.
- PaperCut supports CVE and publishes to NVD¹⁹ for issue reporting.

¹⁶ <https://www.papercut.com/solutions/security/report/>

¹⁷ <https://owasp.org/www-project-dependency-check/>

¹⁸ <https://wiki.openstack.org/wiki/Security/OSSA-Metrics>

¹⁹ <https://nvd.nist.gov/>



Process and procedure

- PaperCut has a formal incident response plan.
- PaperCut will completely comply with applicable breach notification laws.
- PaperCut requires background checks on all data operations employees prior to employment.
- PaperCut requires new employees to acknowledge data operations policies.
- PaperCut has a documented data security policy.
- PaperCut regularly reviews and updates data security and access policies according to IAM.
- PaperCut has internal audit procedures and full audit trails²⁰ for all cloud data access.

Conclusion

PaperCut MF's Scan to Cloud Storage is built with security as the foundation of every step your document takes. At PaperCut, one of our core values is Customer First, and a significant component of this is protecting what is important to them. We have taken great care (okay, and we kinda geek out) to implement the most advanced security required to keep customer information private, secure, and accessible.

²⁰ <https://cloud.google.com/logging/docs/audit/>



Glossary

Term	Description
Amazon Web Services (AWS)	A suite of cloud computing services offered by Amazon.com.
Business Continuity / Disaster Recovery (BC/DR)	A set of best practices and procedures to ensure a business can recover from and continue to operate after a crisis.
Cloud Security Alliance Security Trust Assurance and Risk (CSA STAR)	Emphasizes key principles of transparency, rigorous auditing, and harmonization of standards.
Common Vulnerabilities and Exposures (CVE)	A standardized list of publicly known cybersecurity vulnerabilities
Data at Rest	Data stored on a device or backup medium in any form.
Data in Transit	Data that is moving through networks (hard-wired, WiFi, mobile, etc.)
Denial of Service attack (DoS)	A cyber-attack where a bad actor attempts to overwhelm a computer or network causing it to be unresponsive.
Family Educational Rights and Privacy Act (FERPA)	A USA government law protecting privacy of student education records.
General Data Protection Regulation (GDPR)	Regulation to harmonize data privacy laws across Europe.
Health Insurance Portability and Accountability Act (HIPAA)	A USA government law stipulating how PII should be protected from fraud and theft by the healthcare and healthcare insurance industries.
Health Information Trust Alliance Common Security Framework (HITRUST CSF)	A privacy and security controls framework.
Information Security Registered Assessors Program (IRAP)	A framework for assessing an organization's security controls for Australian government.
International Organization for Standardization and the International Electrotechnical Commission (ISO/IEC)	Organizations that develop, maintain and promote standards for IT



Term	Description
Man in the Middle attack (MitM)	Where an attacker secretly relays and possibly alters the communications between two parties.
Multifunction Device (MFD)	A printer that also makes copies, scans and donuts.
National Health Service (NHS)	Healthcare providers in the United Kingdom.
National Institute of Standards and Technology (NIST)	A USA government agency promoting standards and best practices for US cybersecurity.
Open Authorization (OAuth)	An open standard for access delegation for Internet users to grant applications access to their information.
Optical Character Recognition (OCR)	The electronic conversion of images of typed, handwritten or printed text into computer readable text.
Open Web Application Security Project (OWASP)	A non-profit organization working to improve software security.
Protected Health Information (PHI)	Any information about health status, provision of health care or payment for health care that can be linked to a specific individual.
Personally Identifiable Information (PII)	Any information relating to an identifiable person.
Amazon Simple Storage Service (S3 Bucket)	An object storage service hosted by AWS.
System and Organization Controls (SOC)	A suite of service offerings CPAs may provide in connection with system-level controls of a service organization or entity-level controls of other organizations.
Transport Layer Security (TLS)	A network layer protocol providing authentication, confidentiality, and data integrity between computer applications.



Thank you

Authors

David O'Hara (Solutions Architect, PaperCut Software)

Contributors

Amir Khassaia (Senior Product Engineer, PaperCut Software)

Bryce Smith (Product Manager, PaperCut Software)

Geoff Smith (Distinguished Product Engineer, PaperCut Software)

Sonja McShane (Product Designer, PaperCut Software)

PaperCut HQ

sales@papercut.com

papercut.com

Support

papercut.com/support