PaperStream AI: Security & Architecture Overview

A Comprehensive White Paper

Disclaimer: This document describes the technical architecture, security posture, and operational practices of PaperStream AI, a paid Software as a Service (SaaS) application extension and PaperStream Capture Pro/ Pro Premium, a desktop capture application, both offered by PFU America, Inc. It was compiled with the utmost care and to the best of our knowledge at the time of publication. The information provided is subject to change as technology evolves and the platform is enhanced. It is recommended to regularly refer to official documentation and consult with our technical and security teams for the most current information regarding PaperStream AI and PaperStream Capture Pro/ Pro Premium. This document is provided for informational purposes only and does not constitute a contractual commitment, warranty, service level agreement, or guarantee regarding the performance, availability, security, compliance, or functionality of PaperStream AI.

PaperStream

Contents

Objectives

PaperStream AI was designed to address the challenges organizations face when turning semi-structured or unstructured business content into actionable data. This white paper serves the following objectives:

- To provide a thorough understanding of PaperStream AI's technical architecture, the design principles behind it, and the rationale for the choices made.
- To describe the security properties, controls, and practices implemented to protect customer data, ensure platform integrity, and deliver a robust, enterprise-grade environment.
- To describe the operational practices that surround PaperStream AI — including monitoring, vulnerability management, and disaster recovery.
- To give security-conscious buyers, procurement teams, and information security reviewers the information to evaluate PaperStream AI against their internal requirements, risk frameworks, and regulatory obligations.
- To equip administrators with the information needed to understand and manage PaperStream AI effectively within their own organizations.

Introduction

In the evolving landscape of intelligent document capture and processing, PaperStream AI emerges as a modern, secure, and scalable solution, purpose-built for a multitude of business environments. PaperStream AI is built on Microsoft Azure, combining the reliability, global reach, and data handling of a world-class cloud infrastructure provider with deep expertise in data management.

PaperStream AI is designed first and foremost as an application extension that supports intelligent document processing for all types of workflows. Combined with PaperStream Capture Pro / Pro Premium, the platform's architecture and operational model are built to help customers meet their demanding requirements for availability, security, and compliance. PaperStream AI is deployed on Microsoft Azure to provide the underlying compute, storage, and network services required to operate at scale; however, the security posture and customer experience described in this paper are driven by PaperStream AI's application design and PFU America's operational practices.

This white paper explores the architecture, security controls, performance characteristics, and administrative capabilities that make PaperStream AI a trusted choice for organizations with demanding secure and robust workflow requirements.

“

A modern, secure, and scalable solution, purpose-built for a multitude of business environments.

User Benefits

PaperStream AI allows organizations to focus on their core business problems while simultaneously transforming their data. There are several benefits that organizations derive through PaperStream AI's use.

Accessibility

PaperStream AI is accessible from any device with an internet connection and an installation of PaperStream Capture Pro / Pro Premium (with version 6.0 or above). It is designed so that administrators and end users can access the platform from the office, at home, or in the field. The platform strives to maintain consistent uptime, enabling access to critical workflows. For organizations with hybrid or remote workforces, this accessibility supports reduced location-based limitations associated with traditional systems.

Scalability

PaperStream AI is architected to scale seamlessly with organizational growth. Whether onboarding new users, expanding to additional locations, processing increasing volumes of data, or integrating with new downstream systems, PaperStream AI accommodates growth without requiring infrastructure changes or operational interruptions. New PaperStream Capture Pro / Pro Premium instances can be connected quickly, supporting faster deployment of distributed environment.

Automatic Updates

PaperStream AI is delivered as a managed SaaS solution. All updates including security patches, bug fixes, performance improvements, and new features are applied automatically by PFU Limited. Customers can access the latest available capabilities and updates, helping to reduce the administrative effort required for maintenance.

Cost-Efficiency

PaperStream AI's subscription-based model is designed to reduce the need for additional upfront capital investment in hardware or on-premises infrastructure. Costs are recognized as operating expenditures, simplifying budget management and improving financial predictability. Organizations can adjust their subscription tier on a yearly basis and add additional credits as their needs evolve, ensuring cost alignment with actual usage. The platform's cloud-native architecture can help reduce the certain costs associated with on-premises systems — including hardware maintenance, refresh cycles, and some of the IT resources required to support dedicated infrastructure.

Architecture Overview

The application platform is composed of two parts. PaperStream AI is one, PaperStream Capture Pro/ Pro Premium the other. Knowing how these parts are structured, and how they are kept separate from one another allows administrators and security reviewers to understand where the platform's protections sit before tracing how a document actually moves through it. The PaperStream AI tenant is hosted on Microsoft Azure, which provides the compute, storage, networking, and identity services the platform depends on.

Application Design

PaperStream Capture Pro/ Pro Premium is the core application: it handles capture, indexing, extraction, and routing. PaperStream AI is built as a set of independent containerized components rather than a single monolithic service to extend PaperStream Capture Pro / Pro Premium's capabilities.

Where Data Is Stored and Protected

Two layers hold data inside the environment. The database stores tenant data, user information and licensing permissions. A separate file store holds documents and intermediate processing artifacts only for the duration of a job. Data written to either layer is encrypted at rest and the database is backed up incrementally on a regular schedule.

Identity and Access

Access to the platform is centralized. Users authenticate through OAuth 2.0, with administrators having multi-factor authentication available to enforce an even stronger security posture. Furthermore, what each user can do is governed by role-based access control. Administrative actions and changes to access are recorded in audit logs.

Data Flow

The components described in the architecture overview work together to deliver actionable data to business processes. Understanding that sequence helps administrators and security reviewers trace how a document moves through the system, identify where data resides at each stage, and confirm that the platform's behavior aligns with their organization's requirements.

A document's journey through PaperStream AI begins with ingestion. A user initiates a capture job in PaperStream Capture Pro/ Pro Premium, which accepts documents from connected scanners, digital file imports, email attachments, and shared folders. The ingestion service authenticates the request and validates the input before accepting any data. All communication at this entry point is encrypted over TLS 1.2 or higher.

Once accepted, the document is passed to the processing engine and queued as an asynchronous job. The engine performs optical character recognition (OCR), classifies the document type, and applies AI-driven field extraction. These operations draw on cloud-based AI components. While processing is underway, intermediate outputs are written temporarily to tenant storage.

After processing is completed, the extracted data is returned to PaperStream Capture Pro/ Pro Premium. The destination connectors then route that data to the appropriate downstream system namely, an ERP platform, ECM repository, or other connected business applications. Each connector enforces the authentication and transmission requirements of the destination system.

Once delivery is confirmed, the document data held in transient storage is automatically deleted. Audit records of the transaction are retained in the database layer as part of the platform's operational log.

Security Overview

PaperStream AI's security posture is built on the principle of defense in depth. Multiple independent layers of control work together to protect customer data. This section describes each layer, specifically how network communication is secured, how data is encrypted, how users are authenticated, how tenant boundaries are enforced, and how data is handled and retained throughout its lifecycle.

Ports, Connectivity, and Transport Layer Security

PaperStream AI uses HTTPS exclusively for all communication between clients and the platform. Port 443 is the standard entry point for all administrative and API traffic. TLS 1.0 and TLS 1.1 are disabled at the platform level and cannot be negotiated by clients, as they are antiquated protocols. The minimum accepted version is TLS 1.2.

For document delivery to downstream systems, the specific ports used depend on the destination connector's configuration. Each connector specifies its required outbound connections as part of its setup. Administrators should review those requirements with their network or security team before deploying a connector to ensure that any necessary firewall or proxy rules are in place.

Encryption

PaperStream AI enforces AES-256 encryption on object storage, the tenant database, database backups, and virtual machine disks. This standard applies to content regardless of its sensitivity classification.

All data in transit between clients and the platform is protected by TLS 1.2, at a minimum. Together, these controls ensure that customer data is encrypted both while it moves across the network and while it rests in storage.

User Authentication

Access to the PaperStream AI Dashboard is managed through OAuth 2.0-based authentication. Multi-factor authentication (MFA) is supported at the platform level and should be enabled by administrators to require a second form of verification beyond username and password.

Administrators manage role assignments through the PaperStream AI administrative interface. User permissions are governed through role-based access control (RBAC), carrying the minimum permissions required for its function. Users are unable to access capabilities or data outside the scope of their assigned role.

Administrative actions are recorded and are available to authorized administrators to help support compliance and security review requirements.

Tenant Isolation

PaperStream AI is a multi-tenant platform. Each customer organization is assigned a unique Tenant ID. Data, configuration, and processing activity is scoped to that specific identifier. Tenant isolation controls are designed to help prevent unauthorized access between tenants.

Tenant separation is enforced at the application layer through the platform's access control architecture. Data boundaries between tenants are maintained through strict, logical isolation.

Data Handling and Retention

Document data in PaperStream AI is ephemeral by design. Files ingested for processing are held in storage for the duration of the processing job.

Extracted field data and metadata are returned to the customer's environment through PaperStream Capture Pro/ Pro Premium. Once delivered, that data resides in the customer's own systems and is subject to the customer's own data governance policies. It is then disseminated via the configured destination connectors to the appropriate downstream workflow.

Tenant configuration, user records, and licensing data are stored in a database and are retained for the life of the customer's subscription. Data is protected by AES-256 encryption and backed up incrementally on a regular schedule.

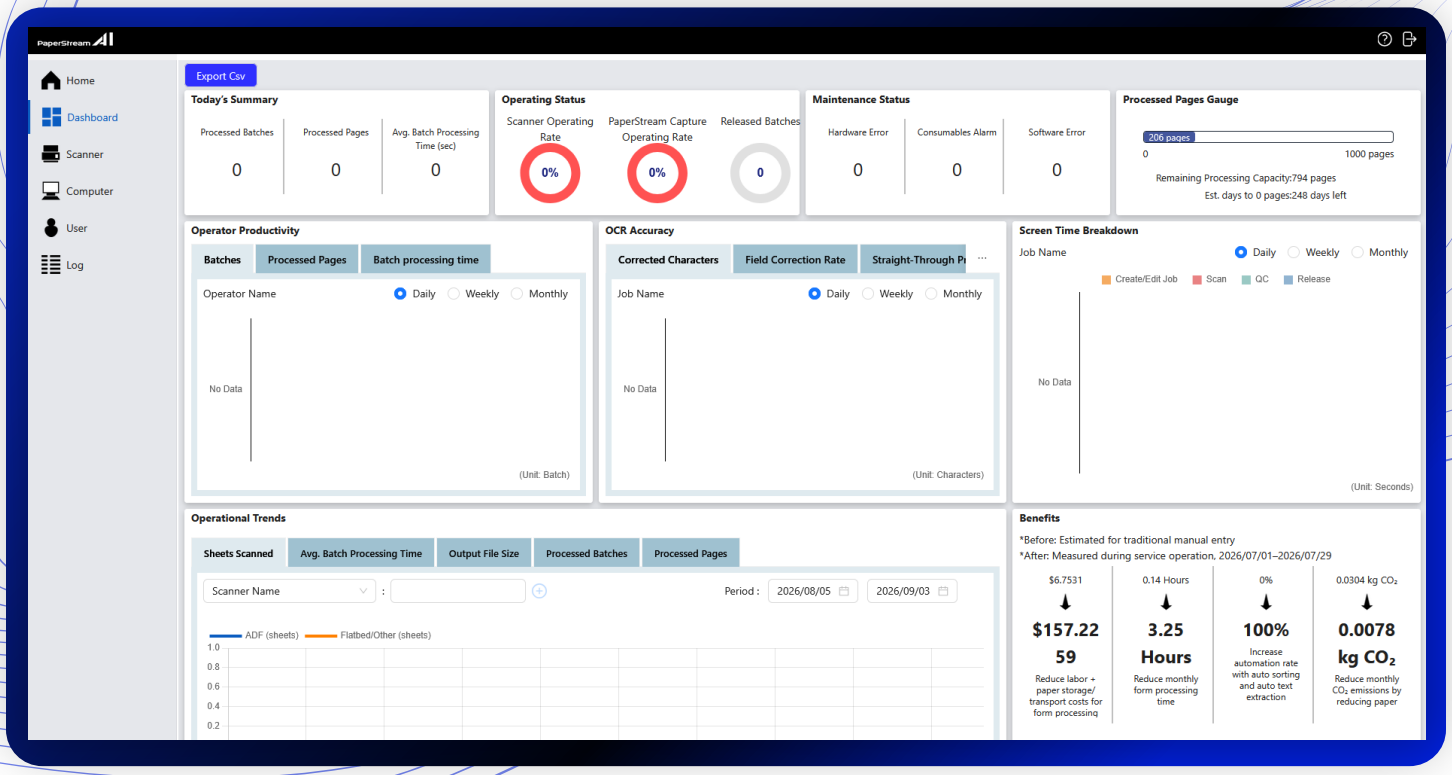

Performance

PaperStream AI is designed to maintain consistent processing throughput across a wide range of workloads. The platform's cloud-native architecture provides the underlying compute capacity and redundancy to meet demand at scale. Performance is driven by two primary design choices: automatic load distribution and dynamic resource scaling, both of which operate without manual intervention from administrators or customers.

Load Balancing

Processing workloads in PaperStream AI are distributed automatically across available compute resources. No single server or component carries the full processing load; instead, incoming jobs are distributed to balance utilization across the system. This prevents bottlenecks during periods of elevated demand. Load balancing operates transparently. Administrators do not configure or manage it directly; it is handled by the platform's infrastructure layer.

Dynamic Scaling

PaperStream AI scales its compute resources dynamically in response to actual workload. When processing volume increases, the platform allocates additional capacity. When demand subsides, those resources are released. Administrators are not required to provision or pre-size infrastructure in anticipation of peak loads.

Centralized Management

PaperStream AI consolidates platform administration into a single interface, which can reduce the operational overhead that typically comes with managing document processing across multiple locations or user groups. Administrators can monitor processing activity, track usage, and manage users from one place, without requiring direct access to the underlying infrastructure.

Rapid Deployment

New PaperStream Capture Pro/ Pro Premium instances can be connected to PaperStream AI quickly. The Document AI Quick Start feature guides administrators through initial configuration with minimal upfront effort, enabling the platform to begin processing documents without requiring custom template development. PaperStream AI includes built-in support for common document types including invoices, purchase orders, quotations, delivery notes, and receipts. Organizations working with these forms can begin processing almost immediately after setup.

Each individual PaperStream Capture Pro/ Pro Premium installation connects to the same PaperStream AI environment, following a many-to-one architecture. For organizations expanding to new sites or onboarding additional users, the process is consistent and does not require separate infrastructure provisioning for each new instance, only for the administrator to add additional PaperStream Capture Pro/ Pro Premium licenses via the PaperStream AI operational dashboard.

IT Infrastructure Reduction

PaperStream AI is a fully managed service. PFU Limited manages the OCR engine, AI processing components, and related service operations. Customers are not required to provision or maintain dedicated servers, storage, AI or OCR processing software within their own infrastructure.

Platform updates, including security patches, feature additions, and performance improvements, are deployed on an as-needed basis.

Administrative Interface

PaperStream AI includes an operational dashboard that gives administrators a consolidated view of platform activity. The dashboard displays processing volumes, credit consumption, accuracy metrics, and processing status across connected workflows.

From the dashboard, administrators can review processing history, monitor usage against their subscription tier, and access user configuration settings. The interface is designed for use by both IT administrators and operations managers; it does not require specialized technical knowledge to navigate.

Information for Administrators

PFU Limited maintains a system status page that provides real-time information on PaperStream AI availability. PFU America also maintains a support resources page that includes user and administrator guides for PaperStream Capture Pro / Pro Premium and PaperStream AI, trouble-shooting guidance, FAQs, and notices regarding planned maintenance and upcoming changes. Administrators are encouraged to bookmark these pages and consult them as a first step when investigating unexpected platform behavior or potential service disruptions.

System Status

Support Contact

For immediate support concerns, please contact our technical assistance center at (800) 626-4686 or email techsupport@pfu-us.ricoh.com.

Support Resources

PFU America provides technical support resources for PaperStream AI administrators.

Learn More Here

About the Author

Paul Gramolini, SecAI+, is a Product Manager for PFU America, Inc. He bridges technology, product strategy, and customer needs to deliver innovative solutions that improve productivity and transform information management.

Conclusion

PaperStream AI addresses a practical and persistent problem. Organizations accumulate large volumes of semi-structured and unstructured business content. Converting that content into usable data has historically required significant manual effort, dedicated infrastructure, or purpose-built extraction platforms. PaperStream AI changes that equation by embedding AI-driven classification and data extraction directly into existing PaperStream Capture Pro/ Pro Premium workflows — without requiring new infrastructure, replacing familiar tools, or undertaking a large-scale technology transition.