Nightfall

Guide to Data Loss Prevention (DLP) in Confluence

Data security in Confluence requires a flexible approach

Confluence is one of Atlassian's most popular collaboration tools. It's a team workspace where users can coordinate on a variety of project types. Over the least year, the rise of remote work has meant many companies have hosted their internal information hubs on Confluence. As more organizations adopt Confluence, the platform has evolved to have more structure – encouraging users to share even more data than ever before.

Each organization has their own unique definition of what constitutes sensitive data, depending on their compliance needs, past security incidents, and other factors. In some cases, security leaders could determine that certain sensitive data needs to be shared within Confluence (e.g. data sharing is required amongst a particular team). On the other hand, some types of data might always be considered sensitive no matter how it appears (e.g. Social Security Numbers, or names alongside ICD10 codes).

It's essential for security leaders to provide safeguards to prevent confidential information from being accessed inappropriately (e.g. externally), without preventing acceptable data-sharing. Atlassian relies on third-party apps like Nightfall AI to provide data loss prevention (DLP) functionality within their apps like Confluence.

How are secrets exposed in Confluence?

- Typical uses for Confluence include hosting public documentation for projects like bug tracking for open-source projects or release notes. Development teams create this information to host publicly on purpose, but sometimes sensitive data can be inadvertently exposed to anyone who navigates to the correct URL.
- Confluence sites tend to be more frequently open to public view. In a sample of 5,000 Jira Software Cloud sites, there were 273 Jira sites with publicly viewable issues and 1,214 Confluence sites with publicly viewable spaces. The larger number of viewable Confluence sites compared to Jira is likely explained by companies having documentation and release notes on a publicly viewable space.
- Information leaked through public Confluence spaces has included customer portal passwords, customer requests, email conversations, meeting notes, and product order data — all of which can contain data that companies must protect.

What types of data are at risk of exposure in Confluence pages and spaces?

Credentials & secrets are most frequently exposed in Confluence, due to its common use as an internal wiki for product and engineering teams. However, other types of unexpected sensitive data have been identified in Confluence by customers using Nightfall. Here are a few types of data that may be exposed in Confluence:

- API keys & access tokens for third party services, e.g. AWS, Stripe, Twilio, etc.
- Cryptographic keys (SSH, PGP, etc.)
- Certificates (SSL, TLS, etc.)
- Passwords and login credentials
- Database credentials
- UUIDs, cookies, etc.
- Credit card numbers
- Customer PII

The need to identify sensitive data in Confluence often arises from a business change, such as merging or separating business units and the subsequent data cleanup. Security leaders should also be aware of their specific industry requirements for compliance, and ensure data security standards are met within Confluence — for example, HIPAA compliance may require that protected health information (PHI) is secured within Confluence.

Permissions settings in Confluence is only half the battle

Proper permissions settings can help reduce data exfoliation risk in Confluence. However, this is just one step you must take to protect sensitive information when collaborating in this space. Data security best practices for Confluence should include these permissions settings:

Include stakeholders as Confluence Admins
Identify stakeholders who will help manage your Confluence accounts and enforce proper user hygiene. Active admins will be your first line of defense in securing your spaces and enforcing data policies.

Understand how permissions levels and restrictions work together
Confluence provides permissions controls at multiple levels and understanding this is key to making sure that no one is authorized to view, add, modify, export, or delete data in Confluence. Security leaders must ensure that the admins who have access to their Confluence admin console understand how permissions function in Confluence.

Monitor logs to track permissions changes across your spaces

Confluence provides product-specific audit logs that administrators within those respective services can access. Reviewing these logs can provide insight into who created, edited, or deleted a space, as well as who is making changes to groups and user permissions. This step will help in the moderating of Confluence spaces and the Atlassian organizations they're part of.

Properly onboard and offboard members

Confluence allows Admins to securely onboard new team members by only inviting users with a designated domain name in their email account. Organization admins who manage org-wide permissions should also be sure to remove users who have left the company.

Key takeaway

All paid instances of Confluence have three levels of permissions:

- **Global permissions** are broad and site-wide.
- **Space permissions** uniquely apply to the space specified by an administrator (usually the space creator).
- **Page restrictions** allow admins to restrict the view or editing of specified pages by specific groups or users.

What is data loss prevention (DLP)?

Data loss prevention is an access control that ensures confidential information is kept on a need-to-know basis. DLP scans for content within messages and files to determine whether an unauthorized disclosure of business-critical information has occurred and can provide automated remediation on the basis of your established data security policies. Additionally, DLP can provide alerts and analytics that help organizations understand risk and employee behavior over time.

Since Confluence is document and media heavy, it can be hard to detect when and where sensitive info could be leaked within the platform. This can lead to increased data exposure risk as well as data compliance violations.

Organizations need to use tools like DLP in order to put into place controls that will help enforce data security best practices by preventing unauthorized parties from accessing documents and folders with sensitive information.

Why is data loss prevention (DLP) essential for protecting data in Confluence?

There are no mature cloud-native DLP products for Confluence. Atlassian does not have a native DLP product, and many CASBs cannot support Atlassian apps. And while some CASBs purport to connect to Confluence, they can only see data in transit and are unable to find pre-existing sensitive content.

Nightfall is the only Confluence DLP solution on the market that enables customers to scan their entire existing Confluence instance. With Nightfall, you can flexibly configure multiple different DLP policies, and apply them to particular locations (Spaces and Pages) within Confluence. This leads to a comprehensive DLP approach in which appropriate data sharing won't be flagged, reducing false positives and noise.

With Nightfall, users can create multiple detection rules that specify whether data is deemed sensitive in any instance, or whether it is deemed sensitive only in combination with other data. This provides granular control over the organization's unique definition of what constitutes sensitive data, further reducing false positives and noisy alerts.

What is Nightfall?

Nightfall is a platform to discover, classify, and protect sensitive data across cloud SaaS & cloud infrastructure. Nightfall supports compliance efforts with a number of industry standards like PCI-DSS, GDPR, HIPAA, CCPA, and much more. Additionally, Confluence is just one of the many platforms Nightfall secures. You can protect data across all your SaaS apps with our native integrations for Confluence as well as Slack, GitHub, and more — or build completely custom solutions for other systems with the Nightfall Developer Platform.

Nightfall works by continuously monitoring data flowing in and out of data silos and classifying that data with machine learning. Data marked as sensitive can be automatically quarantined, deleted, and redacted with workflows.

Key benefits

- Leverage 150+ pre-tuned, machine-learning trained detectors out of the box.
- Customizable, configurable, and flexible DLP for all your Confluence Pages and Spaces.
- Deploy a targeted remediation strategy with comprehensive, context-rich scan results that contain direct links to violations in Confluence.

Discover

Continuously monitor data that is flowing into and out of data silos.

Classify

Machine learning classifies sensitive data & PII automatically

Protect

Automated workflows for quarantine, deletion, redaction, alerts, and more.

How does Nightfall differ from existing platforms?

Nightfall DLP is the industry's first cloud-native data loss prevention solution designed to **discover, classify, and protect** sensitive data in cloud environments. Leverage 150+ detectors, including Nightfall's best-in-class core of machine-learning trained detectors, to detect a wide range of sensitive content types such as standard PII, names, ID numbers, financial information, addresses, credentials, secrets, custom regexes and word lists, and more.

Discover sensitive data across all Confluence Spaces (including Personal Spaces), Pages, Blog Posts, Attachments, Comments and Archived items. Secure your Confluence with DLP scans for a wide range of file types — including plaintext, Office (Google Office, Open Office, msft Office), pdf, html, xml, all popular image file types (jpeg, png, etc), compressed files (zip, tar, etc).

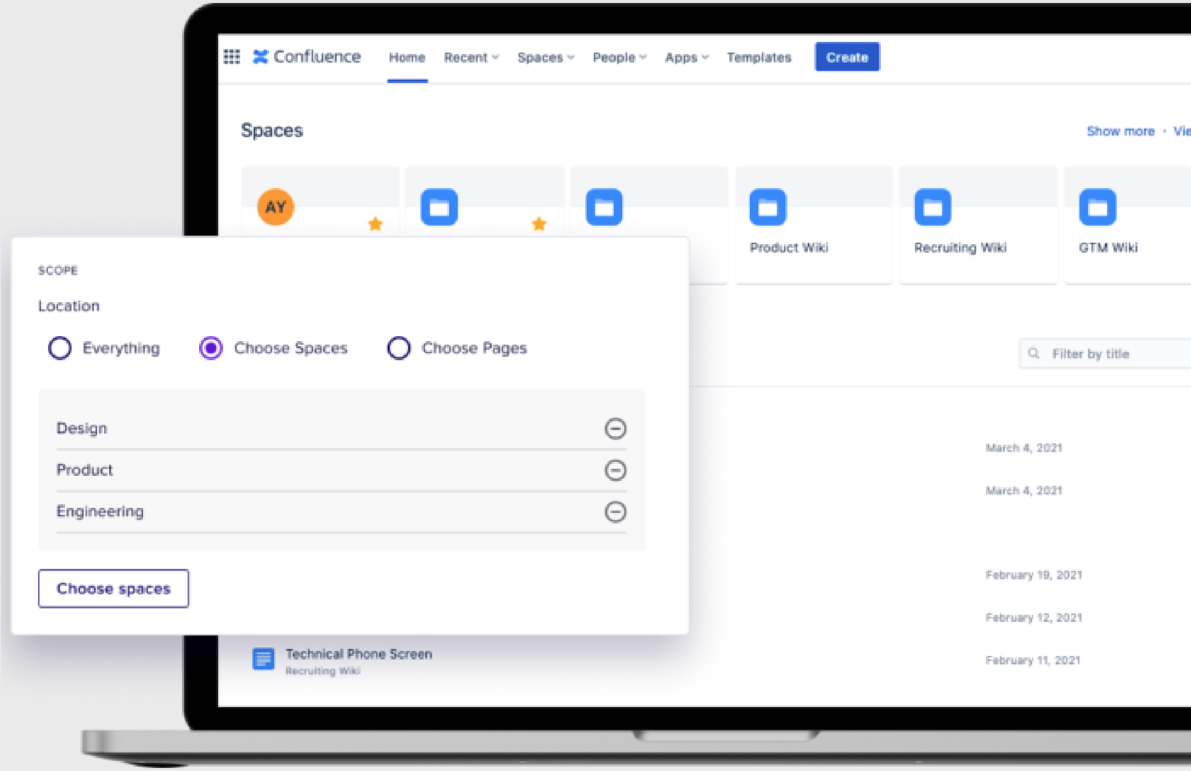

SCOPE

Location

○ Everything ● Choose Spaces ○ Choose Pages

Design

Product

Engineering

March 4, 2021

March 4, 2021

February 19, 2021

February 12, 2021

Technical Phone Screen
Recruiting Wiki

February 11, 2021

Key takeaway

Nightfall DLP is the industry's first deep learning-based platform to detect sensitive data like credentials & secrets in Confluence. We designed our platform to address the low accuracy of tools that rely on traditional methods like regexes or entropy thresholds. Nightfall can be used to discover and protect against both PII and credential leakage across all your Confluence pages and spaces.

What are the key features of Nightfall DLP for Confluence?

Nightfall helps organizations manage DLP in Confluence with these features:

- Quickly and easily connect Nightfall to your Confluence in minutes with our out of the box integration.
- Fully customize your scans with Nightfall's robust detection engine with 150+ detectors, including our proprietary machine-learning trained detectors to detect a wide range of sensitive content types such as standard PII (names, ID numbers, financial information, and addresses) credentials & secrets, custom regexes & word lists, and more.
- Configure granular Detection Rules and set confidence levels within the Nightfall dashboard to determine which data is considered sensitive, either standalone or in combination with other data.
- Build flexible data detection policies based on custom data detectors (e.g. regexes & word lists) and multiple policies to target your DLP scans to certain locations (e.g. Spaces or Pages) or timeframes.
- Discover sensitive data across all Confluence Spaces (including Personal Spaces), Pages, Blog Posts, Attachments, Comments, and Archived items, with context-rich results.
- Nightfall's best-in-class DLP includes machine learning based optical character recognition (OCR) for unstructured data, enterprise-grade security, and high accuracy detection via deep learning — all within a single pane of glass, intuitive UI for configuring your DLP policies.

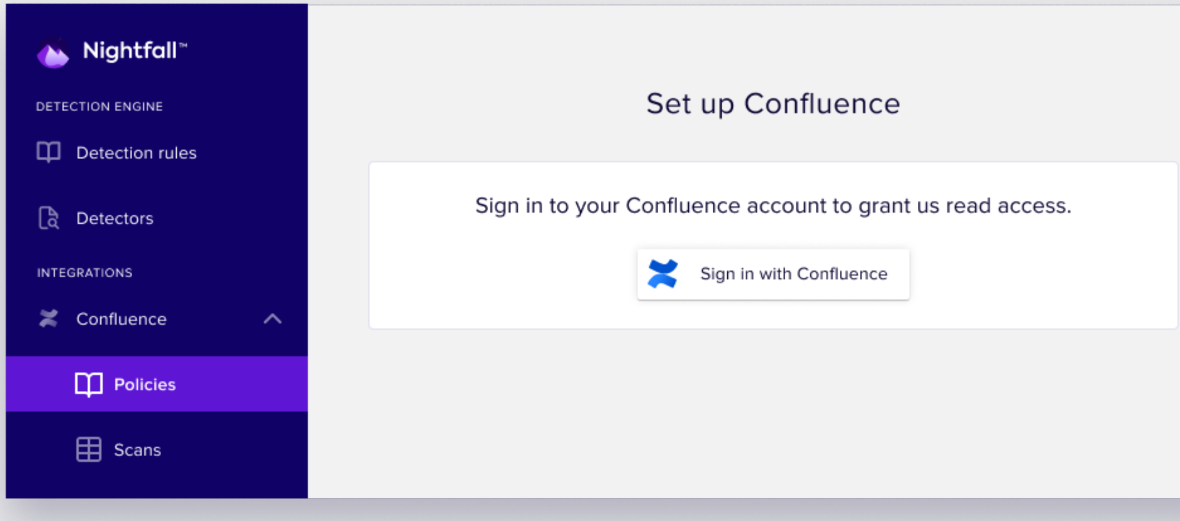

Set up Confluence

Sign in to your Confluence account to grant us read access.

What types of PII does Nightfall detect?

Nightfall can detect the following token types within images via OCR and over 100+ file types, including Google proprietary files types:

Standard PII: Age, Credit Card Number, Email, Ethnic Group, Name, Location, Phone Number

Health: ICD, FDA, DEA, NPI, DOB

Finance: IBAN, SWIFT, CUSIP, Routing Numbers

Crypto: Bitcoin, Ethereum, Litecoin Addresses & Private Keys

Network: IP Address, Hardware ID, MAC Address

Custom: API Keys, your application UUIDs, and much more.

IDs: Driver's License Number, Taxpayer ID, Passport Number, Social Security Number, Vehicle ID

Case study: Amount

Industry: Financial Services
Employees: 660
Don Stewart, Jr., Security Operations Analyst
Profile: Amount develops products and services to help financial institutions grow in the digital realm. Their technology solutions put customer experience first and allow their partners to enter the digital financial space with speed and ease. Amount's requirements for providing for secure banking services must include compliance and data security. Nightfall helps protect Amount's customer data from internal and eternal threats with cloud-native DLP.

The challenges of managing cross-functional data protection

Amount has three goals for protecting customer data: reducing security risks, increasing peace of mind, and securing the company's platform environment. These needs led the information security and I.T. teams at Amount to search for a data security solution that was fast and easy to implement. "We didn't want a heavyweight agent that would bog down our processors," Don says. "We also wanted it to be customizable."

One set of detection rules for all apps

Managing data security and compliance across the entire Amount digital ecosystem is a challenge for Don's teams. Nightfall unites data detection and classification from Slack, Github, and Confluence into one interface for Amount and gives them a deeper understanding of their attack surface and threat risk. With Nightfall, they can apply the same detection rules across all their scans and each of their collaboration tools.

"With Nightfall, our collaboration tools are under one umbrella. It allows us to see everything in a single pane and without having to configure rules in different places."

AMOUNT

Want to learn more about Nightfall?

To get started with Nightfall, request a demo or email us at **sales@nightfall.ai** with any questions.

Request a demo

About Nightfall

Nightfall is the industry's first cloud-native DLP platform that discovers, classifies, and protects data via machine learning. Nightfall is designed to work with popular SaaS applications like **Slack**, **Google Drive**, & **GitHub** as well as IaaS platforms like **AWS**. Developers can build custom solutions in any Saas app with the **Nightfall Developer Platform**.

Nightfall