TECHNICAL WHITEPAPER · FEBRUARY 2023

Protecting Sensitive Data with Policy-Based Selective Encryption

Access Control Carrying Content — a new paradigm for securing unstructured dataacrossmulti-cloudinfrastructure,poweredby standardized cryptographythattravelswiththefileitself.

AUTHOR

Karim Eldefrawy, Ph.D.
Co-founder and CTO, Confidencial, Inc.

FEB 2023
confidencial.io

THE PROBLEM

Challenges Facing Secure Collaboration

Today, it is still challenging to easily and securely collaborate and to selectively shares ensitive unstructured content on a need-to-know basis — content suchas PDF, Word, Excel, and PowerPointdocuments,emails and messages, and images. Current secure sharing solutions fall under one of four classes:

1. Password-locked documents, with documents and passwords sent via email.
2. Cloud-stored documents, with access control at the folder/file level.
3. Versioned documents, with each version containing only a subset of the original content.
4. Contractual statements or compliance policies to mandate that information be kept secure.

These approaches do not provide a long-term, user-friendly, secure solution because once documents are decrypted or downloaded, no protection is provided. In fact, these approaches have not been effective at eliminating leaks and breaches. Executives are looking for better ways to secure unstructured documents across their multi-cloud infrastructure[1,2]:

97% of executives say insider breach risk is a significant concern.

78% of executives think employees accidentally put data at risk in the past year.

75% of executives think employees have put their data at risk intentionally.

90% of enterprises are multi-cloud or have plans to move towards such a mode.

Confidencial's platform is comprised of a set of technologies that automatically discover and protect sensitive information. Our technology enables users to protect parts of unstructured content easily and selectively inside authoring tools (e.g., Word, Excel, PowerPoint, and PDF documents, Outlook emails) as well as within workstream collaboration applications (e.g., Slack and MS Teams). Protection of sensitive information is achieved by using policy-based encryption throughout the life cycle of information. Our use of standardized encryption schemes is modular, and easily customizable and upgradable to the new post-quantum encryption standards that expected in the coming years.

THE PARADIGM

A New Protection Paradigm — AC3

97%

97% of IT Leaders Say Insider Data Breaches are a Major Concern | 2020-02-19 | Security Magazine

90%

90% Of Companies Have A Multicloud Destiny: Can Conventional Analytics Keep Up? (forbes.com)

Our technology enables a new protection paradigm via Access Control Carrying Content (AC3). AC3 travels with content, protects it across multiple clouds, and maximizes collaboration while minimizing the effects of leaks and breaches. AC3 applies proactive leak management via fine-grained cryptographic access control that is inseparable from the content because different parts of the content/data itself are encrypted for different users, groups, or roles. Our protection works inside and outside the enterprise to seamlessly enable inter- and intra-enterprise collaboration with protected content. Our AC3 approach not only encrypts files with a standard encryption scheme, but it also uses multiple standardized encryption schemes (e.g.,RSA-OAEP in the PKCS#1 standard, and the AES-256-GCM standard by NIST) to realize multi-receiver, policy-based protection tailored for specific elements and objects in unstructured content.

See Figures 1 through 4 for an illustration.

- Figure 1: Illustrates the supported generic encryption policies.
- Figure 2: Compares our approach to Attribute Based Encryption (ABE) which our original academic work[3] covered in detail.
- Figure 3: Illustrates how encryption of policy AND clauses occurs, and
- Figure 4: Illustrates how encryption of policy OR clauses occurs.

The original of the technology underlying AC3 was published in the proceedings of the proceedings of the 2022 6th International Symposium on Cyber Security, Cryptology, and Machine Learning [CSCML'22.]

[CSCML'22] Karim Eldefrawy, Tancrède Lepoint, and Laura Tam, "In-app Cryptographically Enforced Selective Access Control for Microsoft Office and Similar Platforms", in proceedings of the 6th International Symposium on Cyber Security, Cryptology, and Machine Learning (CSCML), 2022.

Conference version available at: In-App Cryptographically-Enforced Selective Access Control for Microsoft Office and Similar Platforms | Cyber Security, Cryptology, and Machine Learning

Extended version available at cscml2022-abe-msoffice-extended-version.pdf (keldefrawy.github.io)

ENCRYPTION POLICIES

Supported Encryption Policies

Figure 1: Example of Supported Encryption Policies

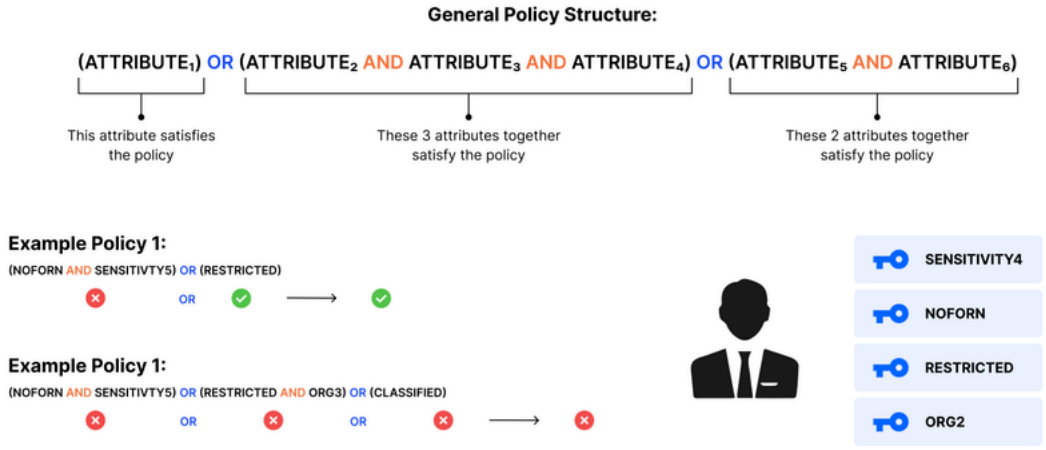

Figure 2: Emulating ABE via Standardized (Public-Key) Encryption

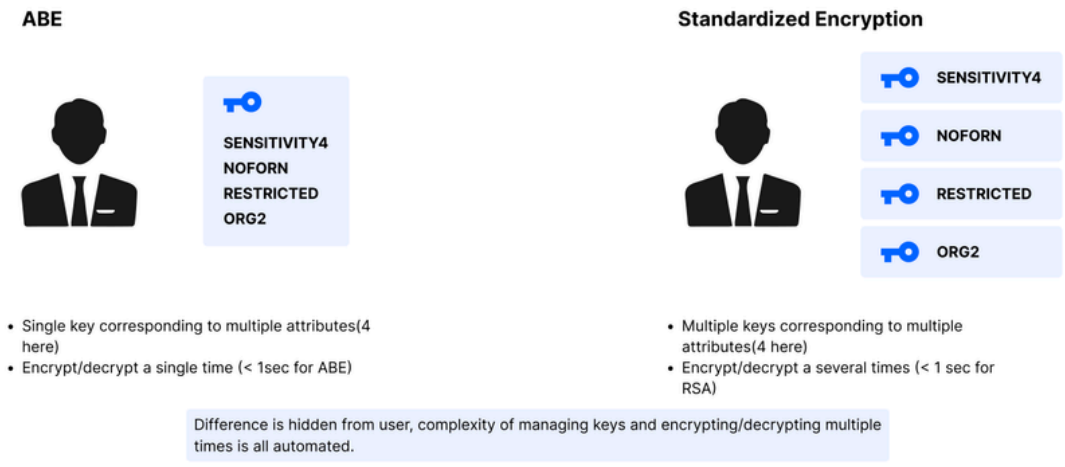

ENCRYPTION POLICIES

Figure 4: Encryption of AND of attributes. Encryptions of shares of the AES key are embedded in documents as meta-data.

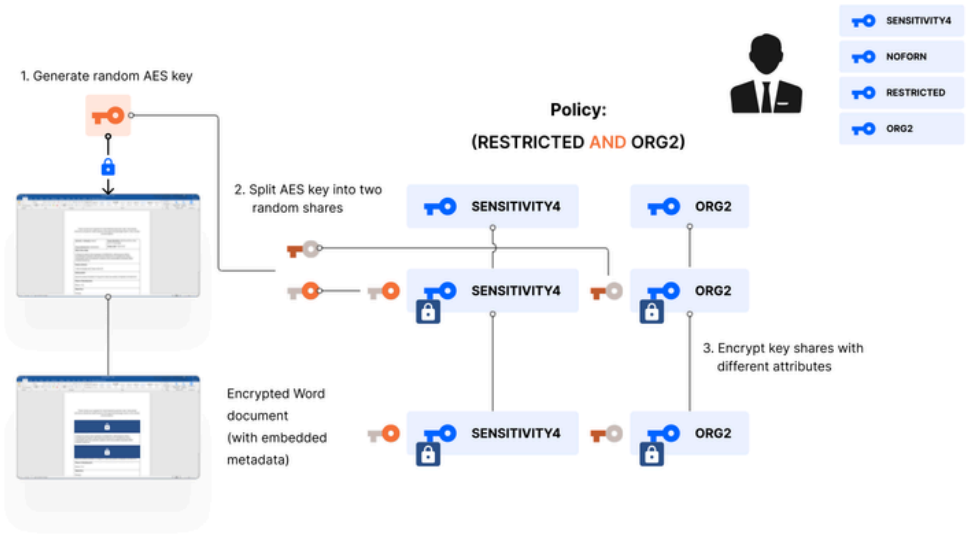

Figure 3: Encryption of OR of attributes. The encryptions of the AES key are embedded in the document as meta-data.

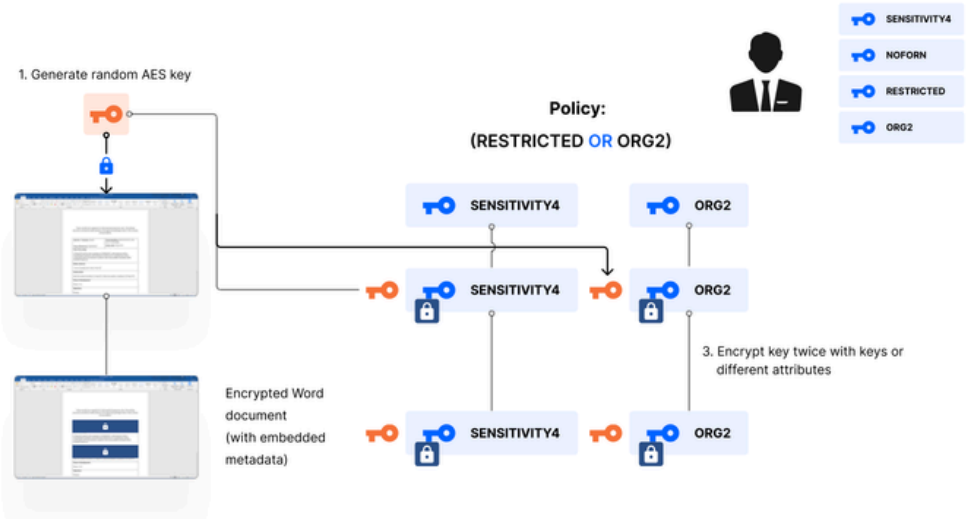

IN PRACTICE

Use Cases

Our platform and the underlying technology are transport- and storage-agnostic and can be applied to a variety of settings (see Figure 5 forillustration ofsteps of operation when securely sharing a document). We are currently focused on the two classes of use cases below, which are already deployed in production at tens of enterprises:

Secure Document Sharing

Users securely share documentsthat are generated from widely used desktop applications and authoring tools (e.g., MS Word, Excel, PowerPoint, and PDF), as well as workstream collaboration tools (e.g., MS Outlook, Slack, and MS Teams).

Secure Document Request

Users securely requestand obtaindocumentsin popular formats (e.g., PDF, MS Word, Excel, PPT, and images such as PNG and JPEG/JPG). The sender of the documents does not have to install any software, nor be a Confidencial user.

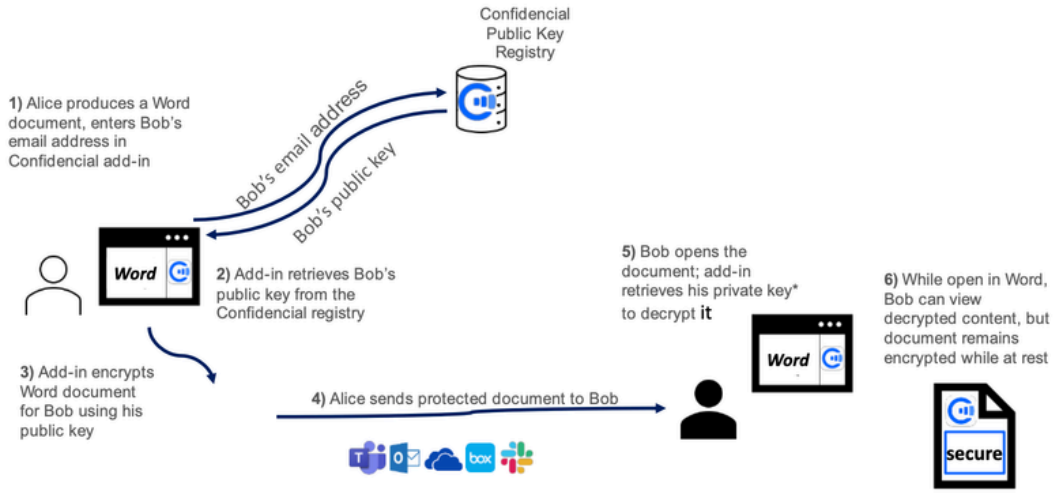

Figure 5: Illustration of steps involved when securely sharing a document using Confidencial's protection with other existing email, messaging, and/or cloud-storage systems. We note that all cryptographic operations happen on a user's device.

INFRASTRUCTURE

Deployment and Architecture

The Confidencial platform architecture has been designed to be modular and flexible and is able to accommodate the following deployment modes:(1)SaaS/cloud-based deployment,(2)on-prem deployment, (3)hybrid deployment.Figures 6 and 7 show two of the most common deployment modes that enterprises request from us. Figure 6 shows a hybrid deployment of Confidencial's platform with an on-prem server for decryption keys and an on-prem server for document events and audit-trail logs. Figure 7 illustrates a cloud-based deployment of Confidencial's platform with keys split (secret shared) between Confidencial's and the enterprise's identity provider tenant. Note that Confidencial does not see encrypted content nor decryption keys; the shares of private keys stored at Confidencial are mathematically proven to be random strings that do not enable decryption nor leak information about the private keys themselves.

We also highlight that our platform could leverage existing enterprise access control and credential management systems such as (Azure) Active Directory, cloud-based solutions like OKTA, or any other SAML- or OpenID-compliant system. The Confidencial platform thus seamlessly integrates into existing enterprise IAM infrastructure.

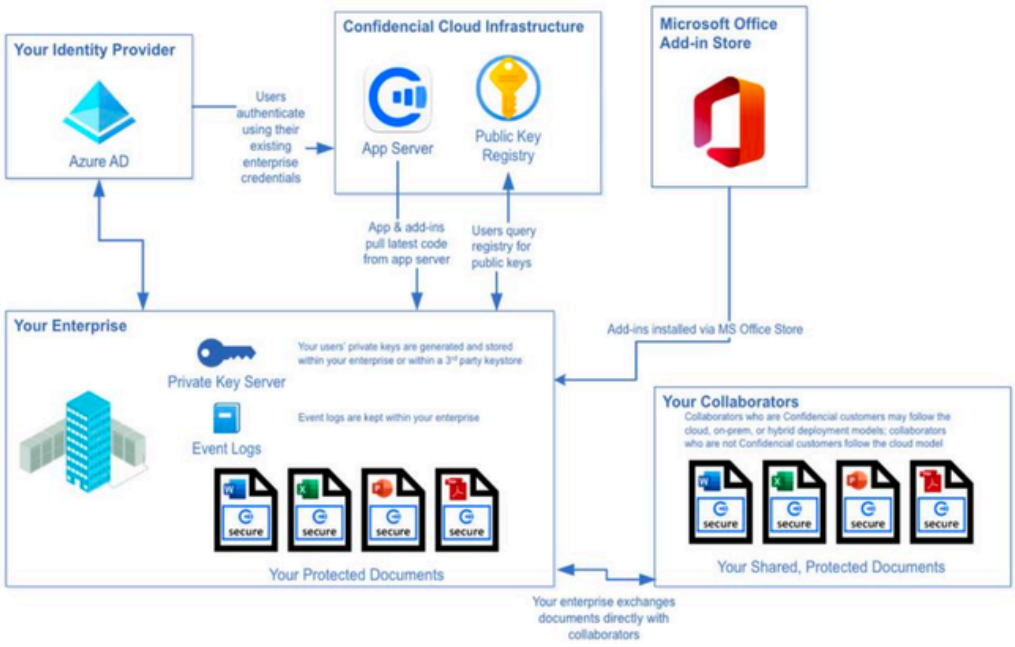

Figure 6: Hybrid deployment of Confidencial's platform with on-prem server for decryption keys and on-prem server for document events and audit-trail logs.

CLOUD DEPLOYMENT

Cloud-Based Split-Key Architecture

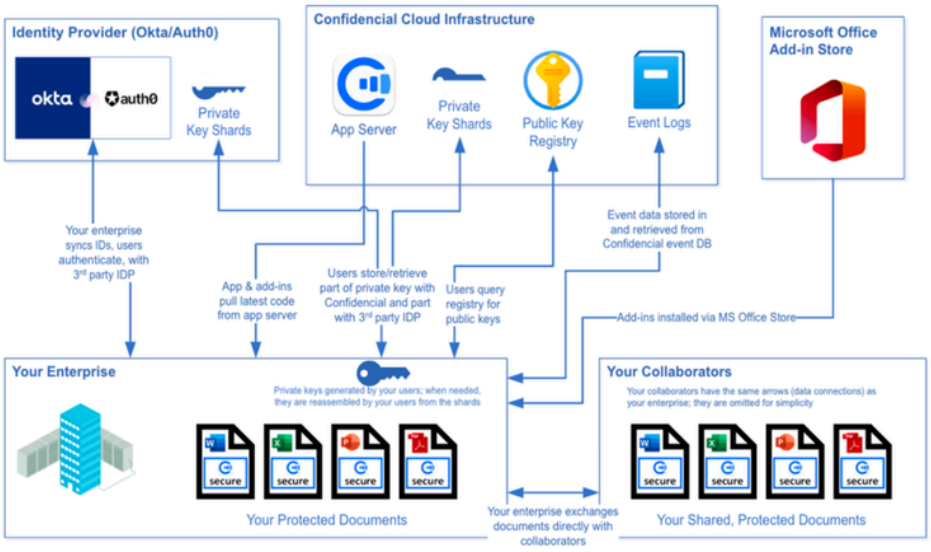

Figure 7: Cloud-based deployment of Confidencial's platform with keys split (secret shared) between Confidencial's and the enterprise's identity provider tenant. Note that Confidencial does not see encrypted content nor decryption keys; the shares of private keys stored at Confidencial are mathematically proven to be random strings that do not enable decryption nor leak information about the private keys themselves.

IP & RESEARCH

Underlying Patents

The following granted SRI patents (exclusively licensed to Confidential, an SRI spin-out) cover the technologies underlying our platform. These patents cover three major areas:

01 Selective Sharing of Content: https://bit.ly/3pDZ0kO

02 Secure and Privacy-Preserving Interactive Computation: http://bit.ly/3dD2wcO

03 Automated Identification of Sensitive Content: http://bit.ly/3dCZAgv, https://bit.ly/2M5rUg2 and https://bit.ly/3aCNFgK

There are *4 additional filed patent applications* that are not yet public. These applications cover our systems and methods for policy-based selective sharing of content using standardized encryption schemes, and secure exchange and management of credentials and cryptographic keys using a resilient, secure, and privacy-preserving overlay network that can be deployed on untrusted clouds.

BENCHMARKS

Performance, Overhead, and Usability

The overhead of AC3 protection is minimal. For illustration, assume an average encryption policy with 10 attributes (that are OR'ed or AND'ed) and that RSA 4096 bits (512B) is used for encrypting the random AES key used to encrypt the content itself. The size of the metadata added to an MS Office file will be 10 x 512B = 5120B ~ 5.2KB. This overhead and growth in size corresponds to an RSA encryption under the key corresponding to each attribute. The expected encryption/decryption speed, assuming serial operation, is ~ 10 x 2msec = 20msec on a typical laptop. Obviously, parallelization could speed it up even more, but we do not expect this to be a bottleneck for users. In addition, we have designed our platform and all plugins with the following in mind to maximize its usability and ease of adoption:

Usability

- Total number of user views to configure and encrypt/decrypt: 2
- Total number of clicks to encrypt/decrypt: 2 to 3
- Templates and high degree of automation help reduce the clicks to just 1

Encryption/Decryption Speed

- Bulk content is encrypted with inexpensive, standardized AES-256-GCM
- Expensive, policy-based public-key encryption only used to encrypt randomly generated AES keys (256 bits)

Overhead in Content Size

- Less than 1% if content size is 10K or more.

Managing Keys

- Inside Enterprise: automatically read enterprise hierarchy from (Azure) Active Directory or SAML- or OpenID-compliant IAM systems
- Across Enterprises: enroll via email and text messages (or other MFA)

Encryption Policies

- Number of encryption keys for medium: less than 1.5 x number of employees (e.g., about 3.5K for 2.5K employees)
- Number of attributes per encryption policy: < 10
- Number of encryption policies: 10s to 100s

The size of the unstructured content will only increase by several KBs for large policies of up to 10s of attributes and it will take < 0.1sec to decrypt and open the document.

WHO WE ARE

About Us

Confidencial Inc is a Menlo Park, California based provider of solutions that help organizations secure their most sensitive information, regardless of whether that information exists inside or is shared outside the organization. Confidencial's patented technology was developed under the Defense Advanced Research Projects Agency (DARPA) Brandeis and RACE programs, which were created to address the military's need to protect the private and proprietary information of its individuals and enterprises. This technology was incubated at SRI International, whose efforts to identify new opportunities, develop products, and create custom solutions, resulting in the spinoff of Confidencial Inc. Our solutions uniquely integrate directly with common desktop applications to deliver a simple point-and-click capability that secures sensitive information in a manner that does not disrupt current business processes or the creation, storage, and distribution of your documents.

Confidencial's team is composed of leading business software executives and cybersecurity experts dedicated to working with organizations across all industries to secure their sensitive documents.

GET IN TOUCH

Contact

For more information about Confidencial please contact us at info@confidencial.io

For more details about the cryptography please contact us at cryptography@confidencial.io

For more details about privacy and sales, or to send any feedback please contact us at privacy@confidencial.io, sales@confidencial.io, or feedback@confidencial.io

GENERAL
info@confidencial.io

CRYPTOGRAPHY
cryptography@confidencial.io

PRIVACY
privacy@confidencial.io

SALES
sales@confidencial.io