WORKING GROUP PAPER

CARNEGIE ENDOWMENT FOR INTERNATIONAL PEACE

PRINCETON UNIVERSITY

CITP CENTER FOR INFORMATION TECHNOLOGY POLICY

APRIL 2019

Likely Future Adoption of User-Controlled Encryption

Encryption Working Group

Likely Future Adoption of User-Controlled Encryption

Encryption Working Group

Carnegie Endowment for International Peace
Publications Department
1779 Massachusetts Avenue NW
Washington, DC 20036
P: + 1 202 483 7600
F: + 1 202 483 1840
CarnegieEndowment.org

This publication can be downloaded at no cost at CarnegieEndowment.org.

CONTENTS

About the Encryption Working Group

The Carnegie Endowment for International Peace and Princeton University have convened a small group of experts to advance a more constructive dialogue on encryption policy. The working group consists of former government officials, business representatives, privacy and civil rights advocates, law enforcement experts, and computer scientists. Observers from U.S. federal government agencies attended a select number of working group sessions. Since 2018, the working group has met to discuss a number of important issues related to encryption policy, including how the relevant technologies and uses of encryption will evolve in the future.

This paper and its companion piece on quantum computing were prepared by Princeton University's Center for Information Technology Policy at the request of the Carnegie Encryption Working Group as briefings to provide insight into future trends related to encryption policy. The papers do not take a position on encryption policy, rather they provide analysis of the future trends related to encryption and how they will shape the issues that policymakers must address.

The Encryption Working Group will continue its efforts to study this important issue and plans on releasing further briefings on aspects of the encryption policy debate around the world.

Members of the Encryption Working Group include:

Jim Baker
Former General Counsel,
Federal Bureau of Investigation

Katherine Charlet
Program Director, Technology and International Affairs, Carnegie Endowment for International Peace

Tom Donahue
Visiting Fellow, George Mason National Security Institute, and former Senior Director for Cyber Operations, National Security Council, White House

Ed Felten
Robert E. Kahn Professor of Computer Science and Public Affairs, Princeton University

Avril Haines
Senior Research Scholar at Columbia University's Columbia World Projects and former Deputy Director, Central Intelligence Agency

Susan Hennessey
Executive Editor, Lawfare, and Senior Fellow in Governance Studies,
The Brookings Institution

Chris Inglis
Managing Director, Paladin Capital Group, and former Deputy Director, National Security Agency

Sean Joyce
US Cybersecurity and Privacy Leader, PwC, and former Deputy Director, Federal Bureau of Investigation

Susan Landau
Bridge Professor of Cyber Security and Policy, Tufts University

Christy Lopez
Distinguished Visitor from Practice, Georgetown Law Center

Alex Macgillivray
Board Member, Data & Society, and former Deputy Chief Technology Officer of the United States

Jason Matheny
Founding Director, Georgetown Center for Security and Emerging Technology, and former Director, Intelligence Advanced Research Projects Activity

Tim Maurer
Co-Director and Fellow, Cyber Policy Initiative, Carnegie Endowment for International Peace

Denis McDonough
Visiting Senior Fellow, Technology and International Affairs, Carnegie Endowment for International Peace, and former White House Chief of Staff

Lisa Monaco
Distinguished Senior Fellow, Reiss Center on Law and Security, New York University School of Law, and former Assistant to the President for Homeland Security and Counterterrorism

Laura Moy
Executive Director, Center on Privacy & Technology, Georgetown Law Center

Michelle Richardson
Director, Privacy and Data Project, Center for Democracy and Technology

Ronald L. Rivest
Institute Professor, Massachusetts Institute of Technology

Ari Schwartz
Managing Director of Cybersecurity Services, Venable LLP

Harlan Yu
Executive Director, Upturn

Denise Zheng
Senior Associate (Non-resident), Technology Policy Program, Center for Strategic and International Studies

Note: This is not a comprehensive list of all members. Some wish to remain anonymous for the time being and to contribute in their personal capacity

Introduction

User-controlled encryption systems, which give the end user or customer sole control over the secret keys needed to recover data, are at the root of the "going dark" phenomenon cited by law enforcement and intelligence community (LE/IC) officials.[1] The impact of user-controlled encryption depends on how widely it will be deployed. Deployment over the next five to ten years can be predicted based on market trends, customer demand (from law-abiding users), and engineering realities. Our analysis assumes that there will be no new laws, regulations, or other mandates that limit the deployment of user-controlled encryption.

We also set aside the question of how often LE/IC officials will be able to defeat or circumvent encryption, without user cooperation, to recover data that are protected by user-controlled encryption. That is an important empirical question for the encryption policy debate, but it is not the topic of this paper. The analysis assumes only that user-controlled encryption can be a meaningful barrier to LE/IC access.

We estimate that developers will eventually incorporate user-controlled encryption into a product whenever there is customer demand for user-controlled encryption and the provider's access to data is not needed for the product's functionality.[2] The design principles behind user-controlled encryption are now well known, and the computational cost of encryption is now low enough, even for voluminous data such as video streams, that performance and developer knowledge will no longer be significant barriers to adoption.

We further estimate that providers' desire to target ads to users will not have much effect on the prevalence of user-controlled encryption. Most online ads are associated with services or content sent to a user by an established commercial party, and that party is able to provide the content to law enforcement. In addition, providers are moving away from ads that target users based on content toward ads that target users based on their identities.

However, some providers will likely want to collect and retain information for use in training machine learning (ML) systems—for product development or other business uses. Because ML technology is developing rapidly, and it may be difficult to predict which data will be useful, at least some providers will stockpile data in case it proves useful for ML applications.

Once a provider has acquired a user's data, most providers will retain the data for as long as it potentially holds future value. The cost of storing data is low and gets lower every year. And there are few laws or regulations that require providers to delete data. Providers who hope that data will be useful for ML will probably retain it. Even if the user deletes an item at the application level, the provider might still retain access to it. For example, an email provider might retain deleted emails unless it promises otherwise.

The commercial factors pushing providers toward collecting and retaining more data will tend to keep data available to law enforcement, where legal justification for access exists, unless user-controlled encryption is used.

Demand for User-Controlled Encryption

There will almost always be customer demand for user-controlled encryption. Some customers will live in, or travel to, countries that lack strong legal protections, and these customers will reasonably want user-controlled encryption to protect them from overreach by the governments of those countries. Even among customers who enjoy legal protections and voluntarily comply with legitimate legal processes, many will want user-controlled encryption as a cybersecurity precaution to prevent their data from being stolen or misused by service providers or others. This is consistent with the Least Privilege Principle, a widely accepted maxim that says systems should be designed so that every party is given the least access privilege they need to do their job.[3] Unless the provider needs access to make the product work, it is safer not to give it to them.

As a result, one can assume that user-controlled encryption will be deployed except when provider access is necessary for a product's operation. And there are two main reasons this need would arise: data recovery and server-side functionality.

Provider Access for Data Recovery

User-controlled encryption necessarily involves some kind of access control functionality—such as a password, personal identification number (PIN), secret key, or token—that is controlled by the user and must be used to access data. If the user cannot recall or locate their password or physical security token, the data will be unrecoverable.

For this reason, many products include a data recovery feature, which allows a provider to recover a user's data on the user's behalf, even if the user does not have the normal access control functionality (for example, if the user has forgotten their PIN or password). Whenever this feature is available, the provider will be able to access data in response to a legitimate legal process, with or without the user's participation.

Provider Access for Server-Side Functionality

Much of today's product functionality is implemented in a client-server fashion: by cooperation between one or more client devices in the user's possession, along with server resources in a provider's data center. For example, client code operating on a user's phone and computer enables the user to access and manage calendars, while server code operating in a company's data center keeps the master copy of each calendar and manages the synchronization of the master copy with separate replicas of each calendar stored on the phone and computer.

Some functions are most efficiently implemented on the server. One reason is the server's capacity—it typically has more storage, faster computation, and higher reliability than any client device. For example, an email application's search function is usually implemented on the server. If the user's mailbox is large, the complete mailbox contents might be stored only on the server (with each client device holding only recently accessed and new messages). The server might also maintain the index used to speed up search queries.

Another reason for implementing a feature on the server side is the server's ability to combine information from multiple users. For example, the best email spam filters use information about spam seen by one user to better recognize spam in other users' inboxes.

Server-side implementation of functions such as these almost always requires provider access to user data.[4] And, as is the case with data recovery, the provider can then turn over the data to LE/IC officials in response to a legitimate legal process.

When Will Encrypted Communications Be User-Controlled?

A message traversing a network is ephemeral by nature while it is in transit. Once the message reaches its destination and is placed into the recipient's storage, the network will forget about the message. The network has no need to keep or recover the message's content while it is in transit, because forwarding and delivering messages does not require any server-side or network-side access to message content, so message transmission protocols will always employ user-controlled encryption (also called end-to-end encryption in this instance). Modern protocols for encrypted communications not only discard messages after they reach their destination, but also discard the old encryption keys that protected past messages (taking care to ensure that they are not recoverable from any other information that is retained).

If, however, an application that is sending or receiving a message has a reason to retain the ability to recover that message, the application will keep a copy of the message in encrypted storage. Accordingly, it is likely in the future that all of the mechanisms that carry data across a network will move toward the use of end-to-end encryption, and any access to message contents will have to be via encrypted storage at a network endpoint and only in the case where an application has chosen to retain a copy of the message.

Program-to-Program Communications

Nowadays, much of the traffic on networks carries communications between software programs rather than between humans. Programs chatter to each other frequently to keep up to date on changes. While these messages are not human communications, they reveal information about the user. For example, an application or service on a user's device might send messages to convey that

the user has traveled to a new location, and potentially where that location is, without any explicit instruction from the user.

A program-to-program message is typically not retained by the recipient program but rather leads to changes in the data held by the recipient program. For example, a message conveying the user's new location might lead the recipient program to update a data record that has the user's last known location or that logs the user's movements over time. As in the case of human-to-human communications, the information will be stored if it might be needed later, and any data recovery would require access to this storage.

When Will Encrypted Storage Be User-Controlled?

As noted above, in future systems, all data are likely to be stored in encrypted form. The question is: who will control the keys needed to decrypt? Will the keys be entirely user-controlled, or will they be accessible to another party, such as a service provider, who can respond to law enforcement requests?

Simple File Storage

Most users need to store files on some storage medium, such as a device owned by the user, cloud storage owned by a service provider, or a more exotic storage technology such as a blockchain. Regardless of the medium, the user might want data recovery to be possible, in case they lose their password, PIN, or token. If so, the provider would need to hold the decryption keys (or be able to recover them without user assistance) and could give law enforcement access to the data. If the user were willing to use a storage medium that does not allow for data recovery, then user-controlled encryption would be used. However, most users will likely want a data recovery capability when storing simple files.

Replicated Cloud Apps

Data will often be replicated in multiple storage systems. For example, a user's email might be stored on a provider's server and on the user's laptop and phone. In such cases, the user will probably employ user-controlled encryption to protect data stored on the laptop and phone, even if the server-side keys are controlled by the provider. This makes sense from a cybersecurity standpoint, because the laptop and phone are at greater risk of being lost or stolen. If the server needs access for data recovery or service functionality, the server will store the data without user-controlled encryption.

A well-managed server will store all data in encrypted form, but if the server, rather than the user, controls the encryption key used on the server, the data cannot be considered protected by user-

controlled encryption. The server operator will be able to recover the data and provide it to law enforcement upon lawful request.

Predictions for Some Example Applications

The likely impact of user-controlled encryption can be evaluated by considering some common applications and using the above analytic framework to evaluate whether they are likely to adopt user-controlled encryption in ways that limit LE/IC access to data.

Email

As discussed above, email is typically replicated in a client-server fashion. For two reasons, it will likely remain accessible on the server in most cases. First, users want to be able to recover their email if they lose their password or other access control information. Second, provider access is necessary for some aspects of email functionality, such as spam detection (using information across multiple users) and the searching of large mailboxes.

Instant Messaging

Instant messages will likely only be accessible to LE/IC officials in some cases. Some users want to retain old messages, so they will want data recovery, which will enable LE/IC officials to gain access. But many other users care little about retaining old messages, so data recovery is not a compelling feature for them. In addition, there is no common feature of instant messaging systems that requires ongoing provider access to instant message content.

Enterprise Messaging

Enterprise messaging services, such as Slack, offer data recovery and promote themselves as repositories of enterprise knowledge. They also offer search capabilities over sometimes large message sets, which requires server-side access. These services will remain accessible to LE/IC officials.

Calendar Management

As discussed above, calendar information is typically replicated in a client-server fashion. It will likely remain accessible to LE/IC officials. Users want to be able to recover their calendars. The need for server-side functionality is less clear, although functions such as suggesting a time that works for multiple meeting attendees might be implemented most efficiently by the server. Regardless, the need for data recovery will require calendar data to be accessible on the server side.

Collaborative Editing

Collaborative editing tools, such as Google Docs and similar products, have the same considerations as calendar management, so calendar data will likely remain available on the server side.

Audio Conferencing

Users typically do not want to recover audio conferencing contents after the end of a session, so data recovery will probably not be a factor. And there is no compelling feature of this application that requires provider access to content. For these reasons, audio conferencing applications are likely to incorporate user-controlled encryption, and LE/IC officials will not have access to the audio content through the provider.

Video Conferencing

At first, video conferencing might seem similar to audio conferencing. However, the data size of video content is much larger than for audio content and can strain network and computational resources. So this requires server-side access to reformat a video stream to fit the capacity of a participant's network and device—thus ruling out user-controlled encryption. Over time, however, networks and devices will become more capable, so the need for this reformatting feature will decrease. In the long run, video conferencing applications are likely to follow audio conferencing in adopting user-controlled encryption.

Future Prevalence of User-Controlled Encryption

Based on this analysis, some significant applications will incorporate user-controlled encryption, and some will not. While use of user-controlled encryption may increase somewhat as encryption technology continues to mature in the marketplace, it will be far from universal.

Acknowledgments

This publication benefited greatly from the suggestions and feedback of several members of the Encryption Working Group.

Notes

1 The alternative terminology, "encryption with user-controlled keys," would be more descriptive but unwieldy. For readability, the more concise term "user-controlled encryption" is used, assuming that readers will remember that it is the decryption keys that are controlled by the user.

2 In some cases, there will be strong consumer demand for user-controlled encryption, yet provider access will be necessary for some desired product features. If these factors are nearly balanced, a provider might give each user the choice between a more functional mode of operation without user-controlled encryption and a less functional mode with user-controlled encryption. Given the engineering complexities involved in this type of product design approach, we expect that this user-choice scenario will likely remain uncommon, and most products will adopt a single solution to functionality versus privacy tradeoffs.

3 The Least Privilege Principle is closely related to the "need to know" principle that is applied to classified information. The idea is that if someone does not need access, then it is prudent to avoid granting them access because granting access can only lead to negative consequences.

4 This is not universally true, because, for some application scenarios, there are "privacy-preserving" approaches that allow limited server-side functionality to be implemented without fully exposing user data to the provider. However, these methods apply only to limited cases, and even when they do apply, providers often choose not to adopt them. Instead, providers choose to use methods that allow more data access, often because of their desire to monetize user data by selling it or by mining it to generate revenue or inform the design of new products.

CARNEGIE
ENDOWMENT FOR
INTERNATIONAL PEACE

1779 Massachusetts Avenue NW | Washington, DC 20036 | P: + 1 202 483 7600

CarnegieEndowment.org