Kenya Cybersecurity and Forensics Association

Research and Innovation Working Group

Microsoft Recall: A Different View

Authors: Muchilwa Lawrence, Dr.Satwinder Singh Rupra

Abstract

This paper examines the introduction of Microsoft's Recall application, launched on March 20, 2024, as part of the Copilot+ PCs initiative. Recall utilizes continuous screenshot capture, a generative AI model, and the Neural Processing Unit (NPU) on user devices to provide a semantically searchable history of user activities. Semantic search enhances search accuracy by interpreting natural language contextually, surpassing traditional keyword searches. Recall facilitates the retroactive retrieval of past activities with precision, ensuring immediate access through advanced indexing and categorization of user actions, visual inputs, and keyboard interactions.

Despite its innovative search capabilities, Recall's current implementation raises significant privacy and security concerns. The application is enabled by default for Microsoft Intune-managed users, storing sensitive data such as passwords and financial information without user consent. This data is stored locally, encrypted with BitLocker on supported systems, yet remains vulnerable to unauthorized access by both administrative and non-administrative users. Moreover, Recall records activities across various browsers, including private modes in Firefox and Tor, raising further privacy issues.

The inherent risks associated with Recall's data collection necessitate critical scrutiny. The UK data watchdog has initiated inquiries into its safety, highlighting the urgent need for enhanced controls and opt-in mechanisms to protect user privacy. This paper concludes with a call for Microsoft to address these security vulnerabilities promptly, recommending that users and organizations implement additional safeguards to mitigate potential threats.

Introduction

Microsoft introduced the Recall application on 20 March 2024 as part of their Copilot+ PCs efforts. It works by continuously taking screenshots of your screen. It uses a Gen AI model and the Neural Processing Unit (NPU) on your device to process that data and make it semantically searchable(mattwojo, 2024). Semantic search enables a user to input vague search queries and get specific results. This type of search is intended to improve the quality of search results by interpreting natural language more accurately and in context (Elastic, 2024). The difference between semantic search and keyword search is that keyword search returns results that match words to words, words to synonyms, or words to similar words. Semantic search looks to match the meaning of the words in the query. In some cases, the semantic search might not generate results with direct word matches, but it will match the user's intent.

The concept enables users to retroactively revisit past activities with precision, facilitating a seamless transition to a specific moment in time through a simple button command. Additionally, it systematically categorizes and indexes nearly all user actions, visual inputs, and keyboard interactions, ensuring instantaneous accessibility through robust search functionality (Beaumont, 2024). Recall is currently in preview status and optimized for select languages. This is to enable the collection of customer feedback, the development of more controls for enterprise customers to manage and govern Recall data, and improve the overall experience for users. Devices that are not powered by a Snapdragon® X series processor will require an installation of a Windows update to run Recal (Microsoft, 2024)l.

The screenshots taken by Recall are stored in the current user's AppData as part of image storage. The NPU processes them and extracts text, into an SQLite database file. Like any SQLite database, it can be accessed by standard database tools like the DB Browser for SQLilte physically or remotely by threat actors once they have a foothold. The NPU is an AI accelerator, deep learning processor, or neural processing unit that consists of a specialized hardware accelerator or computer system designed to accelerate artificial intelligence and machine learning applications

Challenges Posed.

1. Recall is enabled by default globally in Microsoft Intune-managed users, for businesses. You need to enable DisableAIDataAnalysis to switch it off.
2. Screenshots taken by Recall won't hide information such as passwords or financial account numbers. For example, if you log into online banking, your information such as account numbers, balances, purchases, etc will be saved by Recall. All this information is saved in clear text and accessible to any tools, users with access to the system
3. Snapshot retrievals are stored locally on Copilot+ PCs, residing within the confines of the local hard disk infrastructure. These snapshots are safeguarded through data encryption protocols implemented directly on the user's device. For users operating with Windows 11 Pro or an enterprise variant of Windows 11, such as an SKU, additional security measures are applied by leveraging BitLocker encryption technology. This data is accessible by info stealers
4. Nonadministrative users can access Recall data and the SQLite database used to index collected data.
5. Disabling website recording is reported to only work for the Edge browser meaning that all activities done on other browsers will be recorded. This will include disappearing messages
6. Website recording is reported to be active even when browsing in private mode on Firefox and Tor browser
7. The UK data watchdog is making inquiries with Microsoft over Recall and seeking more information on the safety of the product (Imran Rahman-Jones, 2024).
8. 3 months of history saved by default but this will be more based on your disk size

Conclusion.

Recall presents a dual nature wherein it enhances search capabilities while simultaneously expanding the potential attack surface for users. While it promises to elevate user experience, its current iteration prioritizes this enhancement at the expense of user privacy and security.

The data harvested by Recall possesses the inherent risk of exploitation, posing threats to both user privacy and security. This vulnerability is underscored by the accessibility of data to both administrative and non-administrative users. The default activation of Recall contravenes industry standards of privacy, as it lacks the opt-in mechanism typically afforded to users.
Moreover, Recall's data collection extends to sensitive information, including ephemeral messages, passwords, and PINs, further complicating privacy and security concerns. This clandestine data accumulation undermines existing protective measures established for such scenarios.
It is imperative that the Microsoft team attentively considers user feedback and promptly addresses the identified security and privacy loopholes. Organizations and individual users utilizing Recall are urged to exercise vigilance regarding its functionalities and implement appropriate safeguards to fortify their privacy and security measures.

Bibliography

Beaumont, K. (2024, May 21). *How the new Microsoft Recall feature fundamentally undermines Windows security*. Medium. https://doublepulsar.com/how-the-new-microsoft-recall-feature-fundamentally-undermines-windows-security-aa072829f218

Elastic. (2024, May 31). *What is Semantic Search?* | *A Comprehensive Semantic Search Guide*. What Is Semantic Search? https://www.elastic.co/what-is/semantic-search

Imran Rahman-Jones. (2024, May 22). *Microsoft Copilot+ Recall feature "privacy nightmare."* https://www.bbc.com/news/articles/cpwwqp6nx14o

mattwojo. (2024, May 22). *Recall Overview*. https://learn.microsoft.com/en-us/windows/ai/apis/recall

Microsoft. (2024, May 31). *Shop Copilot+ PCs* | *Microsoft*. Windows. https://www.microsoft.com/en-gb/windows/copilot-plus-pcs