China's Microsoft Hack May Have Had A Bigger Purpose Than Just Spying

npr.org/2021/08/26/1013501080/chinas-microsoft-hack-may-have-had-a-bigger-purpose-than-just-spying

Data Stolen in Microsoft Exchange Hack May Have Helped Feed China's AI Project
China broke into tens of thousands of email accounts in January. Now officials fear the breach wasn't just about spying. It was to build the next generation of artificial intelligence.

Investigations

August 26, 20215:00 AM ET

Heard on All Things Considered
Dina Temple-Raston

China's Microsoft Hack May Have Had A Bigger Purpose Than Just Spying

Listen · 7:58 7:58

- **Download**
- `<iframe src="https://www.npr.org/player/embed/1013501080/1031412236" width="100%" height="290" frameborder="0" scrolling="no" title="NPR embedded audio player">`
- **Transcript**

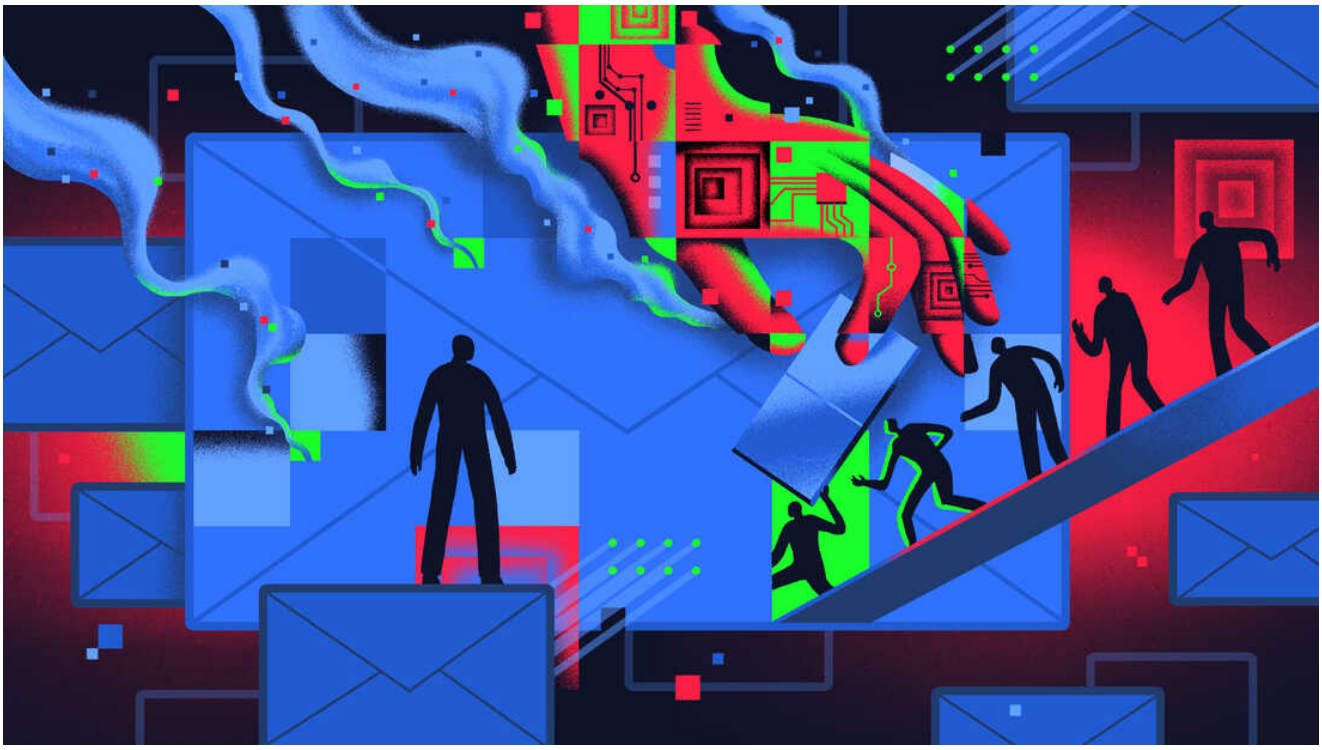

Enlarge this image

When investigators discovered the hack on Microsoft Exchange servers in January, they thought it was about stealing emails. Now they believe China vacuumed up reams of information in a bid to develop better artificial intelligence, or AI. **Matt Chinworth for NPR**

hide caption

toggle caption

Matt Chinworth for NPR

When investigators discovered the hack on Microsoft Exchange servers in January, they thought it was about stealing emails. Now they believe China vacuumed up reams of information in a bid to develop better artificial intelligence, or AI.

Matt Chinworth for NPR

Steven Adair hunts hackers for a living. Back in January, in a corner-of-his-eye, peripheral kind of way, he thought he saw one in his customer's networks — a shadowy presence downloading emails.

Adair is the founder of a cybersecurity company called Volexity, and he runs traps to corner intruders all the time. So he took a quick look at a server his client was using to run Microsoft Exchange and was stunned to "see requests that we're not expecting," he said. There were requests for access to specific email accounts, requests for confidential files.

He followed all this requested information to a virtual server off-site. "The hair is almost rising on my arms right now when I think about it," Adair told NPR later. "This feeling of like, oh, crap this is not what should be going on."

What Adair discovered was a massive hack into Microsoft Exchange — one of the most popular email software programs in the world. For nearly three months, intruders helped themselves to everything from emails to calendars to contacts. Then they went wild and launched a second wave of attacks to sweep Exchange data from tens of thousands of unsuspecting victims. They hit mom-and-pop shops, dentist offices, school districts, local governments — all in a brazen attempt to vacuum up information.

Both the White House and Microsoft have said unequivocally that Chinese government-backed hackers are to blame.

NPR's months-long examination of the attack — based on interviews with dozens of players from company officials to cyber forensics experts to U.S. intelligence officials — found that stealing emails and intellectual property may only have been the beginning. Officials believe that the breach was in the service of something bigger: China's artificial intelligence ambitions. The Beijing leadership aims to lead the world in a technology that allows computers to perform tasks that traditionally required human intelligence — such as finding patterns and recognizing speech or faces.

"There is a long-term project underway," said Kiersten Todt, who was the executive director of the Obama administration's bipartisan commission on cybersecurity and now runs the Cyber Readiness Institute. "We don't know what the Chinese are building, but what we do know is that diversity of data, quality of data aggregation, accumulation of data is going to be critical to its success."

The break-in

The intruders broke into Exchange by finding a handful of coding errors that gave them entry into Exchange servers and then allowed them to take control. Vulnerable systems just needed to satisfy two conditions: They had to be connected to the internet and controlled locally by the company's IT department, something known in cyber lingo as "on premises," or "on prem." (Microsoft's Office 365 wasn't swept up in the breach because it runs in the cloud, which offers more protection.)

The hack was fairly straightforward: Once the attackers locked onto a target and slipped into the exposed Exchange servers, they planted code that essentially tricked it into requesting information — emails, documents, PDFs — and then any servers on the other end assumed the request was legitimate.

Enlarge this image

Steven Adair, the founder of a Virginia-based cybersecurity company called Volexity, was the first to discover the Microsoft Exchange hack in the wild. "The hair is almost rising on my arms right now when I think about it," he says. **Claire Harbage/NPR hide caption**

toggle caption

Claire Harbage/NPR

Steven Adair, the founder of a Virginia-based cybersecurity company called Volexity, was the first to discover the Microsoft Exchange hack in the wild. "The hair is almost rising on my arms right now when I think about it," he says.

Claire Harbage/NPR

"It was like a conversation in which the receiving server was saying, 'Oh, you're the Exchange server, you're a trusted entity, you're allowed to do this,' " Adair said, "and basically it doesn't check that this is a completely unauthenticated request."

As soon as Adair saw that, he reached out to Microsoft.

"A relatively routine report"

These days most companies run Exchange in the cloud so Microsoft takes care of data security. Some banks, big corporations and defense companies run a hybrid system, using the cloud for a lot of their day-to-day operations but maintaining servers in-house to store

proprietary information they'd prefer to control.

Companies running their own Exchange servers tend to be small and medium-size firms, places with small IT departments that, until recently, didn't spend much time worrying about being targeted in a cyberattack. But that's exactly what happened — because if their email server was connected to the internet it meant any bad guy could hit it.

"At the time it was perceived as a relatively routine report of a couple of vulnerabilities," Tom Burt, a vice president at Microsoft who manages the digital crimes unit, told NPR. "It was in just a couple of dozen entities worldwide and just a handful in the U.S. We and the rest of the defender community see this activity happening all the time."

Microsoft has a Threat Intelligence Center, called MSTIC, that is responsible for investigating and responding to attacks. It tracks dozens of nation-state hackers and has specialists who follow particular groups. So it didn't take MSTIC long to determine who was roaming around in Exchange servers: a group of Chinese government hackers known as Hafnium.

Hafnium, Burt says, is relatively new; Microsoft has only been tracking it regularly since June 2020. It has an M.O. — it tends to target information at government agencies, medical companies and universities.

Typically, hackers find targets by scanning the internet. They look for systems that haven't been updated or patched. Investigators believe that in this case the hackers scanned the internet for companies that were running Exchange locally.

The second step of the hack was a bit more perplexing. The attackers seemed to have a weirdly specific piece of data ready to deploy: the exact email addresses of various people running Exchange servers around the world. That struck Burt as odd, because those email addresses "would be different for every single company and organization around the world," he said. "And that's not public information. So when we looked at this we thought: How is this happening?"

While on its face that was troubling, for most of January and February, the breach appeared manageable — the hack hadn't been widely deployed, and Microsoft was already at work on a patch to correct the coding errors that let the hackers in in the first place. The plan was to release it on its regularly scheduled patch day — known as Patch Tuesday, the first Tuesday of every month. But something unexpected happened: The hack went viral.

Enlarge this image

The hackers were part of a group out of China that Microsoft calls Hafnium. Tom Burt, a vice president at Microsoft who manages the digital crimes unit, says Hafnium emerged on the scene in June 2020. **Jovelle Tamayo for NPR hide caption**

toggle caption

Jovelle Tamayo for NPR

The hackers were part of a group out of China that Microsoft calls Hafnium. Tom Burt, a vice president at Microsoft who manages the digital crimes unit, says Hafnium emerged on the scene in June 2020.

Jovelle Tamayo for NPR

"All of a sudden we saw hundreds a day and then that continued to escalate until we were seeing north of several thousand a day," Burt said. "It was a very significant and noisy escalation. And as we watched that happen, we actually saw a number of different known Chinese actors and a wide range of unknown groups operating from China, all using this exploit."

John Lambert, the head of the MSTIC team, likened it to "the moment before a firecracker goes off. You know something's going to happen and you want to know: How loud is this going to be?"

Clearly, this was going to be loud — and waiting for Patch Tuesday was no longer an option.

All hands on deck

The day software fixes go out to customers is actually the end of the cycle for Microsoft's patch team. Members of the team have spent the whole month leading up to it trying to understand a vulnerability or tweak some functional problem in the software. Patch Tuesday is when the world gets to see what they've been working on and apply it to their systems.

"It's like tax day for us, but it's the runup to tax day for customers," Chang Kawaguchi, director and chief information security officer for Microsoft 365, told NPR. "That's why having a Patch Tuesday, having a consistent expectation on the customer's part, is so important to them, so they can plan for it."

If you have to release a fix anytime before a Patch Tuesday, Kawaguchi said, you ruin somebody's weekend. Instead of going to a movie, they need to be in the office testing whatever Microsoft has created to make sure it doesn't somehow lock up something else they have running on their network.

But the metastasis of the Exchange attack at the end of February meant Kawaguchi's team couldn't wait. It had to build a fix, release it and push it out to customers right away.

What made this difficult is that, initially, those in-house Exchange servers around the world weren't something Microsoft could see or had access to. If the affected servers had been in the cloud, the company could have just pushed out a patch and applied it itself. But because they weren't, Microsoft had to find a way to convince some 350,000 IT administrators running Exchange locally to stop whatever they were doing and patch right away. And that was proving to be hard.

Even putting all that aside, patches are like a ticking time bomb. They don't just protect systems, they alert criminals around the world how to get into unpatched systems. "Going public you can't just tell the good guys," Kawaguchi said. "When we release a patch, the bad guys start reverse engineering it immediately. So we always know when we release that's the starting gun of a race."

A government response

Meanwhile, anxiety about the hack was beginning to ripple through the highest levels of the Biden administration. National security adviser Jake Sullivan tweeted out a message urging IT departments to install the patches. The Cybersecurity and Infrastructure Agency released an emergency directive that warned if the malicious activity was left unchecked, it could "enable an attacker to gain control of an entire enterprise network."

The White House convened a task force — in fact, Microsoft's Burt was on it — to figure out ways to impress upon the nation's Exchange administrators just how serious this was.

Even the FBI got involved. It secured a court order so it could legally scan the internet, find servers the Chinese had breached and then proactively remove whatever they might have left there — all without informing the victims first.

"This is an active threat," press secretary Jen Psaki told reporters at the White House while all this was going on. "Everyone running these servers — government, private sector, academia — needs to act now to patch them."

Kawaguchi said later, "I think this was probably the first time a tool we built was specifically pointed to in a White House press release. There were aspects of this incident and this campaign that were definitely novel."

Enlarge this image

"This is an active threat," White House press secretary Jen Psaki, pictured here in March, told reporters as the hack started to spread. "Everyone running these servers ... needs to act now to patch them." **Samuel Corum/Getty Images hide caption**

toggle caption
Samuel Corum/Getty Images

"This is an active threat," White House press secretary Jen Psaki, pictured here in March, told reporters as the hack started to spread. "Everyone running these servers ... needs to act now to patch them."

Samuel Corum/Getty Images
Kawaguchi said in his nearly 20 years at Microsoft, he'd never seen an attack scale up so quickly. And the breadth of it seemed out of character; nation-state hackers tend to have very targeted goals, he said — they know what they want and they gather it up quietly. In this case, the Chinese were acting like cybercriminals seemingly unconcerned about who or what might get caught up in their attack.

"A lot of customers probably have felt 'I'm too small a fish,' " Kawaguchi said. " 'Nobody, no nation-state, is going to go after me.' And what we're seeing is because there's so much connection between organizations, you can go after a small fish to get to a big fish. And so everybody's having to up their game."

China, for its part, has denied any responsibility for the Microsoft Exchange attack.

"The Chinese are very Shop-Vac-oriented"

The Microsoft Exchange hack was the latest in a long list of Chinese-sponsored cyberattacks. The tally in just the four years between 2014 and 2018 is head-spinning. There was the Office of Personnel Management attack in which hackers spent some time in OPM networks and then whisked away 21.5 million records from the federal government's background investigation database.

There was also a breach at the health care insurer Anthem Inc. in which cyberthieves swiped 78 million names, birth dates and Social Security numbers. Two years after that, credit reporting agency Equifax Inc. announced that hackers stole the credit information of 147.9 million Americans. And then there was the break-in at Marriott's Starwood hotels. In 2018, Starwood announced that someone had cracked into its reservations database and stolen reservation, credit card, passport and other travel information from some 500 million people.

U.S. officials said Beijing-backed hackers were behind every one of those attacks.

"If you look, just look at the Equifax breach alone, which I consider one of the greatest counterintelligence successes by the Chinese Communist Party, they have all the financial data for every single American adult," said William Evanina, former director of the National

Counterintelligence and Security Center. "The Chinese have more data than we have on ourselves."

Evanina is now the founder and CEO of the Evanina Group, a risk consultancy company, and he said he spends much of his time fielding calls about Chinese breaches. "We've had so many, we've grown numb to it," he said. When it comes to information, he said, "the Chinese are very Shop-Vac-oriented."

Enlarge this image

It's been an open secret for years among intelligence officials that China has been on a campaign to steal massive amounts of data. The Justice Department charged Chinese government-based hackers thisyear with intellectual property theft. **Toby Scott/SOPA Images/LightRocket via Getty Images hide caption**

toggle caption

Toby Scott/SOPA Images/LightRocket via Getty Images

It's been an open secret for years among intelligence officials that China has been on a campaign to steal massive amounts of data. The Justice Department charged Chinese government-based hackers thisyear with intellectual property theft.

Toby Scott/SOPA Images/LightRocket via Getty Images

China's appetite for America's private data has been one of the biggest open secrets of modern intelligence. Intelligence officials estimate that China has now stolen all the personal identifiable information of about 80% of Americans, and it has a good start on collecting information on the remaining 20%. And while the individual breaches and numbers are worrying, the real issue is how all this information can be woven together to build on itself.

"Let's play spycatcher for a second here," says Evanina, who used to do counterintelligence for the CIA. "So you have the OPM data breach, so you have an entire security clearance file for someone, you have Anthem records, you have his Marriott point record, credit cards, Equifax, his loans, his mortgages, his credit score."

Imagine how a Chinese intelligence officer can leverage that data to get someone talking or to make a connection that can be used for intelligence purposes. "They know everything about you before they even bump you on a cruise or on a vacation," he said.

A new moonshot

For a long time, what the Chinese government intended to do with all this information was a bit of a mystery, but now, some analysts said the Microsoft Exchange hack offers some new clues. For example, remember those IT administrator emails the Chinese needed to get into the Exchange servers? Microsoft's Burt thinks they got them during an earlier Chinese hacking operation.

"What we've heard directly is they've accumulated vast quantities of data about Americans," Burt said. "And they must have created a massive database that included the actual email of who are the Exchange server administrators."

China may have been leveraging information it had stolen in other attacks. But this is the first time players have actually spoken publicly about how that happened. Intelligence officials told NPR this attack seemed more reckless in that respect.

But that may have been only a piece of a grander plan. Back in 2017, the Chinese Communist Party announced it would be making the development of world-class artificial intelligence a national priority — akin to America's race to the moon. And to do that China made clear it would focus on two things: developing computer scientists who can write algorithms, and amassing information that world-class algorithms need to learn from.

Microsoft's Burt says a specialized piece of information was needed to make the Exchange hack work — the specific email address of local Exchange server administrators. Officials say they believe the Chinese got those addresses during a previous cyberattack. **Jovelle Tamayo for NPR hide caption**

toggle caption

Jovelle Tamayo for NPR

In 2017, Chinese scholars were writing more research papers on AI than any other country in the world. China has more than 1,000 AI firms, second only to the U.S., and its universities are churning out computer scientists at breakneck speed.

China has built-in advantages in the information race. It has more than 1 billion people it can (and does) collect information about, and U.S. officials said it has been supplementing all that with large-scale data heists. (The Justice Department indicted four Chinese military hackers this year over intellectual property theft and economic espionage.)

The Cyber Readiness Institute's Todt said, against that backdrop, the second phase of the Exchange hack — when hackers hoovered up emails and information from tens of thousands of companies — shouldn't be a surprise.

Stealing information from small- and medium-size businesses out in the American heartland doesn't immediately suggest espionage. Instead, officials believe the Chinese gather this information to help them construct the informational mosaic they need to build world-class AI.

It explains their tendency, Todt said, "to gather and aggregate data and as much as possible and not discriminating where that data comes from."

The reason we should care about that is because of the role AI plays in our everyday lives. It is becoming the mechanism by which insurance rates are calculated, credit is given, mortgages are approved and health care data is calculated. And Todt said Americans should take a moment to reflect on what it would mean to have a technology that will touch our lives in a myriad of ways built by someone else and, more specifically, China.

"As it builds out its AI, China can social engineer to its priorities, to its mission," she said. "And that mission may be different from ours."