Security Now! #1062 - 01-27-26

VoidLink: AI-Generated Malware

This week on Security Now!

• CISA's uncertain future remains quite worrisome. • Worrisome is Ireland's new "lawful" interception law. • The EU's Digital Rights organization pushes back. • Microsoft acknowledges it turns over user encryption keys. • Alex Neihaus on AI enterprise usage dangers. • Gavin confesses he put a database on the Internet. • Worries about a massive podcast rewinding backlog. • What does the emergence of AI-generated malware portend?

This is what was once called “Yankee Ingenuity”

Security News

CISA's Future

I've said for years that I've been very pleasantly surprised by the success and effectiveness of CISA. Since its creation in 2015, CISA has been a huge win for our nation's cybersecurity. My default belief is that government has a difficult time getting out of its own way. So CISA was a welcome and well-needed exception to that. As this podcast has covered since CISA's inception, they've been able to mandate that government agencies pay needed attention to many specific critical security problems that would have otherwise fallen through the cracks. And they've been empowered to set deadlines. CISA's creation of KEV, the Known Exploited Vulnerabilities catalog, was a brilliant means of focusing the always-limited and readily-distracted bureaucratic resources where they were needed.

But CISA was not created permanently. Sadly, the Constitution of the United States is completely silent regarding the need for a permanent cybersecurity watchdog agency within the Federal government. This means that politicians created CISA and politicians are required to keep CISA funded and authorized. Last week, The Record updated us on the state of CISA's continuation, writing:

> *Congressional leaders on Tuesday released a compromise government funding bill that would, once again, temporarily extend the life of two key cybersecurity laws. The bipartisan legislation would reauthorize the 2015 Cybersecurity and Infrastructure Security Act (CISA) and the State and Local Cybersecurity Grant Program through the end of September.*
>
> *The extension in the $1.2 trillion funding deal is the latest short-term solution in a monthslong saga for CISA 2015, which provides liability protections to encourage private companies to share digital threat information with the federal government.*

We were talking about the legislative side of this last week. This is the political and finding side:

> *Both statutes received widespread support from the cybersecurity community and the Trump administration prior to their expiration last year. They received temporary reprieves in the continuing resolution that reopened the government in November. The House did approve a bill to extend the grants effort, but there's been no action in the Senate. Meanwhile, several proposals have been introduced to reauthorize 2015 CISA long-term. The House Homeland Security Committee last year passed legislation to renew it for a decade with minor updates but it hasn't been scheduled for a floor vote. A bipartisan Senate duo introduced a bill that would extend the law for 10 years and provide retroactive protections for companies that shared cyber threat data even after the law lapsed. But Senator Rand Paul (R-KY), chair of the Senate Homeland Security Committee, has drafted a bill that would trash the legal protection outlined in the original statute.*

Oh oh. :-(

> *House leaders plan to hold a vote later in the week on the spending deal, which boosts defense funding to over $839 billion. Lawmakers have 10 days to clear the package for President Donald Trump's signature before federal funding is set to lapse for the programs it covers. With the Senate in recess this week, the upper chamber will need to approve the legislation when they return next week if Congress is going to head off another funding lapse and a partial government shutdown.*

I don't know what's up with Senator Rand Paul. He's always something of a wildcard, and pretty much a pain in everyone's butt. But after first being elected to the Senate in 2010 he's been reelected twice since, every six years. So he seems to be what his state of Kentucky wants. But in this case, it will be very very bad if he gets his way. As we noted last week, the executives of the nation's private infrastructure agencies consider their vulnerability and breach disclosure protections to be a critical and crucial feature of this legislation. No one can make them disclose what they don't want to. So if the government wants to know what's going on, as it should, then protecting those who are voluntarily disclosing is the entire point of this aspect of the reauthorization. We should know in another week whether the politicians have now screwed up what had been a surprisingly well-designed and well-working system.

Ireland's New Lawful Interception Law

The first half of this next piece of news is that Ireland has just passed a new lawful interception law, granting the government significant new powers. The short blurb that carried that news read:

> *The Irish government has passed a new lawful interception law. The new legislation grants law enforcement and intelligence agencies the power to surveil any type of modern communications channel. It also grants the agencies the right to use covert software for their operations, such as spyware. The new law will also require communication service providers to work closely and aid any government operation.*

Add this to other recent news of pending and enacted legislation, we are clearly witnessing an accelerating trend in governments legislating themselves sweeping rights to intercept, monitor and eavesdrop upon pretty much anything they wish. Last week we noted the legislation that is now pending in Germany. This week we have similar legislation that has now passed in Ireland.

https://www.gov.ie/en/department-of-justice-home-affairs-and-migration/press-releases/minister-jim-ocallaghan-strengthens-lawful-interception-powers/

I scanned last Tuesday's press release from the Irish government from which I'll excerpt some notable pieces. The first point talks about the clear need for an update to their very old law from 1993. That's noncontroversial. But point #2 says that the new law includes *"a clear statement of the general legal principle that lawful interception powers needed to address serious crime and security threats are applicable to all forms of communications."* To call this "sweeping" would be an understatement. They write:

> *The Minister proposes an updated legal framework which is flexible and includes comprehensive principles, policies and definitions to allow for lawful interception powers to be applied to* ***any digital devices or services*** *which can send or receive a communications message e.g. the "internet of things" and email / digital messaging services. The legislation will provide for a clear statement of general principle that* ***lawful interception powers apply to all forms of communications, whether encrypted or not****, and can be used to obtain either content data (the substance of a communication) or related "metadata" (data that provide information about a communication but not its content, such as phone call or email time and date, the sender / receiver of a communication, the geolocation of an electronic device or the source and destination IP address). The legislation will also apply to parcel delivery services.*
>
> *The Minister's view is that effective lawful interception powers can be accompanied by the*

necessary privacy, encryption and digital security safeguards.
In June 2025, the EU Commission published a "Roadmap for lawful and effective access to data for law enforcement", which stated that terrorism, organised crime, online fraud, drug trafficking, child sexual abuse, online sexual extortion, ransomware and many other crimes all leave digital traces. Around 85% of criminal investigations now rely on electronic evidence. Requests for data addressed to service providers tripled between 2017 and 2022 and the need for these data is constantly increasing. The Commission paper includes proposals to deliver a "technology roadmap" on encryption issues with expert input and emphasises the need to reconcile technology and lawful access concerns through industry standardisation activities. This EU initiative complements the Minister's proposed approach to reforming the law on interception in Ireland and will inform the development of the General Scheme.

I have much to say about this. But I want to first also share the pre release's 4th point regarding *"the inclusion of a new legal basis for the use of covert surveillance software as an alternative means of lawful interception to gain access to electronic devices and networks for the investigation of serious crime and threats to the security of the State."* They write:

*The Minister also proposes to provide **a legal basis** for the use of covert surveillance software as an alternative means of lawful interception to gain access to electronic devices and networks for the investigation of serious crime and threats to the security of the State. This is used legally in other jurisdictions for a variety of purposes when necessary, such as gaining access to some or all of the data on an electronic device or network, covert recording of communications made using a device or disrupting the functioning of a personal or shared IT network being used for unlawful purposes. The Minister proposes to take into account a 2024 report from the European Commission for Democracy through Law to the Council of Europe (the Venice Commission) on this subject, "Report on a Rule of Law and Human Rights Compliant Regulation of Spyware".*

In other words, conduct that has historically been denied by governments and which no State agency would admit to using or doing, is now being ratified into law and made explicitly legal. I believe that any objective observer who has witnessed the earlier saber rattling and more recently both pending and enacted legislation would have to conclude that we are currently in an environment of slowly eroding privacy protections.

Encryption happened, and along with everyone else who appreciated knowing that their communications were private – just as a sort of an *"okay thanks, that's nice to have"* – bad guys started using it along with everyone else and they soon discovered that it protected them from law enforcement. While encryption was not created to protect criminals, the privacy it affords doesn't know or care whether you're doing good or breaking the law.

So when bad guys began hiding behind the same encryption everyone else was using, law enforcement quite reasonably asked providers for the contents of the bad guys' encrypted messages and was told that the system had been deliberately designed with to provide absolute communications content privacy for all of its users, regardless of use, and that they were unable to comply with lawful court orders to turn over their users' data. They did not have that data and they had no means for obtaining it.

That stumped the world's governments for a few cycles until someone had the bright idea to simply require the world to work differently. They said: We all agree that citizens have fundamental rights to privacy, except in cases where that privacy is being abused and is not in the public interest. So we've decided that we will determine when and where people should not

have privacy, and since we're a nation of laws, we're going to make it legal to do whatever we need to in order to obtain the privacy-violating access to our citizens' communications that we have determined we need to have ... always in support of the greater good, of course ... and besides, think of the children.

That objective observer I mentioned before would see that we're currently in a period of transition. The truth of encryption caught the world's governments off guard. They've all seen the same movies the rest of us have. Those movies uniformly depict both hackers and intelligence services *"cracking the encryption"* whenever it was really necessary to do so. Everyone knows that the *"really good encryption"* just takes somewhat longer to crack, right? The politicians just assumed that was all true. Why wouldn't they? They believed that was the way encryption really worked – right up until they encountered the truth of today's encryption, they really didn't understand that modern encryption is absolutely unbreakable. Period.

So it took them several rounds of stumbling and failed legislation to figure out that what's actually needed is for them to outlaw any encryption that no one can break. That's right. They want the encryption we have in the movies; and they're going to keep writing and rewriting legislation until they get it. So the formal legalization of the use of spyware is just the next step along that path. We're not there yet, but our objective observer would have to conclude that the world's governments, their law enforcement and intelligence agencies, will not be satisfied until it's possible to obtain access to the communications of anyone they desire.

The EDRi pushes back

The day after Ireland proudly enumerated the various features of its newly passed expansive legislation, EDRi, the European Digital Rights organization, perhaps in response to Ireland's announcement, posted under the headline *"EDRi launches new resource to document abuses and support a full ban on spyware in Europe"* It seems that the European Union has their own equivalent of the EFF and it's the EDRi. Their posted piece begins by stating *"The context: Europe's spyware crisis remains unresolved"* and they write:

> *Spyware remains one of the most serious threats to fundamental rights, democracy, and civic space in Europe. Over the past years, repeated investigations have shown that at least 14 EU Member States have deployed spyware against journalists, human rights defenders, lawyers, activists, political opponents, and others. These cases have revealed the reality of an opaque, dangerous market that thrives on exploiting vulnerabilities and endangering us, and the States' reluctance to provide any accountability or justice for victims.*
>
> *Despite the findings of the European Parliament's PEGA Inquiry Committee in 2023, and the push from human rights organisations, the European Commission has so far refused to propose binding legislation to prohibit spyware. Not only that: it has done nothing. Right now, no EU-wide red lines exist against the use of spyware.*

Quite the reverse, in fact. As we just noted, the day before this was posted Ireland announced their formal legalization of the practice. The EDRi's posting continues:

> *This means that victims lack effective remedies, authorities face no scrutiny, and commercial spyware vendors continue to operate with near-total impunity, enriching themselves by violating human rights, and even benefiting from European public funding.*
>
> *At the same time, this political inaction is increasingly being challenged. Investigative*

journalists, researchers, and civil society organisations have continued to expose spyware's human impacts, and the opaque markets behind its development and deployment. A broad coalition of civil society and journalism organisations has openly called on EU institutions to end their inaction and to adopt a full ban on commercial spyware. Adding to this push, EDRi has also adopted a comprehensive position paper calling for a full ban on spyware in the European Union as the only possible path forward from a human rights perspective.

Our collective refusal to accept the normalisation of the use of spyware is also visible inside the European Parliament. On 21 January 2026 in Strasbourg, an informal Interest Group against spyware was launched, bringing together MEPs from across political groups with the aim of maintaining scrutiny and challenging the Commission's inaction. While this does not replace legislative action, it signals that political pressure is growing, instead of fading.

Yeah. We'll see. Ireland might take exception.

The Spyware Document Pool that this posting introduces is a really terrific piece of work. I'm only going to share a tiny piece of it. But I've dropped a link to the entire pool into the show notes: https://edri.org/our-work/spyware-document-pool/

The piece I wanted to share addressed the nature and size of the commercial spyware market. They wrote:

The commercial spyware market has grown rapidly over the past decade. This market is now worth billions of euros, driven by the sale of these tools to governments, law enforcement agencies, and sometimes private actors.

Its growth is fueled by an ecosystem that combines technological sophistication with near-total opacity, allowing companies to operate across borders and evading accountability. This makes spyware a highly profitable yet extremely dangerous sector, where abuses remain hidden until uncovered by researchers or investigative journalists.

- *The global spyware industry is estimated to be worth on the order of 12 billion euros per year.*
- *More than 80 governments have contracted commercial spyware, according to the UK's cybersecurity agency.*
- *In 2023, there were at least 49 distinct vendors, along with dozens of subsidiaries, partners, suppliers, holding companies, and hundreds of investors across the supply chain.*
- *56 of the 74 governments identified by the Carnegie Endowment procured commercial spyware from firms either based or connected to Israel.*
- *The Israeli firm Paragon was acquired in 2024 by an investment firm in a deal worth up to 900 million euros.*

And this is the market that Ireland has just taken out of the shadows and made legal for their own use. The other piece of data that I thought our listeners would find interesting was about the market for the vulnerabilities that enable the creation and deployment of this spyware. They write:

The buying and selling of zero-day vulnerabilities is closely linked to the spyware market, as these flaws allow spyware to bypass security protections and operate undetected. The vulnerabilities market is dangerous because:

- *It magnifies risk: A single zero-day can compromise millions of devices. Once a vulnerability is found, the risk is anyone can exploit it.*
- *It drives innovation in spyware: Spyware vendors continuously adapt their tools to exploit newly discovered vulnerabilities.*
- *It lacks accountability: Vulnerabilities are traded secretly, with minimal regulation, creating an ecosystem with no rules that poses a risk to all of us.*
- *Concentration multiplies risk: Many people are using only two OS (Android and iOS), and some apps are globally used (Whatsapp, Gmail…). Once someone breaks into one of these systems, they can have access to hundreds of millions of devices.*

A 0-day vulnerability costs, via brokers, between:
- *5 and 7 million dollars for exploits targeting iPhones;*
- *up to 5 million for Android phones;*
- *up to 3 and 3.5 million for Chrome and Safari respectively;*
- *and 3 to 5 million dollars for WhatsApp and iMessage.*

In 2024 the Google Threat Analysis Group reported that
- *20 out of the 25 vulnerabilities found on their products (Android, Gmail), in 2023 were used by Spyware vendors to perform their attacks.*
- *As of June 2025, more than 21,500 new vulnerabilities had already been published (133 new vulnerabilities per day).*

Even though at least 14 EU countries are reported to have used commercial spyware, regulation in Europe remains entirely absent.

To say that the future of encryption currently exists in a state of tension and uncertainty would be no overstatement. Given the reality of the overwhelming power of the world's governments and the necessity for vendors to abide by their laws, much as I wish it were not the case, I do not see the interests of the EFF and EDRi ultimately winning out here. Governments are never going to be satisfied until and unless they're able to intercept and monitor the communications of specific groups of individuals under the order of their courts. That's clearly the path we're on. And as for the legal use of absolute encryption? Enjoy it while it lasts. Eventually, only criminals will be able to use unbreakable encryption. But its use will have been criminalized so that those who use it will be guilty of at least that.

Microsoft will provide its users' BitLocker keys to law enforcement

On the heels of Ireland's news and EDRi's pushback comes Microsoft's admission that they provided BitLocker keys to the FBI when asked. The headline of Thomas Brewster's piece in Forbes was *"Microsoft Gave FBI Keys To Unlock Encrypted Data, Exposing Major Privacy Flaw"* with the tag line *"The tech giant said it receives around 20 requests for BitLocker keys a year and will provide them to governments in response to valid court orders. But companies like Apple and Meta set up their systems so such a privacy violation isn't possible."* Here's what we learn from Forbes' reporting:

Early last year, the FBI served Microsoft with a search warrant, asking it to provide recovery keys to unlock encrypted data stored on three laptops. Federal investigators in Guam believed the devices held evidence that would help prove individuals handling the island's Covid unemployment assistance program were part of a plot to steal funds.

The data was protected with BitLocker, software that's automatically enabled on many modern Windows PCs to safeguard all the data on the computer's hard drive. BitLocker scrambles the data so that only those with a key can decode it. It's possible for users to store those keys on a device they own, but Microsoft also recommends BitLocker users store their keys on its servers for convenience. While that means someone can access their data if they forget their password, or if repeated failed attempts to login lock the device, it also makes them vulnerable to law enforcement subpoenas and warrants. In the Guam case, Microsoft handed over the encryption keys to investigators.

Microsoft confirmed to Forbes that it does provide BitLocker recovery keys if it receives a valid legal order. Microsoft spokesperson Charles Chamberlayne said "While key recovery offers convenience, it also carries a risk of unwanted access, so Microsoft believes customers are in the best position to decide how to manage their keys." He said the company receives around 20 requests for BitLocker keys per year and in many cases, the user has not stored their key in the cloud making it impossible for Microsoft to assist.

The Guam case is the first known instance where Microsoft has provided any encryption key to law enforcement. Back in 2013, a Microsoft engineer claimed he'd been approached by government officials to install backdoors in BitLocker, but had turned the requests down. Senator Ron Wyden said in a statement to Forbes: "It is simply irresponsible for tech companies to ship products in a way that allows them to secretly turn over users' encryption keys. Allowing ICE or other Trump goons to secretly obtain a user's encryption keys is giving them access to the entirety of that person's digital life, and risks the personal safety and security of users and their families."

Law enforcement regularly asks tech giants to provide encryption keys, implement backdoor access or weaken their security in other ways. But other companies have refused. Apple in particular has repeatedly been asked for access to encrypted data in its cloud or on its devices. In a highly publicized showdown with the government in 2016, Apple fought an FBI order to help open phones belonging to terrorists who shot and killed 14 in San Bernardino, California. Ultimately, the FBI found a contractor to hack into the iPhones.

Privacy and encryption experts told Forbes the onus should be on Microsoft to provide stronger protection for consumers' personal devices and data. Apple, with its comparable FileVault and Passwords systems, and Meta's WhatsApp messaging app also allow users to backup data on their apps and store a key in the cloud. However, both also allow the user to put the key in an encrypted file in the cloud, making law enforcement requests for it useless. Neither Apple nor Meta are reported to have turned over encryption keys of any kind in the past.

Matthew Green, cryptography expert and associate professor at the Johns Hopkins University Information Security Institute said: "This is private data on a private computer and they made the architectural choice to hold and retain access to that data. They absolutely should be treating it like something that belongs to the user. If Apple can do it, if Google can do it, then Microsoft can do it. Microsoft is the only company that's not doing this," he added. "It's a little weird. The lesson here is that if you [meaning Microsoft] have access to [its users] keys, eventually law enforcement is going to come."

Jennifer Granick, the ACLU's surveillance and cybersecurity counsel, raised concerns about the breadth of information the FBI could obtain if agents were to gain access to data protected by BitLocker. She said: "The keys give the government access to information well beyond the time frame of most crimes, everything on the hard drive. Then we have to trust that the agents only look for information relevant to the authorized investigation, and do not take

advantage of the windfall to rummage around."
In the Guam case, the court docket shows the warrant was successfully executed. The lawyer for defendant Charissa Tenorio, who pleaded not guilty, said the information provided to her by the case's prosecutors included information from her client's computer and that it included references to BitLocker keys that Microsoft had provided the FBI. The case is ongoing.

Both Matthew Green and Jennifer Granick said Microsoft could have users install a key on a piece of hardware like a thumb drive, which would act as a backup or recovery key. Microsoft does allow for that option, but it's not the default setting for BitLocker on Windows PCs.

Without the encryption keys from Microsoft, the FBI would've struggled to get any useful data from the computers. BitLocker's encryption algorithms have proven impenetrable to prior law enforcement attempts to break in, according to a Forbes review of historical cases. In early 2025, a forensic expert with ICE's Homeland Security Investigations unit wrote in a court document that his agency did "not possess the forensic tools to break into devices encrypted with Microsoft BitLocker, or any other style of encryption." In one previous case, federal investigators obtained keys by discovering that a suspect had stored them on unencrypted drives.

Now that the FBI and other agencies know Microsoft will comply with warrants similar to the Guam case, they'll likely make more demands for encryption keys, Green said. "My experience is, once the U.S. government gets used to having a capability, it's very hard to get rid of it."

Okay. So the first takeaway from this is obvious – and it doesn't involve any sort of moral or ethical judgement either way. It's just the facts: Because encryption is absolute and unforgiving, ***it can be super useful*** to have a backup plan of some kind. Someone who will never forget. Someone to hold onto one's emergency encryption backup keys. There's no doubt about that. Only if you are willing to take sober and full responsibility for never forgetting how to log in would it make sense to have no backup whatsoever, anywhere. That said, one option is to allow Microsoft to be the entity to hold onto your keys in the event of an emergency. They're certainly the default easy choice. The only downside to that – and, again, without any judgement – is that they will also turn your keys over to law enforcement after a judge approves their request. And may not be a bad thing if you're certain that this would never become an issue. But if that's a concern, it's a good thing you're now aware that Microsoft cannot be a trusted guardian of your privacy. They will capitulate and now all global law enforcement and intelligence services know that. So it might be better to entrust those secrets to a close friend who law enforcement would never think to ask.

But I said this was the first takeaway. There's another, and it's much more subtle. But I very much want to point it out to our listeners. This Forbes article reminded me of that previous instance, 13 years ago back in 2013 when a Microsoft engineer claimed he'd been approached by government officials to install backdoors in BitLocker. My recollection was that it was more than a claim, and that it was more than once. For one thing, there were multiple people involved. So it wasn't just hearsay. I have no problem with the FBI asking. As the saying goes: "Well... you can ask..."

To set this up for our listeners I need to share the first portion of Mashable's coverage of this incident from 2013. Mashable's coverage of this story was introduced with the leading question headline *"Did the FBI Lean On Microsoft for Access to Its Encryption Software?"* they wrote:

The NSA is not the only government agency asking tech companies for help in cracking technology to access user data. Sources say the FBI has a history of requesting digital backdoors, which are generally understood as a hidden vulnerability in a program that would, in theory, let the agency peek into suspects' computers and communications.

In 2005, when Microsoft was about to launch BitLocker, its Windows software to encrypt and lock hard drives, the company approached the NSA, its British counterpart the GCHQ and the FBI, among other government and law-enforcement agencies. Microsoft's goal was twofold: get feedback from the agencies, and sell BitLocker to them.

However, the FBI, concerned about its ability to fight crime – specifically, child pornography – apparently repeatedly asked Microsoft to put a backdoor in the software. A backdoor - or trapdoor – is a secret vulnerability that can be exploited to break or circumvent supposedly secure systems. For its part, the FBI categorically denies asking for such access, telling Mashable that the Bureau does not ask for backdoors, and that it only serves companies lawful court orders when it needs to access users' data. (And, legally, it would still need a warrant even if a backdoor did exist.)

Peter Biddle, the head of the engineering team working on BitLocker at the time, revealed to Mashable the exchanges he had with various government agents. Biddle told Mashable "I was asked multiple times.", confirming that a government agency had inquired about backdoors, though he couldn't remember which one. He said: "And at least once the question was more like 'If we were to officially ask you, what would you say?' "

According to two former Microsoft engineers, FBI officials complained that BitLocker would make their jobs harder. An FBI agent reportedly said: "It's going to be really really hard for us to do our jobs if every single person could have this technology. How do we break it?"

The story of how the FBI reportedly asked Microsoft to backdoor BitLocker to avoid "going dark" – the FBI's term for a potential scenario when encryption makes it impossible to intercept criminals' communications or break into a suspect's computer – provides a snapshot into how U.S. government agencies try to persuade tech companies to weaken their security products, or even poke a hidden hole to make them wiretap-friendly.

Last week [meaning 13 years ago in 2013], The New York Times, ProPublica, and The Guardian revealed that one of the ways the NSA circumvents Internet cryptography is to ask companies to put backdoors in their products. The FBI is reportedly doing the same in the name of fighting crime, and its persuasion techniques appear to be very similar. According to reports, both the NSA and the FBI are subtle in their requests, which are never formal, written requests, but are usually uttered during casual conversations, almost jokingly.

Nico Sell, the founder of the privacy-enhancing app Wickr, was approached by an FBI agent after speaking at the RSA security conference at the end of February, as first reported by CNET. According to Nico Sell, the agent asked: "So are you gonna give us a backdoor?" She declined, and after pressing the agent – asking him to explain if he had a written request and to reveal his boss – the agent backed down. Cryptography and security expert Bruce Schneier said he's heard of these same types of tactics from others the government has approached seeking technological backdoors.

Bruce told Mashable: "It's never an explicit ask. It's an informal, oblique mention, joking conversation, where you're felt out as to whether you're amenable to it. If you're amenable to it the conversation continues, if you're not, it's like it never happened."

Despite the requests being informal, Schneier and other surveillance experts are concerned. "A request is a request," and despite not being illegal, he said, "it's coercive."

In the case of Microsoft, according to the engineers, the requests came in the course of multiple meetings with the FBI. These kinds of meetings were standard at Microsoft, according to both Biddle and another former Microsoft engineer who worked on the BitLocker team, who wanted to remain anonymous due to the sensitivity of the matter.

Biddle said: "I had more meetings with more agencies that I can remember or count." He said the meetings were so frequent, and with so many different agencies, he doesn't specifically remember if it was the FBI that asked for a backdoor. But the anonymous Microsoft engineer we [Mashable] spoke with, confirmed that it was, in fact, the FBI.

During a meeting, according to Biddle and the Microsoft engineer who were both present at the meeting an agent complained about BitLocker and expressed his frustration, saying "You guys are giving us the shaft." Though Biddle insisted he didn't remember which agency he spoke with, he said he did recall this particular exchange. And Biddle wasn't intimidated. He replied: "No, we're not giving you the shaft, we're merely commoditizing the shaft."

Biddle, a believer in what he refers to as "neutral technology," never agreed to put a backdoor in BitLocker. And other Microsoft engineers, when rumors spread that there was one, later denied that was ever a possibility. Niels Ferguson, Microsoft's cryptographer and principal software development engineer wrote: "The suggestion is that we are working with governments to create a backdoor so that they can always access BitLocker-encrypted data. That will happen over my dead body."

For Biddle, this was proof of a fundamental paradox facing government agencies and security software. How do you get secure software you can rely on, while also retaining the ability to break into it if people use it to commit or cover up their crimes? Biddle said: "I realized that we were in this really interesting spot, sort of stuck in the middle between wanting to do a much better job at protecting our users' information, and at the same time realizing that this was starting to make government employees unhappy."

Despite Microsoft's refusals to backdoor its product, the engineers kept working with the FBI to teach them about BitLocker and how it was possible to retrieve data in case an agent needed to get into an encrypted hard drive. At one point, the BitLocker team suggested the agency target the backup keys that the software creates. In some instances, BitLocker prompts users to print out a piece of paper with the key needed to unlock the hard drive, to prevent loss of data if a user forgets his or her key. The anonymous Microsoft engineer said: "As soon as we said that, the mood in the room changed dramatically. They got really excited."

In that instance, law enforcement agents wouldn't need a backdoor after all. As the engineer suggested, all they would need was a warrant to access a suspect's documents and retrieve the document that would unlock his or her hard drive.

And this finally brings me to the point I wanted to make. Mashable quotes Christopher Soghoian, writing:

For Christopher Soghoian, a privacy and surveillance expert at the ACLU, whether or not BitLocker has a backdoor or not isn't even that relevant, since it's a feature that very few Windows users employ or even have access to. It's not included in most Windows versions,

and it's not a default setting – something that, Soghoian said "is not an accident." He told Mashable: "The impact is minimal, because so few people use BitLocker, but it does speak to a friendly relationship between the companies and the government." He said: "If you want to keep your data out of the U.S. government's hands, Microsoft is not your friend. Microsoft is unwilling to really make the government go dark. They are never really willing to protect their customers from the government. They are willing to take some steps but they don't want to go too far."

I wanted to share that last bit because I think that's wrong.

First of all, this is a reminder about the way the world has changed during the intervening 13 years. At the time Christopher Soghoian was quoted, he correctly noted that BitLocker was a non-issue since it was so infrequently used. I doubt he would feel the same way about BitLocker today. It will not be enabled on machines that have been upgraded to Windows 11 if earlier Windows was not using it. But most modern PCs that ship with Windows 11 preinstalled, even Home Edition with its simpler Device Encryption, will have their hard drives encrypted out of the box.

To me, this seems entirely pro-consumer. Microsoft doesn't have to push BitLocker encryption. It certainly causes some pain and annoyance for both them and their Windows users. But it's extremely good for the privacy of their users that someone cannot remove their machine's drive and mount it on another machine to dump its entire contents.

Might this mean that the FBI needs to obtain a court order to compel Microsoft to disclose the encryption keys of someone whom they have convinced a judge may have evidence crucial to a crime they are working to solve? Yes, that might be necessary. But all of that is only necessary because Microsoft is pushing everyone to encrypt their drives in the first place – and there's no way ***any*** law enforcement agency anywhere is happy about that. If Microsoft were not defaulting to using BitLocker, things would still be as they were 13 years ago when Christopher Soghoian said *"Who cares anyway? No one uses BitLocker!"* Today, most new systems do. And most Windows users obtain the huge benefit of having their drive's data much more securely protected from casual non-governmental attack than if it were not encrypted.

I won't disagree that Microsoft might be able to do more, and that they may be shortchanging their user's privacy when push comes to shove. If they can provide unlocking keys to their users in an emergency, then they can also provide them to law enforcement when under order to do so. On balance, I would venture that many many many more Windows users' data have been **saved** by this policy than have been **compromised** by law enforcement subpoena. I'm sure that Microsoft would always require a court order before disclosing. It's not as if anyone can just ask for someone else's decryption keys. So the data does have the same protection as does our other personal property and possessions. The protection is not absolute, no. And, yes, if users were capable to taking full responsibility it could be. But at least in the United States, the protection is in line with what U.S. citizens enjoy in the other areas of their lives.

So... anyone who objects to turning their keys over to Microsoft is now forewarned that their keys may be disclosed upon legal demand. If that's a concern, Windows Pro, Enterprise and Education edition users can disable Microsoft's default key escrowing by setting two policies:

- Do NOT allow BitLocker recovery information to be stored in Azure AD
- Do NOT allow BitLocker recovery information to be stored in Microsoft account

After doing that, their Bitlocker recovery keys can be rotated so that Microsoft never obtains the updated keys.

Listener Feedback

Alex Neihaus — Alex shares some tremendous business-centric perspective on what the application of AI will mean to the enterprise and what pitfalls it is more than likely to offer:

Gents, Security Now! has spent a lot of time over the last couple of months on AI and how its first, most "natural" application is software development. After a recent experience packaging up a hobby script as a public open source PowerShell module, I couldn't agree more that the development toolset is rapidly changing. But…(there's always a but)…I worry about "mechanically" produced code, particularly in enterprise systems that deal with financial and personal information at scale (think a brokerage or a multi-state healthcare system).

If we look at the historical waves of management thinking about the development costs of crucial enterprise systems, we see an endless push to reduce those costs. That has inevitably led to declines in quality, reliability and most of all security. In the early days of enterprise development, engineers worked in-house. I started as a developer at Mass. General Hospital in the 1980s when "outsourcing" wasn't in the lexicon. Yet MGH developed MUMPS in-house which today is the core database environment of the largest electronic medical records vendor in the USA.

Outsourcing became all the cost-cutting rage among enterprises, followed-on quickly by the offshoring of enterprise development. Executives, based on then-current business consulting doctrine, decided that IT wasn't their core business. They thought their businesses just made widgets or sold products or provided a service. IT was orthogonal to their core function. History shows what a strategic mistake that thinking was. It led directly to the situation we find ourselves in today: a race to the bottom to procure development resources that cost a fraction of in-house resources. Security Now! regularly documents the results: breaches, botnets, system failures and worse.

*Over time, enterprises discovered their IT systems **are** the business. You can't make a widget, sell it or service it without an enterprise system. Unfortunately, many businesses continue to dismantle their core capabilities with a massive, mistaken "shoemaker's children" syndrome.*

AI's rapid development means yet another giant epoch in computing technology is just starting. We're about to live through another turn of the wheel of technological progress.

Having used AI in a simple vibe coding project, it's clear to me that AI cannot replace developers. But that's how it's being pitched to the same enterprises that previously committed the "not our core function" mistake. Amazon and Microsoft are both insisting that replacement is the primary benefit. Enterprises, now completely beholden to these hyperscalers, take their cues from them. I use AWS and Azure cloud products every day in real client situations; I can tell you quality isn't their north star. Anyone who's ever struggled with Microsoft's automated API documentation can tell you it isn't worth the electrons used to display it.

IOW, we should not repeat the mistaken business assumptions that drove the outsourcing debacle. Instead, we can upscale developers' skills while still retaining the focus on human development. Imagine the benefits if we used AI not to replace those $25/hour outsourced developers who produce the worst code I've ever seen, but instead train them to use AI to write tests, to check the level of every included NPM or PyPI package against the CVSS database before recompiling or to fuzz their functions. (That last is the hallmark of outsourced code vulnerabilities: it works but only on the happy path.)

We're still in the early hype-cycle days of AI. But the hype sometimes becomes at least the partial reality. Thinking about AI only as replacement for developers makes the same mistake we made with outsourcing, only magnified many times.

You both know how much I love the show — and your tenacity in producing compelling podcasts, week after week, decade after decade. Thanks for that. SN!'s #1 fan, Alex

I wanted to share Alex's perspective because I think it is super valuable and exactly correct. It also makes such complete sense. My own perspective tends to get wrapped up in the technology, so the effect AI will be having on internal enterprise development isn't the sort of thing I tend to focus upon. So thank you, Alex. I imagine this may give many of our enterprise-centric listeners something to think about and perhaps discuss with their peers and managers.

Gavin (has a confession to make)

Hi Steve, I have a confession to make: I have knowingly opened up public database access in production systems. Here's how this came about.

A few years ago I became the sole software developer at a small UK ISP, following the departure of several senior team members. The company had a plethora of legacy systems scattered among various cloud service providers. Following COVID-19, sales plummeted with many customers shutting up shop and very few businesses investing in connectivity products. We had to start cutting costs or risk going under ourselves.

One of our biggest costs was managed database instances. My predecessors had spun up individual database servers (mainly mysql and postgres) for each of our many applications across different clouds (AWS, Digital Ocean, etc) and multiple different accounts within each. In terms of application isolation, this was a great approach, but it was costing a small fortune.

My task was to consolidate as many of these databases as possible, which would bring our costs down quickly (amounting to thousands of pounds per month). There were three possible ways of accomplishing this after first migrating all application databases to a central instance:

- *Move all of the applications from the various cloud accounts into a single account and VPC so they could all access the new database instances privately.*
- *Give each of our applications static IPs and set up security rules in front of the new database instances to limit access*
- *Open up full public access and make the passwords as strong as possible*

The problem with moving all of our applications is that it would take a very long time and many of them relied on cloud provider-specific utilities and network configurations for which we would need to find alternatives and rewrite large swathes of legacy code.

The problem with the static IP solution is that they became quite expensive and some of our platforms (Digital Ocean Apps for example), at the time, didn't offer them at all.

So, reluctantly, the third option was chosen and management were happy to take on the additional risk, which I explained to them, especially in light of the immediate expected cost savings.

And so, for about 4 years, we were running with our main databases publicly exposed to the

Internet. But now, today, after a lot of work, all of our apps are in private subnets and linked VPCs, and, thankfully, our databases are no longer exposed.

I also know that my company was not alone in doing this sort of thing, having spoken with other devs in the industry. So here it is, a real world example of how this can happen and not through negligence, rather through unfortunate resource pressures.

Luckily, we were never compromised (to our knowledge) and we just about managed to bounce back as a business. And I'm making sure this sort of thing cannot happen again.

All the best, Gavin, listening since 2018

First of all, Gavin, confession is always good for the soul. Second, I have no problem whatsoever with what you needed to do, because none of it was done blindly or without thought and a clear understanding and balancing of the costs and potential consequences. So I would judge that to be, while not entirely safe, at least entirely responsible. You were not irresponsible. And given the constraints you were operating under, the interim solution you adopted was the best you could achieve. No one should find fault with that.

John David Hickin

John's email subject was: "On ISP's Selling Your DNS Data." He wrote:

Can't you just set up the ISP's modem/router as an edge router (turning off WiFi as well) and connect another router or more behind that? An old solution of yours repurposed. Cheers, John.

I received other similar questions, so I wanted to take a minute to examine the IPS's advantage. What we need to consider is that our ISPs – mine is Cox Cable – knows exactly who we are by household name, address and payment information, and they are in the very special and privacy sensitive position of having direct access to our individual Internet traffic. We've invested endless podcasts examining cookies and fingerprinting and tracking beacons and all manner of privacy breaching and privacy protecting solutions and technologies. And there, among it all, sits our ISPs through which every bit, byte, kilobyte, megabyte and gigabyte of our traffic flows. No one else on the entire planet enjoys such direct access to exactly what those in our household are doing from moment to moment.

Before Let's Encrypt's great "Encrypt the Internet" TLS revolution, our ISPs were often privy to the detailed content of everything we did. In retrospect it was quite bracing. Today with everything encrypted, ISPs are unable to see into our connections, but they can still see where we're connecting and if we are not encrypting our DNS, they can also track every domain name anyone in our household, and any of our home's IoT devices, looks up.

Tracking the remote IP we're connecting to is much less useful today than it was years ago due to the massively widespread use of multi-domain hosting. Cloudflare has a large pool of IP addresses which provides services to their large pool of customer websites. Among those there's a many-to-many relationship. So having an ISP only able to see a customer's traffic destination is far less useful today than it was 20 years ago.

However, there's still the problem that the SNI – Server Name Indication – that's carried in the TLS Client Hello handshake is only encrypted when both ends support and negotiate TLS v1.3. TLS v1.3 introduced ECH – Encrypted Client Hello – for the express purpose of preventing

anyone who might be examining Internet traffic from picking up the destination domain of any new connection. At this point in time, at the start of 2026, more than half of all Internet traffic is using TLS v1.3. So the privacy leakage that has continued to occur during the TLS handshake is draining away slowly over time.

After destination IP and TLS handshake leakage, DNS is the remaining potential privacy leak. Firefox users are now being automatically protected thanks to an agreement between Mozilla and Cloudflare to use Firefox's built-in DNS over HTTPS with Cloudflare's DNS resolvers by default.
By default, the Chromium browser family will upgrade to DoH if the DNS provider the user has manually configured for unencrypted DNS also supports DoH. This is called "Opportunistic DoH". But since most users have not manually reconfigured their DNS and just run with whatever DNS their ISP provides, that will be unencrypted DNS over UDP. So only people using Firefox will have their DNS lookups masked from ISP snooping.

One increasingly popular solution is to use or obtain a home router that can perform its own remote DoH lookups and configure it to use one of the major free CDN DNS solutions offered by Cloudflare, Google, OpenDNS, NextDNS or whatever service you choose. After that, all of the internal network's DHCP-configured devices will be using standard DNS to your router, and all queries for domain names will be encrypted and handled by your router.

With everything using TLS, and TLS moving to v1.3 with Encrypted Client Hello to mask the target domain, and your browser or router using DoH, the only remaining privacy concern is an ISP able to observe the destination of your traffic. As I noted, thanks to the increasing use of CDNs and cloud hosting which aggregate many domains among IP addresses, that's far less certain than it once was. But for anyone desiring absolute privacy from ISP snooping, the final steps would be to use some form of traffic tunneling. So that means TOR or a VPN. Using one of those means that the ISP is able to determine nothing beyond the fact that you're using TOR or a VPN. That they can see. But nothing else.

Circling back around to John's question, it should be clear that as long as an ISP is the carrier of a subscriber's traffic, nothing the user might do inside their own network, such as placing one router behind another, would change the nature of the traffic which emerges from their LAN to pass under their ISP's watchful and perhaps curious traffic logging or monitoring, if any such existed.

Troy Shahoumian

Steve, OMG! I just realized how many podcast files I have that probably need rewinding! Can you recommend a program to do this?

I feel your pain, Troy. And I can only imagine the size of the podcast rewinding backlog you might be facing down. I know that many of our listeners also listen to other podcasts. So your burden might be even larger. But just taking Security Now!, since this is podcast 1,062, if you've been listening from the start, or if you started late then went back to catch up from the beginning – and if, God help you – you have not previously rewound **any** of those before now, well, whoof ... it's not going to be pretty. I don't envy the corner you've painted yourself into. What you really need is some sort of mass, gang, parallel, podcast rewinder. Once I finish reworking GRC's eCommerce facility, and Lorrie and I get moved into our new place, I plan to readdress GRC's ValiDrive freeware. And while I didn't have it on my roadmap, I may see whether I might be able to sneak in some sort of mass podcast rewinding facility as a side feature. I can certainly see the clear need. Until then, try to hang in there and we'll keep you in our prayers!

VoidLink: AI-Generated Malware

What we've been expecting has happened – and it's every bit as bad as we worried it would be. Last Tuesday, CheckPoint Research published their analysis of a newly discovered malware which they named *"VoidLink"*. Their research was titled *"VoidLink: Evidence That the Era of Advanced AI-Generated Malware Has Begun"* What we all knew had to be coming has arrived. CheckPoint summarized this news with five key points:

- *Check Point Research (CPR) believes a new era of AI-generated malware has begun. VoidLink stands as the first evidently documented case of this era, as a truly advanced malware framework authored almost entirely by artificial intelligence, likely under the direction of a single individual.*

- *Until now, solid evidence of AI-generated malware has primarily been linked to inexperienced threat actors, as in the case of FunkSec, or to malware that largely mirrored the functionality of existing open-source malware tools. VoidLink is the first evidence based case that shows how dangerous AI can become in the hands of more capable malware developers.*

- *Operational security (OPSEC) failures by the VoidLink developer exposed development artifacts. These materials provide clear evidence that the malware was produced predominantly through AI-driven development, reaching a first functional implant in under a week.*

- *This case highlights the dangers of how AI can enable a single actor to plan, build, and iterate complex systems at a pace that previously required coordinated teams, ultimately normalizing high-complexity attacks, that previously would only originate from high-resource threat actors.*

- *From a methodology perspective, the actor used the model beyond coding, adopting an approach called Spec Driven Development (SDD), first tasking it to generate a structured, multi-team development plan with sprint schedules, specifications, and deliverables. That documentation was then repurposed as the execution blueprint, which the model likely followed to implement, iterate, and test the malware end-to-end.*

We've been rejoicing over the surprising jump Claude Code, for example, has made in enabling end-to-end creation of applications. As they say "Everyone's doing it!" Unfortunately, we've known that "everyone" would eventually include malware authors. That has now happened and it's as bad as we worried it would be. I'm not going to examine this in depth because what's the point? There will be another one tomorrow. This is clearly the beginning of an entirely new problem domain. Nevertheless, CheckPoint's Introduction is worth sharing. They wrote:

When we first encountered VoidLink, we were struck by its level of maturity, high functionality, efficient architecture, and flexible, dynamic operating model. Employing technologies like eBPF and LKM rootkits and dedicated modules for cloud enumeration and post-exploitation in container environments, this unusual piece of malware seemed to be a larger development effort by an advanced actor. As we continued tracking it, we watched it evolve in near real time, rapidly transforming from what appeared to be a functional development build into a comprehensive, modular framework. Over time, additional components were introduced, command-and-control infrastructure was established, and the project accelerated toward a

full-fledged operational platform.

In parallel, we monitored the actor's supporting infrastructure and identified multiple operational security (OPSEC) failures. These missteps exposed substantial portions of VoidLink's internal materials, including documentation, source code, and project components. The leaks also contained detailed planning artifacts: sprints, design ideas, and timelines for three distinct internal "teams," spanning more than 30 weeks of planned development. At face value, this level of structure suggested a well-resourced organization investing heavily in engineering and operationalization.

However, the sprint timeline did not align with our observations. We had directly witnessed the malware's capabilities expanding far faster than the documentation implied. Deeper investigation revealed clear artifacts indicating that the development plan itself was generated and orchestrated by an AI model and that it was likely used as the blueprint to build, execute, and test the framework. Because AI-produced documentation is typically thorough, many of these artifacts were timestamped and unusually revealing. They show how, in less than a week, a single individual likely drove VoidLink from concept to a working, evolving reality.

As this narrative comes into focus, it turns long-discussed concerns about AI-enabled malware from theory into practice. VoidLink, implemented to a notably high engineering standard, demonstrates how rapidly sophisticated offensive capability can be produced, and how dangerous AI becomes when placed in the wrong hands.

The general approach to developing VoidLink can be described as Spec Driven Development (SDD). In this workflow, a developer begins by specifying what they're building, then creates a plan, breaks that plan into tasks, and only then allows an agent to implement it.

Artifacts from VoidLink's development environment suggest that the developer followed a similar pattern: first defining the project based on general guidelines and an existing codebase, then having the AI translate those guidelines into an architecture and build a plan across three separate teams, paired with strict coding guidelines and constraints, and only afterward running the agent to execute the implementation.

VoidLink's development likely began in late November 2025, when its developer turned to TRAE SOLO, an AI assistant embedded in TRAE, an AI-centric IDE. While we do not have access to the full conversation history, TRAE automatically produces helper files that preserve key portions of the original guidance provided to the model. Those TRAE-generated files appear to have been copied alongside the source code to the threat actor's server, and later surfaced due to an exposed open directory. This leakage gave us unusually direct visibility into the project's earliest directives.

In this case, TRAE generated a Chinese-language instruction document. These directives offer a rare window into VoidLink's early-stage planning and the baseline requirements that set the project in motion.

TRAE (spelled T.R.A.E.) is a creation of ByteDance, the famous Beijing-based creator of TikTok. It's been around since last February and it has been maturing. What makes TRAE appealing is that it's an IDE – Integrated Development Environment – centric solution. TRAE's documentation explains:

Trae IDE is your powerful AI-powered code editor from ByteDance, featuring Claude 3.5 and

GPT-4 and DeepSeek integration. It's designed to be your coding companion, offering AI-assisted features like code completion, intelligent suggestions, and agent-based programming capabilities. When developing with Trae IDE, you can collaborate with AI to boost your productivity. Trae IDE provides essential IDE functionality including code editing, project management, extension management, version control, and more. It supports seamless migration from VSCode and Cursor by importing your existing configurations.

During coding, you can engage in real-time conversations with the AI assistant for help, including code explanations, documentation generation, and error fixing. The interface is fully optimized for both English and Chinese users. The AI assistant understands your code context and provides intelligent code suggestions in real-time within the editor. Simply describe your requirements to the AI assistant in natural language, and it will generate appropriate code snippets or autonomously write project-level and cross-file code.

Tell the AI assistant what kind of program you want to develop, and it will provide relevant code or automatically create necessary files based on your description. With support for multiple programming languages and a rich plugin ecosystem, Trae IDE helps you build complete projects efficiently.

I want to give everyone a sense for what's happening in this segment of the world. So here's an independent review/posting made last May, three months after TRAE's release wrote:

Meet Trae AI: A Free AI Coding Agent With Model Context Protocol (MCP)

AI code assistants are flooding the market, but most still feel like chatbots taped onto an editor. Trae AI takes a different route. It ships an Integrated Development Environment (IDE) with a built-in agent framework that parses your entire codebase, talks to outside tools through the Model Context Protocol (MCP), and—crucially—costs nothing to install. If you are still paying for a $20 monthly subscription, Trae AI is an AI coding agent that offers "local-first" setup and a zero-dollar price tag, making it worth a test drive.

So what is Trae? Trae AI is a free AI coding agent with model context protocol (MCP) that offers itself as a collaborative partner for software engineers. It's designed to fit into a developer's existing coding environment, not as a replacement, but as an intelligent AI assistant. Trae Provides:

- *Budget relief: The main editor and completion model are free, removing the line-item that has kept many finance and ops leaders from green-lighting AI pair-programming pilots.*
- *Agentic workflow: Instead of a single, do-everything helper, Trae lets you spin up specialist agents—one for refactoring, another for writing tests, a third for documentation, with each AI agent getting its own prompt, tools, and guardrails.*
- *Enterprise-style data rules without enterprise pricing: Code stays on your machine; any files briefly sent for indexing are wiped after embeddings are created. Regional hosting (US, Singapore, Malaysia) keeps governance teams calm about residency.*

What does Trae AI bring to the table?

- *Working Together: Trae's development environment is built to work with existing developer setups. The goal is to improve how developers and AI can cooperate for better outcomes*

and faster project completion.

- *Direct AI Communication: Developers can talk to Trae using straightforward language and simple instructions, such as "@Agent #Context," and they can delegate work, facilitating a more interactive relationship between humans and AI.*

- *Custom AI Assistants: Trae offers a flexible system for setting up specialized AI agents. It comes with a standard agent called "Builder" for everyday tasks. Past that, developers can create their own group of AI helpers, each with specific tools, skills, and ways of working, so the AI can be adjusted to fit precise project requirements.*

- *Connecting to Other Tools: Trae can link up with different external applications; currently, it uses a system known as the Model Context Protocol (MCP), which allows its AI agents to gather information from outside resources to better complete the tasks they're given.*

- *Understanding Project Details: Trae gains a good grasp of a project's specifics by looking at code repositories, information from online searches, and documents provided by users. Developers can also set up custom rules to fine-tune the AI's behavior, making sure it handles tasks exactly as intended.*

- *Smart Code Suggestions: As developers type, Trae offers intelligent code completions as it can anticipate what the developer is trying to write and automatically fill in code segments, helping speed up the writing process.*

The idea is to make the interaction feel natural, allowing developers to assign tasks or ask for help using simple commands. This approach could fundamentally change team dynamics, making AI less of a tool and more of a team member.

Conclusion:

The arrival of free, capable AI coding agents like Trae AI isn't just another tech trend; it shows a maturing of AI into a practical aid for a highly skilled and often costly workforce. Its mix of free pricing, configurable agents, and tight privacy controls offers a low-risk way to explore agentic coding without rewriting procurement rules.

For CTOs and engineering managers, the math is straightforward: swap a paid copilot for a free, locally hosted agent system and redirect budget to GPU credits or headcount. While AI won't be replacing entire development teams anytime soon, tools that augment their abilities, especially free ones, are certainly worth trying.

If your roadmap includes AI-assisted development but your finance team keeps asking for ROI proof, Trae might be the simplest "yes" you can give for the entire quarter.

I don’t mean to suggest that this TRAE IDE-centric AI coding system is in any way super-special. Quite the contrary, in fact. The world is already being flooded with similar and similarly powerful AI-based solutions. I just wanted to share a sample of the tool that happened to be picked by the Chinese language speaker who created this particular “VoidLink” malware.

As is always the case for these sorts of things, my interest in sharing this on the podcast is in giving this event – and the news it brings – some context. Unfortunately, today, I truly fear that the news is worse than bad, and I’m unable to find a silver lining.

We're all familiar with the notion of asymmetric warfare, sometimes referred to as a guerilla war.

The use of malware to penetrate, infect, exfiltrate and encrypt an enterprise's resources is inherently asymmetric. One lone malicious hacker hiding somewhere, anywhere on the Internet – perhaps, literally in his mother's basement – is able to single handedly attack and significantly negatively impact the national economy of the United Kingdom with one well-placed attack on Jaguar Land Rover. That's the very definition of asymmetry.

The problem with this emergence of AI, and its expected application to the empowerment of all forms of coding, is that I believe history and the evidence suggests that the bad guys will be gaining a far greater advantage from their malicious application of AI to create malware than the good guys will be gaining through their use – to do what? It's not at all clear what the good guys can do that isn't already being done. In other words, I cannot see how the benefit from the application of AI to both sides is in any way even close to being symmetric. I believe that AI's value is extremely asymmetric, and that the asymmetric battle that's been waging for the past decade has just become far more asymmetric.

In years past we have observed that hacker talent encompasses a wide range, from the so-called "script kiddies" at the low end to the elite hackers at the high end. And we know that this also takes a pyramid shape, with a great many lower-end wannabe hackers at the bottom and a much more rarified few at the top of the pile.

Recently, we've seen that the followers of this podcast have already been employing AI to create successful solutions that they would have never been able to create otherwise. You, Leo, as a lifelong coder, could have written your newsfeed reader from scratch the old fashioned way. So for you Claude's AI served as a powerful accelerant. But we know from the testimony of our listeners that for the many of them, who were coding-adjacent but not coders, AI has now bridged that gap to allow **them** to create their own functioning tools that never existed before.

So what AI has already done is completely eliminate coder-wannabe script kiddies from the low-end by empowering them to author their own powerful malicious code. They no longer need to follow someone else's script. Any mischief they can think to get up to, an AI will happily manifest, in code, just for them. Consequently, we are almost certainly facing a forthcoming explosion in the volume and variety of malicious attack code.

I would like to be able to imagine some form of silver lining for the defenders in this asymmetric war. But I've been unable to come up with any. What we see is an epidemic of misconfiguration and lazy configuration, communication failures and finger pointing, lingering old designs and practices, and systems that remain online despite not having received any attention for years. We also see employees in positions of trust on internal enterprise networks being tricked into clicking malicious links and inviting malware inside the house. No form of fancy coding AI is going to fix **any** of those things. Every single one of those is a human-factors failure. We already know how to fix every one of those things, but we haven't cared enough to do so. And there's reason to believe that we're about to pay the piper.

A great many of the world's enterprises are sitting ducks and entire new generations of would-be hunters who have been using slingshots have all just been up-armed with advanced cyber-rifles.