

FILED

IN THE UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF VIRGINIA

Microsoft Corporation,

Plaintiff,

v.

Does 1-10 Operating an Azure Abuse
Network,

Defendants.

2024 DEC 19 PM 2:38

Civil Action No. 1:24-cv-2323

FILED UNDER SEAL

COMPLAINT

Plaintiff MICROSOFT CORP. (“Microsoft”) brings this action to protect itself, its customers, and the public from Defendants DOES 1-10’s (“Does” or “Defendants”) malicious scheme to misuse Microsoft systems and technology for improper and illegal purposes, including the unlawful generation of harmful images using Microsoft’s Azure OpenAI Service. By this action, Microsoft seeks to disrupt a sophisticated scheme carried out by cybercriminals who have developed tools specifically designed to bypass the safety guardrails of generative AI services provided by Microsoft and others.

NATURE OF ACTION

1. This action arises under the Computer Fraud and Abuse Act, 18 U.S.C. § 1030 (“CFAA”); the Digital Millennium Copyright Act, 17 U.S.C. 1201 *et seq.* (“DMCA”); the Lanham Act, 15 U.S.C. § 1125 *et seq.*; and the Racketeer Influenced and Corrupt Organizations Act, 18 U.S.C. § 1962(c). This action also involves Microsoft’s claims for trespass to chattels and tortious interference under Virginia state law. Microsoft seeks injunctive and other equitable relief and damages from Defendants for their creation, control, maintenance, trafficking, and ongoing use of illegal computer networks and piratical software to cause harm to Microsoft, its customers, and the public at large.

2. Although the technology implicated by this action is complex, the fundamental

allegations against Defendants are simple. Defendants used stolen customer credentials and custom-designed software to break into the computers running Microsoft's Azure OpenAI Service. Defendants then used Microsoft's computers and software for harmful purposes. Microsoft respectfully seeks the Court's assistance in putting a stop to Defendants' illegal conduct and holding Defendants to account for what they have done.

THE PARTIES

3. Plaintiff Microsoft Corp. is a corporation duly organized and existing under the laws of the State of Washington, having its headquarters and principal place of business in Redmond, Washington. Microsoft is a leading provider of technology products and services, including computer software, Internet services, websites, and email services.

4. Defendant DOE 1 is a natural person with access to and control over instrumentalities used in connection with the violations of law described in this Complaint, including at least the website located at "reentry.org/de3u," the source code repositories located at "github.com/notfiz/de3u," and stolen Azure API keys and other Microsoft customer authentication information. Based on the information it has been able to gather to date, Microsoft is informed and believes, and hereby alleges, that a reasonable opportunity for investigation or discovery will likely yield further evidentiary support showing that DOE 1 resides outside the United States.

5. Defendant DOE 2 is a natural person with access to instrumentalities used in connection with the violations of law described in this Complaint, including the reverse proxy tool located at "https://gitgud.io/khanon/oai-reverse-proxy." Based on the information it has been able to gather to date, Microsoft is informed and believes, and hereby alleges, that a reasonable opportunity for investigation or discovery will likely yield further evidentiary support showing that DOE 2 resides outside the United States.

6. Defendant DOE 3 is a natural person with access to and control over instrumentalities used in connection with the violations of law described in this Complaint, including reverse proxy tool infrastructure like the domain "aitism.net," the Cloudflare tunnel,

and the AWS IP Address used by Defendants for carrying out the scheme alleged in this Complaint. Based on the information it has been able to gather to date, Microsoft is informed and believes, and hereby alleges, that a reasonable opportunity for investigation or discovery will likely yield further evidentiary support showing that DOE 3 resides outside the United States.

7. Based on the information it has been able to gather to date, Microsoft is informed and believes, and hereby alleges, that a reasonable opportunity for investigation or discovery will likely yield further evidentiary support showing that DOES 4-10 are end-users of the illegal technology and services trafficked by DOES 1, 2, and 3. At least one of DOES 4-10 resides outside the United States. DOES 4-10 have each knowingly used infrastructure and technology provided by DOE 1, 2, and 3 to unlawfully access and use Microsoft's software and computers for the purpose of generating harmful content.

8. Defendants collectively operate and/or control infrastructure, software, and technical artifacts used to carry out the violations of law described in this Complaint. To summarize briefly: Defendants illegally procured authentication information from legitimate Microsoft customers with malicious intent, trafficked and used that stolen customer authentication information to bypass Microsoft authentication gates and gain unauthorized access to Microsoft software and computer systems, and then exploited their unauthorized access to Microsoft's software and computers to create harmful content in violation of Microsoft's policies and through circumvention of Microsoft's technical protective measures.

9. Microsoft is unaware and uncertain of the true names and capacities of Defendants sued herein as Does 1-10 inclusive and therefore sues these Defendants by such fictitious names. Microsoft will amend this complaint to allege Defendants' true names and capacities when ascertained through discovery of admissible evidence. Microsoft will exercise due diligence to determine Defendants' true names, capacities, and contact information, and to effect service upon those Defendants.

10. Each of the Defendants is responsible in some manner for the occurrences herein alleged. The injuries to Microsoft, its customers, and others herein alleged have been

proximately caused by such Defendants.

11. The actions alleged herein to have been undertaken by Defendants individually were actions that each Defendant authorized, controlled, directed, or had the ability to authorize, control or direct, and/or were actions each Defendant assisted, participated in, or otherwise encouraged, and are actions for which each Defendant benefited from and is liable. Each Defendant aided and abetted the actions of other Defendants set forth below. Each Defendant had knowledge of those actions, provided assistance, and benefited from those actions, in whole or in part. Each Defendant was the agent of each of the remaining Defendants, and in doing the things hereinafter alleged, was acting within the course and scope of such agency and with the permission and consent of other Defendants.

JURISDICTION AND VENUE

12. The Court has subject matter jurisdiction over this action pursuant to 28 U.S.C. § 1331 because this action arises out of Defendants' violation of the CFAA (18 U.S.C. § 1030), the DMCA (17 U.S.C. § 1201), the Lanham Act (15 U.S.C. § 1125(a), (c)), and the Racketeer Influence and Corrupt Organizations Act (18 U.S.C. § 1962(c)). The Court also has supplemental jurisdiction over Microsoft's state-law claims for trespass to chattels and tortious interference with contract pursuant to 28 U.S.C. § 1367.

13. In carrying out the conduct described in this Complaint, Defendants have availed themselves of the privilege of conducting business in Virginia and have directed acts complained of herein toward the state of Virginia and this judicial district. For example,

- a. Defendants chose a ".org" domain, "rentry.org/de3u," as an access point for the tools used to carry out the misconduct described in this Complaint. Since 1984, the ".org" top level domain ("TLD") has been managed by the Public Interest Registry ("PIR"), based in Reston, Virginia. This means that whenever Defendants or third parties access Defendants' tools through Defendants' "rentry.org/de3u" website, that access depends on PIR's physical domain name servers ("DNS") and DNS routing services. As sophisticated actors who have demonstrated substantial

knowledge of computer networks and the Internet, Defendants necessarily understood that by selecting a “.org” TLD for its website, Defendants would be relying on hardware and services provided by PIR from Reston, Virginia, to distribute the malicious de3u software within the United States in order to effect their scheme.

- b. Defendants chose the “.net” domain, “atism.net,” to act as a node in the network Defendants created to carry out the scheme described in this Complaint. Since 2000, the “.net” TLD has been managed by Verisign, Inc., based in Reston, Virginia. This means that whenever Defendants or third parties access Defendants’ tools through the reverse proxy network Defendants created, that access depends on Verisign’s physical DNS and DNS routing services. As sophisticated actors who have demonstrated substantial knowledge of computer networks and the Internet, Defendants necessarily understood that by selecting a “.net” TLD to support their network infrastructure, Defendants would be relying on hardware and services provided by Verisign from Reston, Virginia, to effect their scheme.
- c. Defendants chose the AWS IP Address, which geolocates to computers physically located in Virginia, as one end of the access tunnel they created into the Azure OpenAI Service in order to carry out their scheme. Defendants, sophisticated actors who have demonstrated substantial knowledge of computer networks, intentionally created software and systems so that images created by the Azure OpenAI Service were copied to a computer physically located in Virginia at the AWS IP Address.

14. In addition, Defendants have acted at all times relevant with knowledge that their acts would cause harm through computers located in Virginia thereby injuring Plaintiff, its customers, and others in Virginia and the United States.

15. Alternatively, to the extent Defendants are not subject to personal jurisdiction under Virginia’s long arm statute, Defendants also have availed themselves of the privilege of

doing business in the United States and have sufficient national contacts with the United States as a whole to subject each Defendant to the Court's jurisdiction consistent with the requirements of due process. For example, Defendants intentionally availed themselves of the privilege of doing business in the United States by:

- a. Intentionally configuring their software and systems to use physical machines, technology, and services provided in and from the United States by Microsoft (including, for example, Microsoft Azure servers, technology, and services);
- b. Intentionally configuring their software and systems to use physical machines, technology, and services provided in and from the United States by Amazon Web Services ("AWS"), a U.S. company;
- c. Intentionally configuring their software and systems to use physical machines, technology, and services provided in and from the United States by Cloudflare, a U.S. company;
- d. Intentionally configuring their software and systems to use physical machines, technology, and services provided in and from the United States by PIR;
- e. Intentionally configuring their software and systems to use physical machines, technology, and services provided in and from the United States by Verisign;
- f. Intentionally configuring their software and systems to victimize Microsoft, OpenAI, AWS, Cloudflare, PIR, and Verisign; and
- g. Intentionally configuring their software and systems to create and distribute harmful materials within the U.S.

16. Accordingly, to the extent Defendants do not have sufficient contacts with Virginia alone to support jurisdiction and venue in this Court, each Defendant is subject to national service of process under Federal Rule of Civil Procedure Rule 4(k)(2) and jurisdiction in this Court comports with due process given Defendants' national contacts with the United States.

17. Pursuant to 28 U.S.C. § 1391(b), venue is proper in this judicial district. A substantial part of the events that give rise to Plaintiffs' claims, and a substantial amount of the

infrastructure used to carry out Defendants' scheme, is situated in this judicial district. Venue is also proper in this judicial district under 28 U.S.C. § 1391(c) because Defendants are subject to personal jurisdiction in this judicial district.

FACTUAL BACKGROUND

Overview

18. Microsoft is the well-known creator and provider of the Windows operating system and a variety of other software and services. Microsoft has invested substantial resources in developing high-quality products and services. Due to the high quality and effectiveness of Microsoft's products and services and the expenditure of significant resources by Microsoft to market those products and services, Microsoft has generated substantial goodwill with its customers, establishing a strong brand and developing the Microsoft name and the names of its products and services into strong and famous world-wide symbols that are well-recognized within its channels of trade. Microsoft has registered trademarks representing the quality of its products and services and its brand, including the Microsoft®, Windows®, and Azure® marks.

19. Microsoft has long been a leader in the field of Artificial Intelligence ("AI"). AI refers to software that imitates human behaviors and capabilities. AI encompasses a wide range of workloads including:

- Machine Learning ("ML") – the way humans "teach" a computer model to make predictions and draw conclusions from data is often the foundation for an AI system.
- Computer vision – Capabilities within AI to interpret the world visually through cameras, video, and images.
- Natural language processing – Capabilities within AI for a computer to interpret written or spoken language and respond in kind.
- Document intelligence – Capabilities within AI that deal with managing, processing, and using high volumes of data found in forms and documents.
- Knowledge mining – Capabilities within AI to extract information from large volumes of often unstructured data to create a searchable knowledge store.


- **Generative AI – Capabilities within AI that create original content in a variety of formats including natural language, image, code, and more. Typically, generative AI applications take in natural language input (“prompts”) and return appropriate responses in a variety of formats including natural language, image, code, and audio.**

20. Since 2016, Microsoft has been committed to building its Azure platform into an AI supercomputer for the world. Commencing in 2019, Microsoft began a partnership with groundbreaking AI company OpenAI through a multiyear investment designed to ensure that AI’s benefits are broadly shared with the world. OpenAI has used Microsoft infrastructure to train its models, which are now deployed in Azure to power generative AI products like GitHub Copilot, DALL·E and ChatGPT. These innovations have captured imaginations and introduced large-scale AI as a powerful, general-purpose technology platform that Microsoft believes will create a transformative impact.

21. In addition to its technological leadership, Microsoft has also long been at the forefront of setting ethical standards for use of AI. Microsoft is committed to the ethical advancement of AI and has identified six principles it believes should guide AI development and use: (1) Fairness – AI systems should treat all people fairly; (2) Reliability and safety – AI systems should perform reliably and safely; (3) Privacy and security – AI systems should be secure and respect privacy; (4) Inclusiveness – AI systems should empower everyone and engage people; (5) Transparency – AI systems should be understandable; and (6) Accountability – People should be accountable for AI systems. Microsoft helps customers put responsible AI principles into action by providing educational resources and technical capabilities, as described further below.

22. Unfortunately, Microsoft’s status as a leader in AI also makes Microsoft a prime target for bad actors who wish to misuse Microsoft’s Azure technology for unlawful purposes. One type of malicious conduct that Microsoft has had to combat concerns malicious use of Microsoft’s Azure OpenAI Service and DALL·E image generation technology to create harmful content.

23. DALL·E is an AI system that can create realistic images and art from a description in natural language. To use an example provided on OpenAI’s website, a user can provide the text input: “a photorealistic image of an astronaut riding a horse” and DALL·E can return images that matches that text input, as depicted in the examples below.



24. According to OpenAI’s website, DALL·E began as a research project and has several built-in safety mitigations in place. For example, OpenAI states that it has limited the ability for DALL·E to generate violent, hate, or adult images by removing the most explicit content from DALL-E training data. OpenAI also used advanced techniques to prevent photorealistic generations of real individuals’ faces, including those of public figures. OpenAI also has a content policy that does not allow users to generate violent, adult, or political content, among other categories and says that OpenAI “won’t generate images if our filters identify text prompts and image uploads that may violate our policies. We also have automated and human monitoring systems to guard against misuse.”

25. As described in more detail below, Microsoft’s Azure-based integration of DALL·E adds further layers of safety and security. Yet despite Microsoft’s and OpenAI’s various safety mitigations, sophisticated bad actors have devised ways to obtain unlawful access to Microsoft’s systems, circumvent safety mitigations, and generate harmful content using

Microsoft systems. This lawsuit involves some of these groups of bad actors.

26. Using stolen API keys and technical circumvention measures, the Defendants in this case gained unauthorized access to the computers and software that provide Microsoft's Azure-based implementation of OpenAI's generative AI models ("Azure OpenAI Service") and used that unauthorized access to circumvent Microsoft's safety measures preventing generation and dissemination of harmful content. Defendants used Microsoft's Azure OpenAI Service to generate thousands of harmful images.

Microsoft Azure Infrastructure and Software

27. In 2008, Microsoft announced Azure as its new cloud computing operating system. Originally targeting businesses and developers, "Windows Azure" was built as an extension of Windows New Technology ("Windows NT") and marked the beginning of Microsoft's Cloud Platform as a Service offering. Windows Azure became commercially available in 2010 and Microsoft has steadily enhanced and expanded its Azure offering ever since. Microsoft eventually changed the name of its cloud offering to "Microsoft Azure" and later launched the Azure Data Lake Store and Azure Data Lake Analytics to provide an end-to-end Big Data and analytics platform on Azure.

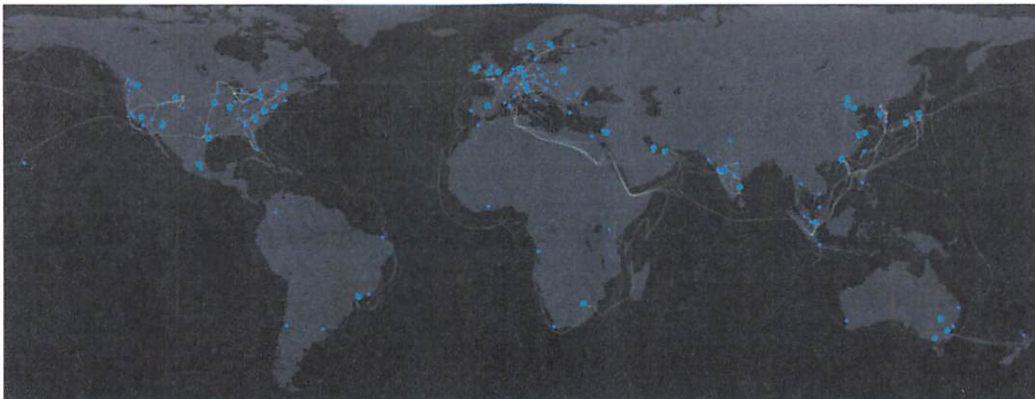
28. Azure enabled Microsoft to become a first mover in the Machine Learning and Artificial Intelligence domains. Microsoft invested heavily in database, Big Data, AI, and Internet of Things ("IoT") technology, and in creating services centered around machine learning and artificial intelligence. As a result, Azure quickly garnered a reputation as a leading platform for AI and rich cloud services operations. Azure currently offers hundreds of services, many of which are industry leading Software as a Service ("SaaS"), Platform as a Service ("PaaS"), and Infrastructure as a Service ("IaaS") solutions. Because of Microsoft's continuing innovation and investment, including over \$1 billion per year of investment in security to protect its customers' data from cyberthreats, many Fortune 500 companies entrust their business to Azure today.

29. Most Azure cloud services fall into four broad categories: infrastructure as a service, platform as a service, serverless, and software as a service. These are sometimes called

the cloud computing “stack” because they build on top of one another. The most basic category of cloud services is infrastructure as a service (“IaaS”), which allows customers to rent IT infrastructure like servers, virtual machines, storage, networks, and operating systems on a pay-as-you-go basis. Platform as a service (“PaaS”) refers to cloud computing services that supply an on-demand environment for developing, testing, delivering, and managing software applications. PaaS is designed to make it easier for developers to quickly create web or mobile apps without worrying about setting up or managing the underlying infrastructure of servers, storage, network, and databases needed for development. Software as a service (“SaaS”) is a method for delivering software applications over the Internet on demand and typically on a subscription basis. With SaaS, cloud providers host and manage the software application and underlying infrastructure and handle any maintenance, like software upgrades and security patching. Users connect to the application over the Internet, usually with a web browser on their phone, tablet, or PC. Microsoft Azure provides all of these types of services, both via its public cloud infrastructure and via hybrid cloud offerings (a type of cloud computing that combines on-premises infrastructure—or a private cloud—with a public cloud).

30. Azure’s global infrastructure includes access-controlled computers and networks that may only be accessed and used by authenticated customers. Azure datacenters and regional hubs are depicted below in Figure 1. In addition to the Microsoft-owned resources depicted in Figure 1, customers use their own computing resources and public Internet infrastructure to connect to and use Azure services.

Fig. 1: Azure Global Datacenters



31. Azure physical infrastructure is supported by proprietary Microsoft software responsible for enabling, *e.g.*, communications routing, system monitoring, load balancing, and security, providing database and user interface functionality.

32. For example, the Azure portal is a web-based, unified console that lets users create and manage their Azure resources. Users can use the Azure portal to build, manage, and monitor everything from simple web apps to complex cloud deployments. The Azure portal is designed for resiliency and continuous availability and has a presence in every Azure datacenter. The Azure portal updates continuously, requires no downtime for maintenance activities, and can be accessed through any supported browser. For each of the end-user facing elements of Azure described, there is associated back-end software authored by Microsoft that is required to provide the described functionality. Microsoft takes steps to protect its copyright interests in Azure software and grants customers licenses to use it under certain conditions.

The Azure OpenAI Service

33. The Azure OpenAI Service is Microsoft's cloud solution for deploying, customizing, and hosting generative AI models created by the groundbreaking company OpenAI. The Azure OpenAI Service provides access to many of OpenAI's cutting-edge models including various versions of OpenAI's GPT and DALL-E models. The Azure OpenAI Service brings together OpenAI's cutting edge models and APIs with the security and scalability of the Azure cloud platform. Microsoft experts in AI research, policy, and engineering collaborate to develop practical tools and methodologies that support AI security, privacy, safety, and quality and embed them directly into the Azure AI platform.

34. As part of Microsoft's commitment to responsible AI, Microsoft designed and operates Azure OpenAI Service with the intention of protecting the rights of individuals and society and fostering transparent human-computer interaction. For this reason, the Azure OpenAI Service is a Limited Access service, and access and use are subject to eligibility criteria determined by Microsoft.

35. There are several ways that Microsoft ensures responsible use of the Azure

OpenAI Service, including by offering tools to moderate generated content, by providing guidance for safely designing applications, by limiting access to certain models to ensure that responsible AI safeguards are working in practice, and by imposing contractual restrictions on access to the Azure OpenAI Service.

Contractual Restrictions on Use of the Azure OpenAI Service

36. The Azure OpenAI Service is made available to customers under the terms governing their subscription to Microsoft Azure Services, including Product Terms for Microsoft Azure Services.

37. A person or company that wishes to use Azure services must first create an Azure account and user profile. Azure users must provide accurate location, name, and contact information and must agree to the Microsoft customer agreement. Among other things, the Microsoft Customer Agreement states:

- a) **Licenses for Products.** Products are licensed and not sold. Upon Microsoft's acceptance of each order and subject to Customer's compliance with this Agreement, Microsoft grants Customer a nonexclusive and limited license to use the Products ordered as provided in this Agreement. These licenses are solely for Customer's own use and business purposes and are nontransferable except as expressly permitted under this Agreement or applicable law.
- b) **Duration of licenses.** Online Services and some Software are licensed on a subscription basis for a specified period of time. Subscriptions expire at the end of the applicable subscription period unless renewed. Some Subscriptions renew automatically until canceled. The Subscription term for Online Services that are billed in arrears based on usage is the same as the billing period unless otherwise specified in the Product Terms. Perpetual Software licenses become perpetual upon payment in full.
- c) **End Users.** Customer will control access to, and use of, the Products by End Users and is responsible for any use of the Products that does not comply with this Agreement.

38. The Microsoft Customer Agreement also includes a "restrictions" section that expressly prohibits several categories of conduct:

- f) **Restrictions.** Except as expressly permitted in this Agreement or Product documentation, Customer must not (and is not licensed to):
 - (i) reverse engineer, decompile, or disassemble any Product or Services Deliverable, or attempt to do so (except where applicable law permits despite

this limitation);

- (ii) install or use non-Microsoft software or technology in any way that would subject Microsoft's intellectual property or technology to any other license terms;
- (iii) work around any technical limitations in a Product or Services Deliverable or restrictions in Product documentation;
- (iv) separate and run parts of a Product or Services Deliverable on more than one device;
- (v) upgrade or downgrade parts of a Product at different times;
- (vi) transfer parts of a Product separately; or
- (vii) distribute, sublicense, rent, lease, or lend any Products or Services Deliverables, in whole or in part, or use them to offer hosting services to a third party.

39. Some Azure services are provided free of charge, but most require payment.

Microsoft provides predictable and transparent pricing models that let customers pay only for the cloud resources they use and scale as they grow.

40. After agreeing to the Microsoft Customer Agreement, users wishing to access and use Azure resources and services must authenticate themselves with valid Microsoft-provided user credentials. There are several ways users can authenticate themselves to gain access to Azure services. For example, users can authenticate themselves to Azure using Microsoft Entra ID, which is a cloud-based identity and access management service.

41. Microsoft Azure provides AI-optimized infrastructure that helps build and train some of the industry's most advanced AI solutions. Azure users can leverage supercomputing performance for some of the most complex generative AI models, with reliability at massive scale with on-demand sizes ranging from eight to thousands of virtual machines interconnected by state-of-the-art networking systems.

Conduct Guidelines and Policies

42. The Microsoft Generative AI Services Code of Conduct ("Code of Conduct") defines the requirements that all customers of Microsoft Generative AI Services must adhere to in good faith. The Code of Conduct requires customers to ensure that all of their applications

built with Microsoft Generative AI Services and Azure AI Content Safety:

- Implement meaningful human oversight.
- Implement technical and operational measures to detect fraudulent user behavior in account creation and during use.
- Implement strong technical limits on inputs and outputs to reduce the likelihood of misuse beyond the application's intended purpose.
- Disclose the synthetic nature of generated voices, images, and/or videos to users such that users are not likely to be deceived or duped – or able to prank others – into believing they are interacting with a real person or that any voice or other generated content is authentic or attributable to a specific individual.
- Test applications thoroughly to find and mitigate undesirable behaviors
- Establish feedback channels.
- Implement additional scenario-specific mitigations as appropriate.”

43. In addition, the Code of Conduct's Usage Restrictions provide that “Customers, users, and applications built with Microsoft Generative AI Services and Azure AI Content Safety must NOT use the services,” for example:

- In any way that is inconsistent with this Code of Conduct, including the Responsible AI mitigation requirements, the Content requirements, and any applicable Limited Access Requirements;
- To generate or interact with content prohibited in this Code of Conduct;
- To present content alongside or to monetize content prohibited in this Code of Conduct;
- To make decisions without appropriate human oversight as part of an application that may have a consequential impact on any individual's legal position, financial position, life opportunities, employment opportunities, or human rights, or may result in physical or psychological harm to an individual;
- To deploy subliminal techniques (e.g., visual, auditory, or other signals beyond a

normal person's range of perception) with the intent to deceive or cause harm;

- To deceive or intentionally misinform, for false advertising, or to manipulate or distort the behavior of a person in a way that causes harm;
- To exploit any of the vulnerabilities of a person (e.g., age, disability, or socio-economic situation);
- To infer people's sensitive attributes such as gender, race, nationality, religion, or specific age (not including age range, position of mouth (e.g., smile or frown), and hair color);
- Except for customers approved for modified content filtering, to identify or verify individual identities based on people's faces, voices, or other physical, physiological, or behavioral characteristics;
- For unlawful tracking, stalking, or harassment of a person;
- To generate content with the purpose of removing or altering content credentials or other provenance methods, marks, or signals ("AI Content Credentials") that indicate that the content was generated by a Microsoft Generative AI Service;
- To generate content with the purpose of misleading others about whether the content was generated by a Microsoft Generative AI Service; or
- To detect AI Content Credentials with the purpose of removing or altering them.

44. The Code of Conduct also imposes Content Requirements that prohibit the use of Microsoft Generative AI Services for processing, generating, classifying, or filtering content in ways that can inflict harm on individuals or society. These content requirements apply to use of features of, and the output of, all Microsoft Generative AI Services and Azure AI Content Safety. This includes, but is not limited to, use of features of Azure OpenAI Service and all content provided as input to or generated as output from all models available in Azure OpenAI Service, such as GPT and DALL·E. These requirements apply to the use of Azure AI Content Safety, including features such as customized categories, and to all content provided as input to the service and content generated as output from the service regardless of content filter settings.

45. Microsoft prohibits content that describes, features, or promotes sexual exploitation or abuse, whether or not prohibited by law. Microsoft further prohibits the creation of erotic, pornographic, or otherwise sexually explicit content. This includes sexually suggestive content, depictions of sexual activity, and fetish content. Microsoft prohibits content that attacks, denigrates, intimidates, degrades, targets, or excludes individuals or groups on the basis of traits such as actual or perceived race, ethnicity, national origin, gender, gender identity, sexual orientation, religious affiliation, age, disability status, caste, or any other characteristic that is associated with systemic prejudice or marginalization. Microsoft prohibits content that targets individual(s) or group(s) with threats, intimidation, insults, degrading or demeaning language or images, promotion of physical harm, or other abusive behavior such as stalking.

46. Microsoft prohibits content that is intentionally deceptive and likely to adversely affect the public interest, including deceptive or untrue content relating to health, safety, election integrity, or civic participation. Microsoft also prohibits inauthentic interactions, such as fake accounts, automated inauthentic activity, impersonation to gain unauthorized information or privileges, and claims to be from any person, company, government body, or entity without explicit permission to make that representation.

47. Microsoft provides further guidance in the Transparency Note: Azure AI Content Safety (“Transparency Note”). The Transparency Note is intended to help users understand how Microsoft’s AI technology works, the choices system owners can make that influence system performance and behavior, and the importance of thinking about the whole system, including the technology, the people, and the environment. Among other things, the Transparency Note advises users to:

- Avoid open-ended, unconstrained content generation.
- Avoid scenarios where use or misuse of the system could result in significant physical or psychological injury to an individual.
- Avoid scenarios where use or misuse of the system could have a consequential impact on life opportunities or legal status.

- Avoid high-stakes scenarios that could lead to harm.
- Carefully consider all generative use cases, because some content generation scenarios may be more likely to produce unintended outputs and these scenarios require careful consideration and mitigations.

48. The Transparency Note also cautions users that large-scale natural language, image, and speech models trained with such data can potentially behave in ways that are unfair, unreliable, or offensive, in turn causing harms.

Technical Measures Protecting the Azure OpenAI Service

49. In addition to the contractual restrictions and the guardrails imposed by the Code of Conduct, the Transparency Note, and Microsoft's AI principles, Microsoft has also developed technical measures controlling access to and enhancing the safety of the Azure OpenAI Service.

50. Microsoft technical measures for protecting the safety of the Azure OpenAI Service include Microsoft's content filtering and abuse detection technologies. Within the Azure OpenAI Service, the OpenAI models are integrated with Microsoft-developed content filtering and abuse detection models. For example, Azure OpenAI Service includes a content filtering system that works alongside core models, including DALL·E image generation models. This system works by running both the prompt and completion through an ensemble of classification models designed to detect and prevent the output of harmful content. The content filtering system detects and takes action on specific categories of potentially harmful content in both input prompts and output completions. The text content filtering models for the hate, sexual, violence, and self-harm categories have been specifically trained and tested on the following languages: English, German, Japanese, Spanish, French, Italian, Portuguese, and Chinese. However, the service can work in many other languages.

51. The content filtering system integrated in the Azure OpenAI Service contains Neural multi-class classification models aimed at detecting and filtering harmful content; the models cover four categories (hate, sexual, violence, and self-harm) across four severity levels (safe, low, medium, and high). Other optional classification models are aimed at detecting

jailbreak risk and known content for text and code; these models are binary classifiers that flag whether user or model behavior qualifies as a jailbreak attack or match to known text or source code. The use of these models is optional, but use of protected material code model may be required for Customer Copyright Commitment coverage.

52. The Azure OpenAI Service includes default safety applied to all models, with some exceptions not relevant here. These configurations provide customers with a responsible experience by default, including content filtering models, blocklists, prompt transformation, content credentials, and others. For example, Azure OpenAI DALL·E also comes with prompt transformation by default. This transformation occurs on all prompts to enhance the safety of an original prompt, specifically in the risk categories of diversity, deceptive generation of political candidates, depictions of public figures, protected material, and others.

53. In the default streaming scenario, completion content is buffered, the content filtering system runs on the buffered content, and – depending on the content filtering configuration – content is either returned to the user if it doesn't violate the content filtering policy (Microsoft's default or a custom user configuration), or it's immediately blocked and returns a content filtering error, without returning the harmful completion content. This process is repeated until the end of the stream. Content is fully vetted according to the content filtering policy before it's returned to the user.

54. Customers can also configure content filters and create custom safety policies that are tailored to their use case requirements. The configurability feature allows customers to adjust the settings, separately for prompts and completions, to filter content for each content category at different severity levels. For example, customers can choose the Asynchronous Filter as an extra option, in which case content filters are run asynchronously, and completion content is returned immediately. Because content filters are run asynchronously, content moderation messages and policy violation signals are delayed, however. While customers retain discretion on certain content filtering configurations, importantly, for Azure OpenAI models, only trusted customers

who have been approved for modified content filtering have full content filtering control.¹

55. In addition to the content filtering system, Azure OpenAI Service performs Abuse Monitoring to detect content and/or behaviors that suggest use of the service in a manner that might violate applicable product terms. Azure OpenAI Service detects and mitigates instances of recurring content and/or behaviors that suggest use of the service in a manner that may violate the Code of Conduct or other applicable product terms. There are several components to Abuse Monitoring:

- **Content Classification:** Classifier models detect harmful language and/or images in user prompts (inputs) and completions (outputs). The system looks for categories of harms as defined in the Content Requirements and assigns severity levels as described in more detail on the Content Filtering page.
- **Abuse Pattern Capture:** Azure OpenAI Service's abuse monitoring looks at customer usage patterns and employs algorithms and heuristics to detect indicators of potential abuse. Detected patterns consider, for example, the frequency and severity at which harmful content is detected in a customer's prompts and completions.
- **Human Review and Decision:** When prompts and/or completions are flagged through content classification and abuse pattern capture as described above, authorized Microsoft employees may assess the flagged content, and either confirm or correct the classification or determination based on predefined guidelines and policies. Data can be accessed for human review only by authorized Microsoft employees via Secure Access Workstations (SAWs) with Just-In-Time (JIT) request approval granted by team managers. For Azure OpenAI Service resources deployed in the European Economic

¹ Some customers may want to use the Azure OpenAI Service for a use case that involves the processing of sensitive, highly confidential, or legally-regulated input data but where the likelihood of harmful outputs and/or misuse is low. These customers may conclude that they do not want or do not have the right to permit Microsoft to process such data for abuse detection, as described above, due to their internal policies or applicable legal regulations. To address these concerns, Microsoft allows customers who meet additional Limited Access criteria and attest to specific use cases to apply to modify Azure OpenAI abuse monitoring features.

Area, the authorized Microsoft employees are located in the European Economic Area.

- **Notification and Action:** When a threshold of abusive behavior has been confirmed based on the preceding three steps, the customer is informed of the determination by email. Except in cases of severe or recurring abuse, customers typically are given an opportunity to explain or remediate—and implement mechanisms to prevent recurrence of—the abusive behavior. Failure to address the behavior—or recurring or severe abuse—may result in suspension or termination of the customer’s access to Azure OpenAI resources and/or capabilities.

56. To detect and mitigate abuse, Azure OpenAI stores all prompts and generated content securely for up to thirty (30) days (no prompts or completions are stored if the customer is approved for and elects to configure abuse monitoring off). The data store where prompts and completions are stored is logically separated by customer resource (each request includes the resource ID of the customer’s Azure OpenAI resource). A separate data store is located in each geography in which the Azure OpenAI Service is available, and a customer’s prompts and generated content are stored in the Azure Geography where the customer’s Azure OpenAI service resource is deployed, within the Azure OpenAI service boundary. Human reviewers assessing potential abuse can access prompts and completions data only when that data has been flagged by the abuse monitoring system. The human reviewers are authorized Microsoft employees who access the data via point wise queries using request IDs, Secure Access Workstations (SAWs), and Just-In-Time (JIT) request approval granted by team managers.

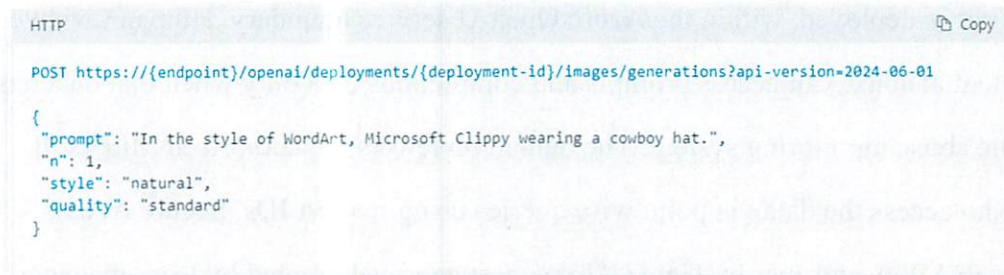
Azure OpenAI Service APIs

57. Microsoft provides access to the Azure OpenAI Service through application programming interfaces, also known as APIs. An API is computer code that enables software applications to communicate with each other. Microsoft’s Azure OpenAI Service APIs and the Microsoft software those APIs are designed to call are original, creative works of authorship and are the product of substantial creative decision making, as well as substantial investment and effort.

58. The Azure OpenAI Service APIs are divided into three categories. First, Microsoft's Azure control plane API is used for things like creating Azure OpenAI resources, model deployment, and other higher level resource management tasks. Azure OpenAI shares software and a common control plane with all other Azure AI Services. Second, Microsoft's data plane authoring API controls software that provides fine-tuning, file-upload, ingestion jobs, batch and certain model level queries. Third, Microsoft's data plane inference API accesses Microsoft software that provides the inference capabilities/endpoints for features like completions, chat completions, embeddings, speech/whisper, and DALL·E.

59. Customers who have entered into the necessary contractual agreements with Microsoft may use Microsoft's APIs to access the Azure OpenAI Service via the Internet using the http protocol.² For instance, the example code in Figure 2 below depicts an API call to the Azure OpenAI Service that requests DALL·E to generate an image of Microsoft Clippy wearing a cowboy hat:

Fig. 2



```

HTTP
Copy
POST https://{endpoint}/openai/deployments/{deployment-id}/images/generations?api-version=2024-06-01

{
  "prompt": "In the style of WordArt, Microsoft Clippy wearing a cowboy hat.",
  "n": 1,
  "style": "natural",
  "quality": "standard"
}

```

60. The API-version field tells Microsoft's system what version of the API the customer is using. The "prompt" field is the text description of the desired image, "n" is the number of images requested, "style" refers to the image style requested, and quality refers to the

² HTTP stands for Hypertext Transfer Protocol, and it's how different parts of the Internet communicate with each other. HTTP is what's known as a "request-response" language because your web browser (Firefox, Safari, etc.) sends an HTTP request to a server that is hosting the web content you want to work with. The server then sends an HTTP response back to your web browser. <https://learn.microsoft.com/en-us/xandr/industry-reference/intro-to-http>

image resolution (e.g., standard or high definition). Only by communicating in the specific format required by Microsoft's API can a customer access the functionality provided by the Azure OpenAI Service API.

61. In response to the API call in Figure 1 above, because there is no prohibited content or abuse detected, the Azure OpenAI Service returns the response depicted in Figure 3 below.

Fig. 3

```
{
  "body": {
    "created": 1698342300,
    "data": [
      {
        "revised_prompt": "A vivid, natural representation of Microsoft Clippy wearing a cowboy hat.",
        "prompt_filter_results": {
          "sexual": {
            "severity": "safe",
            "filtered": false
          },
          "violence": {
            "severity": "safe",
            "filtered": false
          },
          "hate": {
            "severity": "safe",
            "filtered": false
          },
          "self_harm": {
            "severity": "safe",
            "filtered": false
          },
          "profanity": {
            "detected": false,
            "filtered": false
          }
        },
        "url": "https://dalletipusw2.blob.core.windows.net/private/images/e5451cc6-b1ad-4747-bd46-b89a3a3b8b",
        "content_filter_results": {
          "sexual": {
            "severity": "safe",
            "filtered": false
          },
          "violence": {
            "severity": "safe",
            "filtered": false
          },
          "hate": {
            "severity": "safe",
            "filtered": false
          },
          "self_harm": {
            "severity": "safe",
            "filtered": false
          }
        }
      }
    ]
  }
}
```

62. In the response depicted above in Figure 3, the "revised_prompt" field indicates the prompt used by the Azure OpenAI Service to generate the image, and the "url" field is the

uniform resource locator, e.g., the internet address, of the image generated by the Azure OpenAI Service.

63. In the example code above, there is no content filtering called for, so an image is successfully generated and returned to the URL specified in the url field. By contrast, when the Azure OpenAI Service content filtering system detects harmful content, a user receives either an error on the API call if the prompt was deemed inappropriate, or the `finish_reason` on the response will be `content_filter` to signify that some of the completion was filtered.

64. In order to utilize Microsoft APIs to generate an image using DALL·E as described above, users must first authenticate themselves to gain access to the Azure OpenAI Service. The Azure OpenAI Service provides two methods for authentication:

- **API Key authentication:** For this type of authentication, all API requests must include the API Key in the `api-key` HTTP header.
- **Microsoft Entra ID authentication:** Customers can authenticate an API call using a Microsoft Entra token. Authentication tokens are included in a request as the `Authorization` header.

As discussed above, this case involves Defendants' illegal theft, trafficking, and use of stolen API keys.

65. An API key is a unique string composed of 52 randomly generated numbers and letters. API keys are used for data plane (content) requests and may be viewed and managed in the customer's Azure Portal. Key-based authentication is the default type of authentication for most Azure services. For this type of authentication, all API requests must include a valid API key in the `api-key` HTTP header.

66. By design, API keys are difficult to re-create and provide a significant measure of security. However, like any lock-and-key system, API key security is only effective if the key itself is kept secure. For this reason, Microsoft advises its users to adhere to certain best practice regarding API key use and maintenance. For example, Microsoft advises users:

- "Only use API keys if data disclosure isn't a risk (for example, when using

sample data) and if you're operating behind a firewall. Exposure of API keys is a risk to both data and unauthorized use of your search service.”

- “Always check code, samples, and training material before publishing to make sure you didn't leave valid API keys behind.”
- “For production workloads, switch to Microsoft Entra ID and role-based access. Or, if you want to continue using API keys, be sure to always monitor who has access to your API keys and regenerate API keys on a regular cadence.”

67. API keys can be accidentally exposed, for example, when keys are hardcoded into source code that is maintained in publicly accessible source code repositories. Bad actors have been known to create scraping tools designed specifically to search for API keys, and these tools can be applied in any code repository that the bad actor is able to access in order to steal API Keys for malicious purposes. Even when API Keys are maintained in secure environments, they are susceptible to theft by persons gaining unauthorized access to those environments, including during data breaches or the like.³ For this reason, Microsoft counsels against storing API Keys in unencrypted form.

Defendants' Unlawful Access to and Use of the Azure OpenAI Service

68. In late July 2024, Microsoft discovered use of customer API Keys to generate prohibited content. Investigation revealed that the API Keys had been stolen. The precise manner in which Defendants obtained all of the API Keys used to carry out the misconduct described in this Complaint is unknown, but it appears that Defendants have engaged in a pattern of systematic API Key theft that enabled them to steal Microsoft API Keys from multiple Microsoft customers. Multiple customers from whom Defendants stole API Keys are U.S. companies, including companies located in Pennsylvania and New Jersey.

69. Using stolen Microsoft API Keys that belonged to U.S.-based Microsoft customers, Defendants created a hacking-as-a-service scheme—accessible via infrastructure like

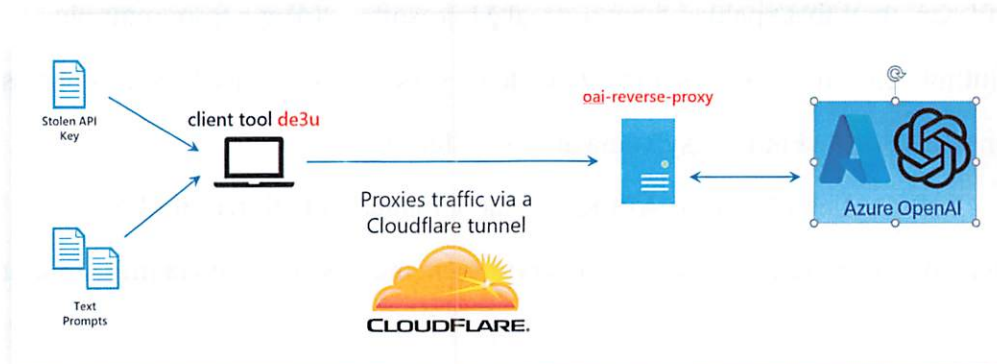
³ Microsoft's investigation to date has uncovered evidence that at least one Defendant has a history of activity on a known data breach information trafficking website.

the “rentry.org/de3u” and “aitism.net” domains—specifically designed to abuse Microsoft’s Azure infrastructure and software.

70. Using de3u and a custom-built reverse proxy service to generate and communicate HTTP requests that included Azure OpenAI Service API calls configured with stolen API Keys, deployment ID, endpoint address and other information configured by the de3u software and oai reverse proxy, Defendants gained unauthorized access to and use of Microsoft computers running Azure OpenAI services software necessary for processing, routing, filtering, executing, and communicating responses to Azure OpenAI Service API calls. Defendants could not have achieved the level of access they achieved without configuring their HTTP requests in a manner designed to circumvent Microsoft’s technological measures for limiting access to and use of the computers and software that comprise the Azure OpenAI Service.

71. Defendants’ malicious service can be described as two related software tools and associated Internet infrastructure used to unlawfully generate images through the Azure OpenAI Service. First, Defendants created a client-side software tool referred to by Defendants as “de3u,” which Defendants make publicly available via the “rentry.org/de3u” domain. Second, Defendants created software for running a reverse proxy service, referred to as the “oai reverse proxy,” designed specifically for processing and routing communications from the de3u software to Microsoft’s systems. Figure 4 below depicts the basic architecture of Defendants’ malicious hacking-as-a-service infrastructure.

Fig. 4

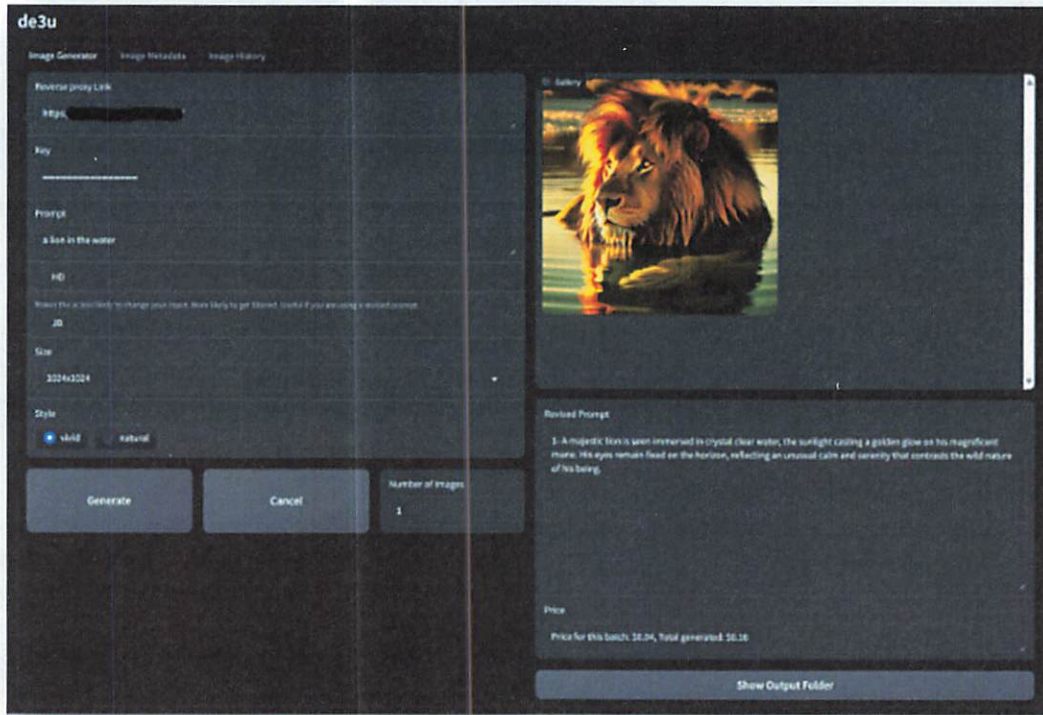


72. **The de3u Software.** At a high level, the de3u software allows users to issue Microsoft API calls to generate images using the DALL-E model through a simple user interface that leverages the Azure APIs to access the Azure OpenAI Service. Using an open-source software package, Defendants built a web application that implements a custom layout and data flow designed specifically for using tools like DALL-E to generate images using text prompts. Defendants' de3u application communicates with Azure computers using undocumented Microsoft network APIs to send requests designed to mimic legitimate Azure OpenAPI Service API requests. These requests are authenticated using stolen API keys and other authenticating information.

73. Defendants' de3u software permits users to circumvent technological controls that prevent alteration of certain Azure OpenAPI Service API request parameters. For example, Microsoft's system is designed so that content generated using a given Microsoft customer's unique API key is only delivered to the endpoint address specified by that customer. Defendants' de3u software and the associated oai reverse proxy service permit Defendants to effectively alter the target endpoint associated with a customer's API key so that the API key is delivered to the de3u user's desired endpoint address.


74. Defendants' de3u software allows simple mapping of the control fields to input and output parameters so that less sophisticated bad actors can leverage stolen API keys without having to write their own code. Figure 5 below is a screen capture of the de3u user interface Defendants created:

Fig. 5



75. Defendants also designed their de3u software to be shared with third parties without the need for a hosting a web server. In a web server configuration, users access software that is running on a computer connected to the internet, as depicted in Figure 6 below.

Fig. 6



76. Defendants' system avoids the need for a web server, relying instead on the "reentry.org/de3u" domain to provide users with access to the code necessary to run the de3u software. This allows Defendants to provide access to their de3u tool—and by extension, access

to the Azure OpenAI Service—to anyone in the world who visits the “reentry.org/de3u” domain.

77. Defendants’ de3u software is designed to try to prevent the Azure OpenAI Service from revising the original text prompt used to generate images, which can happen for example when a text prompt contains words that trigger Microsoft’s content filtering. In addition, Defendants’ de3u software is designed to detect and report whether the Azure OpenAI Service rejected a text prompt because it is considered as violating Microsoft’s content policy. These features, combined with Defendants’ unlawful programmatic API access to the Azure OpenAI service, enabled Defendants to reverse engineer means of circumventing Microsoft’s content and abuse measures.

78. **The “oai” Reverse Proxy Service.** Defendants have implemented and used an “oai” reverse proxy service through which de3u users can access the Azure OpenAI Service. Defendant’s oai reverse proxy service consists of software running on a reverse proxy server that passes communications from de3u user computers through a Cloudflare tunnel⁴ into the Azure OpenAI Service.

79. In general, a reverse proxy server is a server that sits in front of web servers and forwards client (e.g., web browser) requests to those web servers. A reverse proxy ensures that no client ever communicates directly with that origin server.

80. In addition to performing the traditional function of any reverse proxy (e.g., forwarding requests), Defendants’ oai reverse proxy tool processes and alters communications traffic between de3u client computers and the target Azure OpenAI Service. Defendants specifically configured the oai reverse proxy to route traffic to a list of Azure OpenAI Service endpoints.

81. When a de3u user sends a request to the Azure OpenAI Service to generate an

⁴ According to Cloudflare’s public documentation, a “Cloudflare Tunnel provides you with a secure way to connect your resources to Cloudflare without a publicly routable IP address. With Tunnel, you do not send traffic to an external IP — instead, a lightweight daemon in your infrastructure (cloudflared) creates outbound-only connections to Cloudflare’s global network. Cloudflare Tunnel can connect HTTP web servers, SSH servers, remote desktops, and other protocols safely to Cloudflare. This way, your origins can serve traffic through Cloudflare without being vulnerable to attacks that bypass Cloudflare.” <https://developers.cloudflare.com/cloudflare-one/connections/connect-networks/>

image, the de3u software routes the request to the oai reverse proxy address. The oai reverse proxy tool parses the request and forwards it to the Azure OpenAI Service target endpoint through the Cloudflare tunnel.

82. The oai reverse proxy tool also receives and processes responses from the Azure OpenAI service before forwarding responses and other data to the de3u user device. If the de3u user's prompt resulted in generation of an image by the Azure OpenAI service, then the oai reverse proxy tool receives image parameters from the Azure OpenAI service including the URL of the generated image, and the prompt used to generate the image. If no image was generated, the oai reverse proxy tool receives and logs the results of any content filtering.

83. If the de3u user's prompt resulted in generation of an image by the Azure OpenAI Service, then the oai reverse proxy tool retrieves the image from the URL specified in the Azure OpenAI Service return response and saves the image to the computer at the AWS IP Address. The oai proxy service then performs several additional steps including injecting proxy information into the response traffic, setting some HTTP headers, logging events and text prompts, and sending the traffic back to the requesting de3u user client computer.

84. The images saved to the oai proxy server include a C2PA Content Credentials symbol ("CR Icon") inserted by the Azure OpenAI service.⁵ This CR Icon identifies the Azure OpenAI Service as the technology used to generate the image via a metadata field that contains the Microsoft® registered trademark.

85. Although the original image transmitted from the oai reverse proxy to the AWS IP address includes original metadata and the C2PA Content Credentials information, Defendants' system also provides functionality to strip original image metadata and replace it with de3u custom metadata.

⁵ Microsoft is a founding member of the Coalition for Content Provenance and Authenticity ("C2PA"). C2PA addresses the prevalence of misleading information online through the development of technical standards for certifying the source and history (or provenance) of media content. It publishes open source standard that allows for insertion of a digital watermark

CLAIMS FOR RELIEF
FIRST CLAIM FOR RELIEF

Violation of the Computer Fraud & Abuse Act, 18 U.S.C. § 1030

86. Microsoft realleges and incorporates by this reference each and every allegation set forth in the paragraphs above.

87. The computers that provide the Azure OpenAI Service are “protected computers” for purposes of 18 USC § 1030(e)(2).

88. The computers that provide the Azure OpenAI Service are not generally open for access by any user of the internet. Instead, users must authenticate themselves as licensed customers of Microsoft using unique customer identification information in order to gain access to the computers that provide the Azure OpenAI Service.

89. Defendants knowingly and intentionally accessed the Azure OpenAI Service protected computers without authorization, and as a result of such conduct caused damage and loss (18 U.S.C. § 1030(a)(5)(C)). Microsoft has devoted substantial economic and human resources to investigating and remediating Defendants’ conduct.

90. Defendants’ conduct has caused a loss to Microsoft during a one-year period aggregating at least \$5,000. Microsoft’s internal personnel and outside counsel have spent months investigating and working to remediate Defendants’ conduct, which has imposed costs well over the CFAA’s \$5,000 threshold. In addition, the value of the services Defendants fraudulently obtained from Microsoft exceeds \$5,000 dollars.

91. Plaintiff Microsoft seeks injunctive relief and compensatory and punitive damages under 18 U.S.C. §1030(g) in an amount to be proven at trial.

92. As a direct result of Defendants’ actions, Microsoft has suffered and continues to suffer irreparable harm for which there is no adequate remedy at law, and which will continue unless Defendants’ actions are enjoined.

SECOND CLAIM FOR RELIEF

Violation of the Digital Millenium Copyright Act – 17 U.S.C. § 1201(a)(1)

93. Microsoft realleges and incorporates by this reference each and every allegation set forth in the paragraphs above.

94. Microsoft's Azure APIs, the software the Azure APIs interact with, and the software that implements Microsoft's abuse and content filtering policies are creative works of authorship subject to protection under the Copyright Act. For example, the Azure middleware responsible for processing, routing, filtering, executing, and communicating Azure communications is subject to copyright protection. In addition, the Azure OpenAI Service Software that implements Microsoft's abuse and content filtering policies comprises a collection of creative models authored by Microsoft based on, for example, a multitude of nuanced artistic and safety considerations. Microsoft includes copyright headers in the source code for this software.

95. Microsoft controls access to and use of its copyright protected Azure software, including the Azure software responsible for processing, routing, filtering, executing, and communicating in response to Azure OpenAI Service API calls, through use of authentication information that includes API Keys, customer deployment IDs, endpoint information, and token information. Microsoft's API key management system effectively controls access to Microsoft's copyright protected Azure software.

96. In the ordinary course of its operation, Microsoft's API key management system requires application, with Microsoft's authority, of authenticating information to gain access to Microsoft's Azure software.

97. Using HTTP requests containing stolen and dynamically manipulated API Key, deployment ID, end point, and token information, Defendants sent to the Azure OpenAI Service computer commands that mimicked authentic Azure OpenAI Service API calls. These maliciously configured HTTP requests allowed Defendants to circumvent Microsoft's technical measures for controlling access to its Azure software. In addition, Defendants used technical

means to circumvent the normal operation of Microsoft's content filtering systems, gaining unlawful access to portions of Microsoft's Azure software that would have otherwise been denied.

98. Defendants' violations of the DMCA are willful because they were carried out with knowledge of wrongdoing.

99. Defendants' conduct caused Microsoft actual damages in an amount subject to determination at trial.

100. As an alternative to actual damages, Microsoft is entitled to statutory damages in the \$200 to \$2,500 range per act of circumvention for at least 2,500 instances of willful circumvention by Defendants.

101. As a direct result of Defendants' actions, Microsoft has suffered and continues to suffer irreparable harm for which they have no adequate remedy at law, and which will continue unless Defendants' actions are enjoined.

THIRD CLAIM FOR RELIEF

Violation of the Digital Millenium Copyright Act – 17 U.S.C. § 1201(a)(2)

102. Microsoft realleges and incorporates by this reference each and every allegation set forth in the paragraphs above.

103. Does 1-3 have manufactured, imported, offered to the public, provided, and otherwise trafficked in technology and services primarily designed and produced for the purpose of circumventing Microsoft's technological measures for effectively controlling access to works protected under the Copyright Act, including Microsoft's Azure software.

104. Defendants' violations of the DMCA are willful because they were carried out with knowledge of wrongdoing.

105. Defendants' conduct caused Microsoft actual damages in an amount subject to determination at trial.

106. As an alternative to actual damages, Microsoft is entitled to statutory damages in the \$200 to \$2,500 range per act of circumvention for at least 2,500 instances of willful

trafficking by Defendants.

107. As a direct result of Defendants' actions, Microsoft has suffered and continues to suffer irreparable harm for which it has no adequate remedy at law, and which will continue unless Defendants' actions are enjoined.

FOURTH CLAIM FOR RELIEF

False Designation of Origin Under the Lanham Act

108. Microsoft realleges and incorporates by this reference each and every allegation set forth in the paragraphs above.

109. Defendants have generated and distributed unauthorized images containing the Microsoft® mark in metadata identifying Microsoft as the source of such images.

110. Defendants' unauthorized images containing metadata identifying "Microsoft" as the source of such images is likely to cause confusion, or to cause mistake, or to deceive as to the affiliation, connection, or association of Defendants and Microsoft, or of Microsoft's sponsorship, or approval of Defendants' goods, services, or commercial activities.

111. Defendants have distributed within and imported into the United States images containing metadata identifying Microsoft as the source of such images, which is likely to cause confusion, or to cause mistake, or to deceive as to the affiliation, connection, or association of Defendants and Microsoft, or of Microsoft's sponsorship, or approval of Defendants' goods, services, or commercial activities.

112. The Microsoft® mark is famous, distinctive, and widely recognized by the general consuming public of the United States as a designation of the source of goods or services.

113. Defendants' conduct harms Microsoft's reputation and is likely to dilute by tarnishment Microsoft's famous mark.

114. Microsoft is entitled to actual damages in an amount to be proven at trial.

115. As a direct result of Defendants' actions, Microsoft has suffered and continues to suffer irreparable harm for which it has no adequate remedy at law, and which will continue

unless Defendants' actions are enjoined.

FIFTH CLAIM FOR RELIEF

**Violations of the Racketeer Influenced and
Corrupt Organizations Act (RICO) – 18 U.S.C. § 1964(c)**

116. Microsoft realleges and incorporates by this reference each and every allegation set forth in the paragraphs above.

117. Defendants are members of an ongoing association-in-fact enterprise (the "Azure Abuse Enterprise" or "Enterprise") consisting of Does 1-3, who provide hacking-as-a-service software and infrastructure, and Does 4-10, end users who together with Does 1-3 have used de3u, the oai reverse proxy, stolen API Keys, maliciously configured HTTP commands, and other instrumentalities described herein to commit wire fraud and access device fraud in violation of federal law.

118. The Azure Abuse Enterprise's members function as a continuing unit for the common purpose of achieving the objectives of the Enterprise, including the common objectives of wire fraud and access device fraud.

119. Defendants have conducted the affairs of the Azure Abuse Enterprise through a coordinated and continuous pattern of illegal activity in order to achieve their common unlawful purposes.

120. Defendants' pattern of illegal activity is not limited to attacks on Microsoft. Evidence Microsoft has uncovered to date indicates that the Azure Abuse Enterprise has been targeting and victimizing other AI service providers.

121. Microsoft alleges that a reasonable opportunity for discovery will yield evidence that Defendants' pattern of wire fraud and access device fraud predates and postdates the conduct described herein.

122. Through their scheme, Defendants unlawfully accessed Microsoft customer accounts with generative AI service entitlements and purposely altered the capabilities of those services through technical means.

123. Does 1-3 each provided funding, devices, infrastructure, resources, and logistical support needed to conduct the Azure Abuse Enterprise.

124. Does 1-3 sold the Azure Abuse Enterprise's technological capabilities to other malicious actors and provided those other actors with detailed instructions on how to use the Azure Abuse Enterprise's custom tools to generate harmful content.

125. Does 4-10 each provided resources, devices, and prompt engineering needed to conduct the Azure Abuse Enterprise.

126. The Azure Abuse Enterprise has engaged in activities that affect interstate commerce through a pattern of racketeering activity in violation of 18 U.S.C. § 1962(c).

127. Defendants conspired to operate the Azure Abuse Enterprise through a pattern of racketeering activity in furtherance of the common purpose of the Enterprise sometime prior to July 2024. Thereafter, each Defendant took wrongful acts in furtherance of their unlawful agreement by supplying resources to the Azure Abuse Enterprise. Defendants continuously and effectively carried out the purpose of the Azure Abuse Enterprise from at least July to September 2024, causing harm to the business and property of Microsoft and others. Defendants use of the stolen Azure customer credentials referenced herein to gain fraudulent access to Microsoft's systems would have continued beyond September 2024 but for Microsoft's actions to invalidate and replace those customers' credentials. Defendants represent a continuing threat to Microsoft and others and would likely resume their attacks on the Azure OpenAI Service upon coming into possession of additional stolen customer credentials.

128. **Wire Fraud (18 U.S.C. § 1343).** At some point prior to July 2024, Defendants devised a scheme to obtain money or property from Microsoft's paying customers, and to defraud Microsoft, by stealing authentication information from Microsoft customers and misusing that authentication information to gain fraudulent access the Azure OpenAI Service. Defendants understood and intended that their misuse of stolen customer authentication information would deplete the account balances of the paying Microsoft customers whose credentials they stole. Defendants devised their scheme at least in part to avoid paying the costs

of obtaining a license to the Azure OpenAI Service and purchasing the tokens required to use the Service at scale.⁶

129. From July 26, 2024, to at least September 17, 2024, Defendants transmitted and/or caused to be transmitted by means of wire communication in interstate and foreign commerce writings, signals, and pictures for the purpose of executing their scheme to defraud. For example, on numerous occasions between July 26, 2024, and August 18, 2024, Defendants transmitted by means of wire communication in interstate and foreign commerce stolen API Key and token information in order to defraud Microsoft regarding Defendants' identities and authorization to access Microsoft's systems and to deprive Microsoft's paying customers of tokens they had paid for. Defendants continued to use communications transmitted by means of wire communication in interstate and foreign commerce in furtherance of their scheme until at least September 17, 2024, when changes to the oai reverse proxy service were published by one or more of the Defendants.

130. **Access Device Fraud (18 U.S.C. § 1029).** From July 26, 2024, to at least August 18, 2024, Defendants knowingly and with the intent to defraud produced, used, and trafficked in counterfeit access devices including the oai reverse proxy server and de3u computers.

131. From July 26, 2024, to at least August 18, 2024, Defendants knowingly and with intent to defraud trafficked in and used unauthorized access devices, and by such conduct obtained a thing of value aggregating \$1,000 or more during that period.

132. Microsoft is informed and believes, and hereby alleges that discovery is likely to yield evidentiary support showing that Defendants have engaged in similar unlawful conduct in the past and that at least two Defendants are known associates of one another. Defendants' preexisting associations and pattern of unlawful activity makes them a continuing risk for conducting the affairs of the Azure Abuse Enterprise through a pattern of racketeering.

⁶ Tokens refer to the basic units of input and output that the Service processes. Generally, models accessed through the Azure OpenAI Service understand and process text by breaking it down into tokens. Microsoft provides transparent pricing details for input and output tokens at its publicly available website, <https://azure.microsoft.com/en-us/pricing/details/cognitive-services/openai-service/>.

133. The conduct described above has caused harm to Microsoft's business and property in an amount to be computed at trial.

134. The conduct described above was willful and with knowledge of wrongdoing.

135. Microsoft is entitled to and hereby demands treble damages, attorneys' fees, and costs of suit in addition to preliminary and permanent injunctive relief.

SIXTH CLAIM FOR RELIEF

Common Law Trespass to Chattels

136. Microsoft realleges and incorporates by this reference each and every allegation set forth in the paragraphs above.

137. Defendants' actions in abusing the Azure OpenAI Service resulted in unauthorized access to the computers of Microsoft.

138. Defendants intentionally caused this conduct and this conduct was unauthorized.

139. Defendants' actions have caused injury to Microsoft including time, money, and a burden on the computers of Microsoft and Microsoft's customers. Defendants' actions have caused injury to Microsoft's business goodwill and have diminished the value of Microsoft's and its customers' possessory interest in their computers and software.

140. Plaintiff Microsoft seeks injunctive relief and compensatory and punitive damages in an amount to be proven at trial.

141. As a direct result of Defendants' actions, Microsoft has suffered and continues to suffer irreparable harm for which it has no adequate remedy at law, and which will continue unless Defendants' actions are enjoined.

SEVENTH CLAIM FOR RELIEF

TORTIOUS INTERFERENCE

142. Microsoft realleges and incorporates by this reference each and every allegation set forth in the paragraphs above.

143. Microsoft has valid contracts with the customers who have been victimized by Defendants.

144. Defendants had knowledge of Microsoft's customer contracts and intentionally set out to wrongfully use Microsoft's customers' contracts and funds for Defendants' own unlawful purposes.

145. Defendants have interfered with Microsoft's contracts with its customers by stealing customer account information and using that information to deplete customer account funds.

146. Defendants' conduct has impeded the parties to Microsoft's customer contracts' abilities to perform their respective obligations.

147. Microsoft has been damaged by Defendants' conduct in an amount subject to determination at trial.

148. As a direct result of Defendants' actions, Microsoft has suffered and continues to suffer irreparable harm for which it has no adequate remedy at law, and which will continue unless Defendants' actions are enjoined.

PRAYER FOR RELIEF

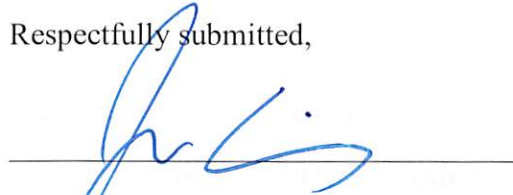
WHEREFORE, Microsoft prays that the Court:

1. Enter judgment in favor of Plaintiff and against the Defendants;
2. Declare that Defendants' conduct has been willful and that Defendants have acted with fraud, malice, and oppression;
3. Enter a preliminary and permanent injunction enjoining Defendants and their officers, directors, principals, agents, servants, employees, successors, and assigns, and all persons and entities in active concert or participation with them, from engaging in any of the activity complained of herein or from causing any of the injury complained of herein and from assisting, aiding, or abetting any other person or business entity in engaging in or performing any of the activity complained of herein or from causing any of the injury complained of herein;
4. Enter a preliminary and permanent injunction isolating and securing the infrastructure, including the domain "atism.net" and the software operating from and through the infrastructure, outside of the control of Defendants or their representatives or agents;

5. Enter judgment awarding Microsoft actual damages in an amount to be proven at trial;
6. Enter judgment in favor of Microsoft disgorging Defendants' profits, and;
9. Order such other relief that the Court deems just and reasonable.

Dated: December 19, 2024

Respectfully submitted,



JOSHUA CARRIGAN (VA Bar No. 96911)

jcarrigan@orrick.com

ORRICK, HERRINGTON & SUTCLIFFE LLP

2100 Pennsylvania Avenue NW

Washington, D.C. 20037

Telephone: + 202 339 8400

Facsimile: + 202 339 8500

ROBERT L. URIARTE (*Pro Hac Vice* forthcoming)

ruriarte@orrick.com

ORRICK, HERRINGTON & SUTCLIFFE LLP

355 S. Grand Ave.

Ste. 2700

Los Angeles, CA 90017

Telephone: + 1 213 629 2020

Facsimile: + 1 213 612 2499

JACOB M. HEATH (*Pro Hac Vice* forthcoming)

jheath@orrick.com

ANA M. MENDEZ-VILLAMIL (*Pro Hac Vice* forthcoming)

amendez-villamil@orrick.com

ORRICK, HERRINGTON & SUTCLIFFE LLP

The Orrick Building

405 Howard Street

San Francisco, CA 94105

Telephone: + 1 415 773 5700

Facsimile: + 1 415 773 5759

LAUREN BARON (*Pro Hac Vice* forthcoming)
lbaron@orrick.com
ORRICK, HERRINGTON & SUTCLIFFE LLP
51 West 52nd Street
New York, NY 10019
Telephone: + 1 212 506 5000
Facsimile: + 1 212 506 5151

Of Counsel:

RICHARD BOSCOVICH
rbosco@microsoft.com
MICROSOFT CORPORATION
Microsoft Redwest Building C
5600 148th Ave NE
Redmond, Washington 98052
Telephone: +1 425 704 0867
Facsimile: +1 425 706 7329

Attorneys for Plaintiff
MICROSOFT CORPORATION