Securing our approach to domain fronting within Azure

microsoft.com/security/blog/2021/03/26/securing-our-approach-to-domain-fronting-within-azure/

March 26, 2021

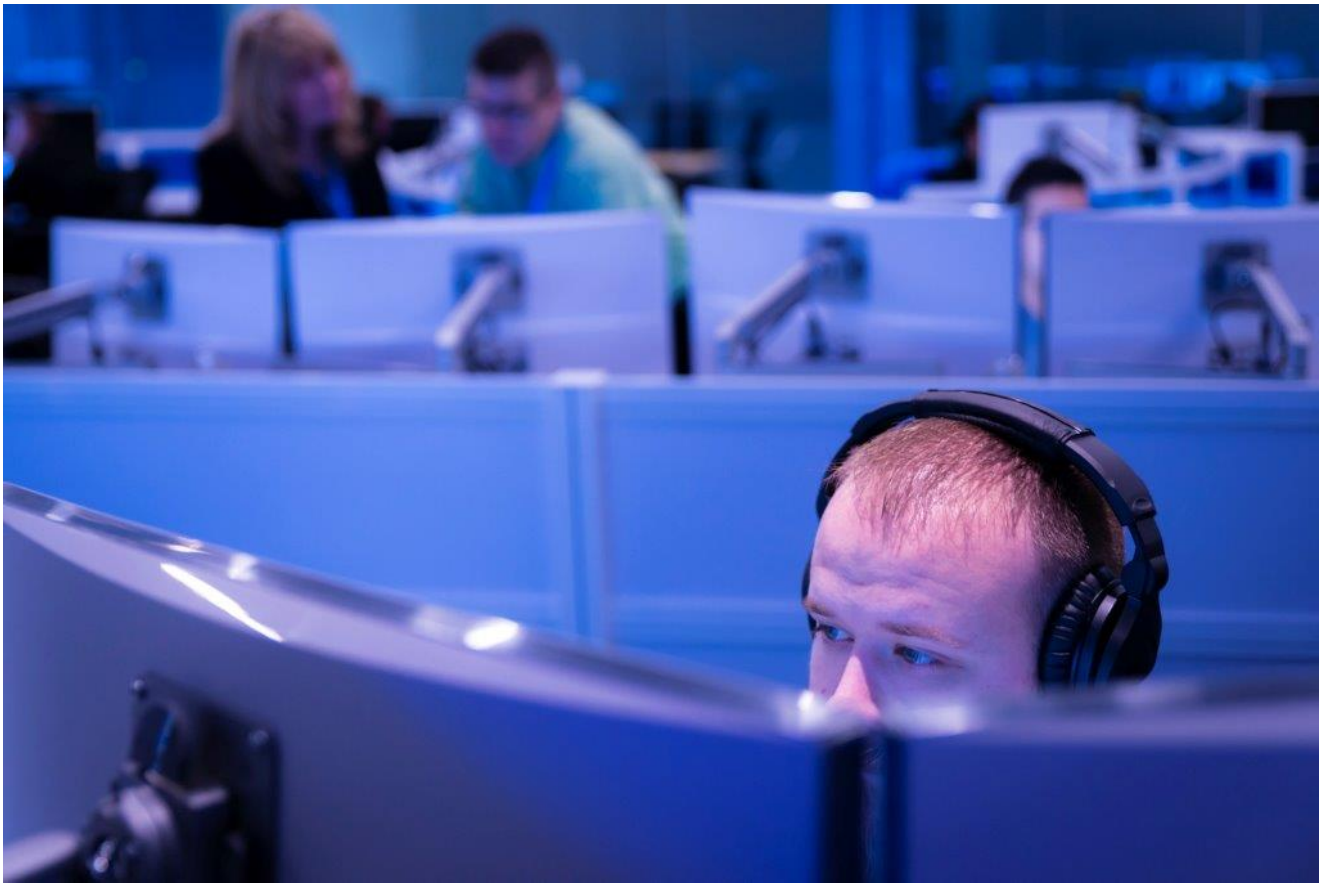

Every single day our teams analyze the trillions of signals we see to understand attack vectors, and then take those learnings and apply them to our products and solutions. Having that understanding of the threat landscape is key to ensuring our customers are kept safe every day. However, being a security provider in a complex world sometimes requires deeper thinking and reflection on how to address emerging issues, especially when the answer is not always immediately clear. Our approach to domain fronting within Azure is a great example of how the ever-changing dynamics of our world have prompted us to re-examine an important and complicated issue—and ultimately make a change.

Let's start with some background. Domain fronting is a networking technique that enables a backend domain to utilize the security credentials of a fronting domain. For example, if you have two domains under the same content delivery network (CDN), domain #1 may have certain restrictions placed on it (regional access limitations, etc.) that domain #2 does not. By taking the valid domain #2 and placing it into the SNI header, and then using domain #1 in the HTTP header, it's possible to circumvent those restrictions. To the outside observer, all

subsequent traffic appears to be headed to the fronting domain, with no ability to discern the intended destination for particular user requests within that traffic. It is possible that the fronting domain and the backend domain do not belong to the same owner.

As a company that is committed to delivering technology for good, supporting certain use cases that support free and open communication are an important consideration when weighing the potential impacts of a technique like domain fronting. However, we know that domain fronting is also abused by bad actors and threat actors engaging in illegal activities, and we've become aware that in some cases bad actors configure their Azure services to enable this.

When it comes to situations like this, Microsoft—as a security company—leads from a place of providing greater simplicity for our customers when they face increased complexity. Our mission is to give our customers peace of mind and help them adapt quickly to a rapidly shifting threat landscape. Therefore, we're making a change to our policy to ensure that domain fronting will be stopped and prevented within Azure.

Changes like this one are not made lightly, and we understand that there will be impacts across a number of areas:

- Our engineering teams are already working to ensure the platform will block anyone from practicing the domain fronting technique on Azure, while also continuing to ensure our products and services provide the highest levels of protection against domain fronting based threats.
- We're continuing to provide clear guidance for penetration testing on our Azure properties, and working closely with security researchers around the world to make sure they have a clear understanding of these changes.

These changes are just another example of the broad impact that security has on our ever-changing world and we'll continue to put the security of our customers and their users at the forefront of everything we do. I'd like to thank my colleagues Nick Carr and Christopher Glyer for their tireless research on Domain Fronting, which helped us to make these policy changes to Azure.

To learn more about Microsoft Security solutions, visit our website. Bookmark the Security blog to keep up with our expert coverage on security matters. Also, follow us at @MSFTSecurity for the latest news and updates on cybersecurity.