CloudFucius Inspects: Hosts in the Cloud

Peter Silva, 2010-18-05

So much has been written about all the systems, infrastructure, applications, content and everything else IT related that's making it's way to the cloud yet I haven't seen much discussion (or maybe I just missed it) about all the clients connecting to the cloud to access those systems. Securing those systems has made some organizations hesitate in deploying IT resources in the cloud whether due to compliance, the sensitivity of the data, the shared infrastructure or simply persuaded by survey results. Once a system is 'relatively' secure, how do you keep it that way when the slew of potentially dangerous, infected clients connect? With so many different types of users connecting from various devices, and with a need to access vastly different cloud resources, it's important to inspect every requesting host to ensure both the user and the device can be trusted. Companies have done this for years with remote/SSL VPN users who request access to internal systems – is antivirus installed and up to date, is a firewall enabled, is the device free of malware and so forth. Ultimately, the hosts are connecting to servers housed in some data center and all the same precautions you have with your own space should be enforced in the cloud.

Since cloud computing has opened application deployment to the masses, and all that's required for access is *potentially* just a browser, you must be able to detect not only the type of computer (laptop, mobile device, kiosk, etc.) but also its security posture. IDC predicts that '*The world's mobile worker population will pass the one billion mark this year and grow to nearly 1.2 billion people – more than a third of the world's workforce – by 2013'* With so many Internet-enabled devices available; a Windows computer, a Linux box, an Apple iteration, a mobile device and anything else with an IP address, they could all be trying to gain access to your cloud environment at any given moment. It might be necessary to inspect each of these before granting users access in order to make sure it's something you want to allow. If the inspection fails, how should you fix the problem so that the user can have some level of access? If the requesting host is admissible, how do you determine what they are authorized to access? And, if you allow a user and their device, what is the guarantee that nothing proprietary either gets taken or left behind? The key is to make sure that only "safe" systems are allowed to access your cloud infrastructure, especially if it contains highly sensitive information and context helps with that.

One of the first steps to accomplishing this is to chart usage scenarios. Working in conjunction with the security policy, it is essential to uncover the usage scenarios and access modes for the various types of users and the many devices that they might be using. The chart will probably vary based on your company's and/or website's Acceptable Use Policy, but this exercise gets administrators started in determining the endpoint plan. Sounds a lot like a remote access policy, huh, with one exception. Usually there is a notion of 'trusted' and 'un-trusted' with remote access. If a user requests access from a corporate issued laptop, often that's considered a trusted device since there is something identifiable to classify it as an IT asset. These days, with so many personal devices entering the cloud, all hosts should be considered un-trusted until they prove otherwise. And as inter-clouds become reality, you'll need to make sure that a client coming from someone else's infrastructure abides by your requirements. Allowing an infected device access to your cloud infrastructure can be just as bad as allowing an invalid user access to proprietary internal information. This is where endpoint security checks can take over. Endpoint security prevents infected PCs, hosts, or users from connecting to your cloud environment. Automatic re-routing for infected PCs reduces Help Desk calls and prevents sensitive data from being snooped by keystroke loggers and malicious programs.

Simply validating a user is no longer the starting point for determining access to cloud systems; the requesting device should get the first review. Pre-access checks can run prior to the actual logon (if there is one) page appearing, so if the client is not in compliance, they won't even get the chance to enter credentials. These checks can determine if antivirus or firewall is running, if it is up-to-date, and more. Systems can direct the user to a remediation page for further instructions to gain access. It's easy to educate the user as to why the failure occurred and relay the possible steps to resolve the problem. For example: "We noticed you have antivirus installed but not running. Please enable your antivirus software for access." Or, rather than deny logon and communicate a detailed remedy, you could automatically send them to a remediation website designed to correct or update the client's software environment, assuring policies required for access are satisfied without any user interaction. Inspectors can look for certain registry keys or files that are part of your corporate computer build/image to determine if this is a corporate asset and thus, which system resources are allowed. Pre-access checks can retrieve extended Windows and Internet Explorer info to ensure certain patches are in place. If, based on those checks, the system finds a non-compliant client but an authorized user; you might be able to initiate a secure, protected, virtual workspace for that session.

As the ever-expanding cloud network grows, the internal corporate resources require the most protection as it's always been. Most organizations don't necessarily want all users' devices to have access to all resources all the time. Working in conjunction with the pre-access sequence, controllers can gather device information (like IP address or time of day) and determine if a resource should be offered. A protected configuration measures risk factors using information collected by the pre-access check; thus, they work in conjunction. For example, Fake Company, Inc. (FCI) has some contractors who need access to Fake Company's corporate cloud. While this is not an issue during work hours, FCI does not want them accessing the system after business hours. The controller can check the time if a contractor tries to log on at 2 AM; it knows the contractor's access is only available during FCI's regular business hours and can deny access.

Post-access actions can protect against sensitive information being "left" on the client. The controller can impose a cache-cleaner to eliminate any user residue such as browser history, forms, cookies, auto-complete information, and more. For systems unable to install a cleanup control, you can block all file downloads to avoid the possibility of the inadvertent left-behind temporary file—yet still allow access to needed cloud applications. These actions are especially important when allowing non-recognized machines access without wanting them to take any data with them after the session.

In summary: First, inspect the requesting device; second, protect resources based on the data gathered during the check; third, make sure no session residue is left behind. Security is typically a question of trust. Is there sufficient trust to allow a particular user and a particular device full access to enterprise cloud resources? Endpoint security gives the enterprise the ability to verify how much trust and determine whether the client can get all the cloud resources, some of the cloud resources, or just left in the rain.

And one from Confucius: *When you know a thing, to hold that you know it; and when you do not know a thing, to allow that you do not know it - this is knowledge.*

ps

The CloudFucius Series: Intro, 1, 2, 3, 4, 5

Digg This

F5 Networks, Inc. | 401 Elliot Avenue West, Seattle, WA 98119 | 888-882-4447 | f5.com

F5 Networks, Inc.
Corporate Headquarters
info@f5.com

F5 Networks
Asia-Pacific
apacinfo@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.com