GARRISON

Endpoint strategies for Zero Trust Architectures

To quote NIST SP.800-207[1], "Zero trust is a cybersecurity paradigm focused on resource protection and the premise that trust is never granted implicitly but must be continually evaluated." As NIST further explains, the objective is – of course – that "authroised and approved subjects (combination of user, application or service, and device) can access the data to the exclusion of all other subjects (i.e., attackers)".

One of the least well understood aspects of implementing a Zero Trust Architecture is the treatment of devices. The criticality of user identity is well rehearsed ("Is access to the resource allowable given the level of confidence in the subject's identity?") But there is often less clarity regarding the importance of the device ("Does the device used for the request have the proper security posture?")

One of the key aspects of Zero Trust is that it is contextual. The "level of confidence in the subject's identity" and the "proper security posture" depend on the level of sensitivity of the resources being accessed. It is desirable to restrict access to the corporate expenses system to genuine employees, using endpoint devices that have not been compromised. But it is critical to restrict access to core systems administrations functions only to genuine systems administrators using high-integrity endpoint devices.

There are therefore three questions that organisations need to ask themselves about endpoint devices.

QUESTION 1

What is the "proper security posture" of an endpoint device for accessing each of the resources (or resource groups) in your Zero Trust Architecture?

It is essential that this question is understood to be separate from questions of user identity: a highly-trusted user may still need to be denied access to a resource if they are trying to access it from a device with an insufficient security posture. Why? Because a compromised device could be in use by two users at the same time: the trusted user, and the attacker. Even identity techniques such as biometrics and MFA cannot fully mitigate this risk – a truth that has been long-established in the world of online banking.

Online banking represents exactly this scenario: an authorised, authenticated user accessing from personal devices whose security posture is usually wholly unknown being used to access critical banking services. Financial services organisations are in a continual fight with criminal gangs who seek to abuse this fact in order to steal money, and it is a fight that can never be fully won. A persistent level of fraudulent activity is enabled by ever-more sophisticated attack techniques such as man-in-the-browser and MFA bypass. The risk is accepted, and a certain level of resulting financial loss is baked in to online banking business models.

1 NIST Special Publication 800-207: Zero Trust Architecture (https://doi.org/10.6028/NIST.SP.800-207)

For some corporate resources, a similar risk profile may be appropriate. Where the potential impact is forecastable and limited, the risk can be accepted: access from even personal devices with wholly unknown security postures can be considered. But for other resources, the business impact may be so potentially dramatic that no CISO could in good faith accept that risk. In a banking context, loss of funds from a single user account may be manageable, while manipulation of core inter-bank accounts could represent an unacceptable (and potentially systemic) risk.

QUESTION 2

How to determine whether an endpoint device has an adequate security posture for the resource being accessed?

This is core to a Zero Trust Architecture, but can be a challenging area for implementation. While user identity is well served with interoperable standards such as SAML and OAuth, device posture is not, and individual organisations need to find their own ways to discriminate.

This is particularly challenging in the context of cloud services, which often include relatively poor support for such device discrimination. Today, many cloud services provide zero support for device discrimination, focusing exclusively on user identity: even those that do typically provide only a very blunt instrument, allowing access to be conditional on the source IP address. This then throws the responsibility back onto the enterprise, to ensure that devices with particular security postures are funnelled through specific IP addresses, with suitable security posture discrimination applied (for example, in a CASB) at that point.

In an ideal deployment, User A might be permitted full administrative rights in Cloud Service X when using a high-security-posture device, and only reduced user rights in that cloud service when using a lower-security-posture device. Until more cloud services support such approaches, for most cloud services the standard approach will be that if User A is not using a high-security-posture device, they will be wholly denied access to that cloud service. (Meanwhile User B, whose rights in the cloud service are only at the user level, may be permitted to access the service even when using a lower-security-posture device).

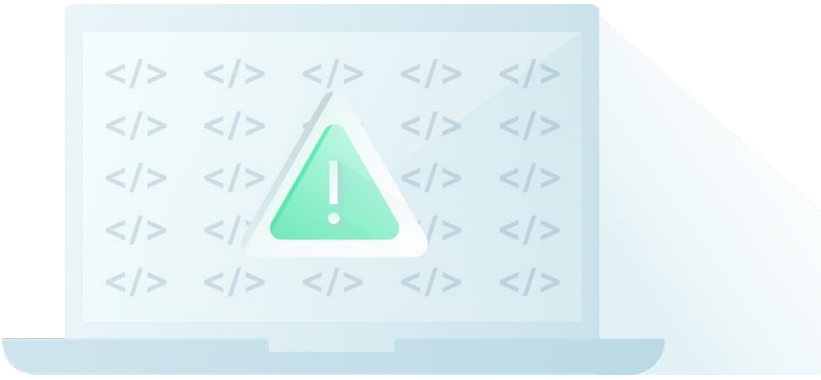

QUESTION 3

Ultimately of course, the key question is the third one: how can an adequate security posture be maintained on the endpoint device?

In this area, guidance on Zero Trust remains largely silent, focusing only on how the security posture can be discriminated. It is left to the separate discipline of Endpoint Security to consider what steps need to be taken to deliver an adequate security posture for the device.

Nonetheless, achieving an adequate endpoint security posture is an essential aspect of a Zero Trust Architecture. Without it, Zero Trust can deliver security only at the expense of business utility – by preventing legitimate users from accessing corporate resources. If we focus only on discriminating endpoint security (rather than delivering endpoint security) then we can only block, not allow. As per NIST: "Assets that are discovered to be subverted, have known vulnerabilities, and/or are not managed by the enterprise may be treated differently (including denial of all connections to enterprise resources) than devices owned by or associated with the enterprise that are deemed to be in their most secure state."

Many elements of endpoint security are of course well rehearsed: the importance of patching, most particularly. Beyond that, the Endpoint Security market is replete with tools such as Antivirus and EDR: but the real challenge for any organisation is to know how effective such tools really are. Particularly for those devices which require a high security posture, how can such a security posture be adequately maintained?

One approach to endpoint security takes the principles of Zero Trust as applied to resources, and applies them to the endpoint. Just as Zero Trust requires "minimising access to resources…to only those subjects and assets identified as needing access", so endpoint security can also minimise access. When endpoints can only access trusted resources, the opportunity to compromise them is drastically reduced, and their security posture dramatically increased.

The most significant minimisation of course applies to Internet-based resources: for endpoint devices requiring more than the lowest security posture, it is normal to restrict access to known malicious resources, as well as to restrict communications protocols (for example, permitting HTTPS only for most resources). But "known malicious" represents a very minimal bar for restriction: the Internet is replete with resources which may not be known malicious, but which certainly represent a very high level of risk. For devices requiring a higher security posture, it would be desirable to restrict access to these also. How can this be done?

In today's Internet, the scope for risk is so high that the most practical answer is to take an inverse approach: to define the resources that represent an acceptable level of risk, and to restrict access to everything else. This allow-list approach is one which is increasingly being considered – if not for all endpoints, at least for those which require a higher security posture, such as those belonging to users who have elevated privileges giving them access to and control over critical systems and data.

The challenge of course is that just as restricting access to corporate resources is a problem from a utility perspective, so too can be restricting access to Internet resources. In today's world, the ability to assimilate information that is just a click away is a critical part of many individuals' working patterns.

Browser Isolation is designed to address precisely this issue: permitting users to access resources which have been restricted in order to maintain the security posture of the endpoint. But it is not hard to discern a paradox: how can access be provided to something which has been restricted? Surely the promise of Browser Isolation must be a fool's paradise – an impossible security concept marketed only to the credulous?

Endpoint strategies for Zero Trust Architectures

And yet some Browser Isolation technologies are trusted for this purpose not only by mainstream enterprises, but also by some of the most sensitive government security agencies – the very agencies who are themselves involved in seeking to penetrate and compromise the endpoint devices of their adversaries. How have they come to believe in the efficacy of such technologies?

CONCLUSION

At the heart of the answer is the concept of a "verifiable pixel gap"

An approach which verifies that the only data which flows from untrusted (and otherwise restricted) resources is a stream of raw pixels (and equivalently, raw audio samples). There are good theoretical reasons why such a pixel gap delivers the seemingly impossible: high-security access to risky resources. Indeed one use of a verifiable pixel gap is to access even the known malicious: a niche requirement, certainly, but a level of security that is difficult otherwise to achieve.

As NIST states, "Implementing a ZTA is a journey rather than a wholesale replacement". It is a journey prompted by the need to maintain the security of enterprise resources while delivering on the business promises of flexibility, mobility and rapid innovation – in an online world with an ever-growing threat level. And as the name itself implies, a Zero Trust Architecture is just that: an architecture rather than a single point solution.

Within a Zero Trust architecture, the verifiable security posture of endpoint devices is just as important as the verification of user identity. Identifying the required security postures, delivering them, and verifying them, is an integral part of any Zero Trust journey.

Email info@garrison.com
UK Telephone +44 (0) 203 890 4504
US Telephone +1 (646) 690-8824

CD00000788 - January 2023