Zero-Trust Architecture

Abstract

The zero-trust architecture was created by John Kindervag in 2010 to address the security flaws of the traditional model and insists that all internal network traffic should also not be trusted by default. Traditional networks are usually designed to have a security perimeter for the incoming traffic coming from the outside world, but not for the incoming traffic coming from inside the network. This makes them vulnerable to attackers who can breach the network without even having to deal with the security perimeter designed for the outside world. The COVID-19 pandemic has made us realise that there is a need for improvements in security within the network, as many healthcare providers are still using vulnerable, outdated legacy systems that can be compromised. Thus, as more users have now started to work remotely and as most of the assets are being moved to the cloud, relying solely on the traditional perimeter approach of using only firewalls and VPNs will be less effective, less efficient and more dangerous. The legacy systems and medical devices that the network administrators cannot manage or control cause restrictions and boundaries in transitioning to a zero-trust security model. The zero-trust architecture also proposes to continuously verify and monitor all communications, as well as encryption of all data, i.e. in transit or at rest.

1. Introduction

The purpose of this study is to identify the problems faced by healthcare organisations in terms of cybersecurity and how a zero-trust approach could solve these problems. The number of IP-enabled connected medical devices is growing, making them vulnerable to breaches that could potentially have an impact on the effectiveness of the device. This vulnerability increases as medical devices and equipment are increasingly becoming connected via the network to other devices, patients and/or healthcare organisation networks. Research has shown that although organisations would benefit from the increased security provided by adopting a zero-trust architecture, many remain hesitant to make the move due to financial implications and/or due to the limitations of the legacy systems.

The zero-trust architecture is a security concept that takes the far-sighted approach to verifying services, devices and individual users, rather than trusting them by default. It asserts that we should verify everything and everyone and trust no one. A zero-trust model supports micro-segmentation, which is a fundamental principle of cybersecurity. It enables us to contain potential threats and not let them spread throughout the enterprise.

Zero-trust network access (ZTNA), which is part of the zero-trust model, uses identity-based authentication to provide access while keeping the network location, i.e. the IP address, hidden. When adopting a zero-trust security model – whether in the cloud or on-premises, it is required to enforce user access policies and have robust authentication mechanisms and tools for creating software-defined security perimeters.

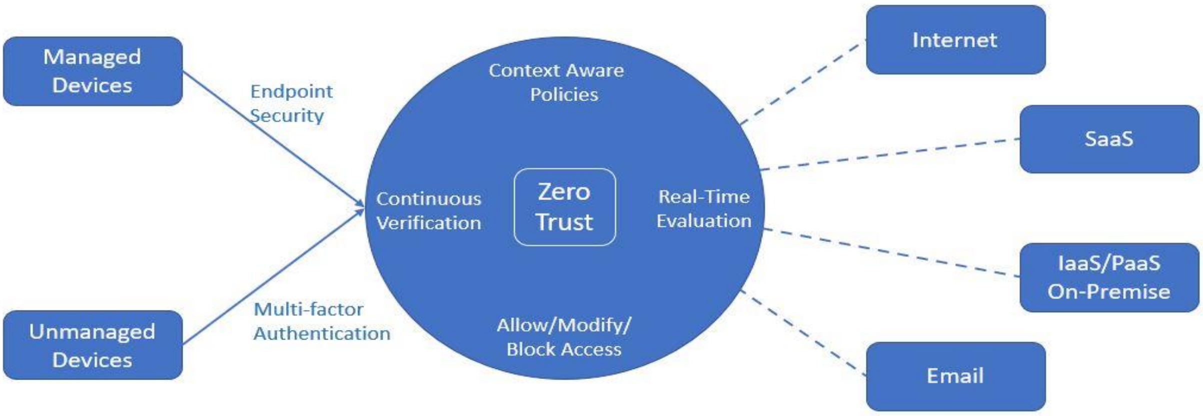

Figure 1: Zero-Trust Architecture

1.1 Scope of a zero-trust model

Zero-trust security cannot be attained through a single tool or a platform, rather it is an approach. The approach usually involves technologies.

- Know what is to be protected – users, devices, data, services and the network.
- Understand the cybersecurity controls already in place.
- Incorporate new tools and modern architecture.
- Apply detailed policy.
- Deploy monitoring and alerting tools.

2. Literature survey

This section contains a literature survey of the relevant and crucial papers related to the problem being researched.

2.1 The drawbacks of the perimeter model

Healthcare organisations that rely on the permitter model (VPNs and firewalls) leave themselves vulnerable to attacks from the inside. Protenus' 2020 report mentions that in 2019, the number of patient records that were breached due to attacks that originated from the inside of the network was up by 26% (3.8 million records) when compared to 2018. An infamous real-world example that can be used to argue for prioritising security within the perimeter is the phishing attack against a Montpellier medical centre. One employee of the Montpellier medical centre opened an email containing a virus that infected more than 600 of their machines. Fortunately, the Montpellier medical centre had independent internal networks, so the virus was restricted to 600 machines in the infected network and did not spread to all of their 6000 machines. A zero-trust approach would have prevented the virus from infecting even the 600 machines that it did. Thus, focusing on the exterior of the network and believing we can trust everyone and every device on the inside is not the way to move forward.

2.2 Micro-segmentation

The first step in implementing zero trust is micro-segmentation of the network. This is a method of logically creating network segments and completely controlling traffic within and between the segments. The traditional micro-segmentation techniques using firewalls, switches and VLAN/ACLs have been rendered inadequate by the emergence of the cloud and advanced cyberthreats, as the traditional techniques cannot safeguard applications in a hybrid and dynamic environment. The future of micro-segmentation is host-based, which means the security barrier must be moved down to the individual hosts. A software-defined perimeter, or SDP as it is commonly known, is a way of enabling host-based micro-segmentation and can be implemented above the network layer without making significant changes to the existing hardware infrastructure. An SDP restricts access to a particular network until the request is properly authenticated and validated.

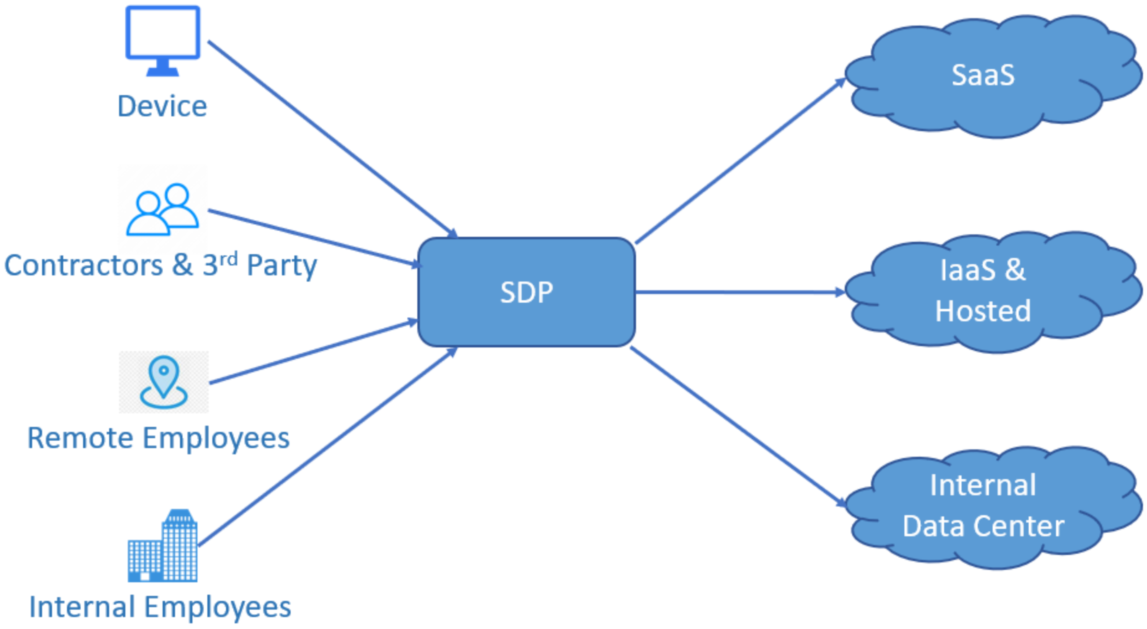

Figure 2: Software-Defined Perimeter

3. Zero-trust implementation in healthcare

Implementation of the zero-trust model involves the following stages:

- Network segmentation – this phase involves making sure that all the connected medical devices can only communicate with systems or devices that are part of their clinical process.
- Block unnecessary communications – this phase involves understanding exactly which communications are needed to maintain the medical device's functionality and designing policies to block any other unnecessary communications. It also involves limiting external communications to the bare minimum.

- Service hardening – this phase involves isolating risks associated with the services used on the individual devices. It is important in this step to require authentication and validation on all communication channels, reduce unnecessary functions and close all the unused ports. Apart from this, all the connected medical devices must be evaluated in order to perform any necessary software upgrades and apply the latest security patches.

4. Conclusion

The SDP architecture can be used to implement the zero-trust model while allowing healthcare organisations to continue using traditional implementations. The important factor to consider while designing a zero-trust framework in the healthcare industry is that it should be secure but not so restrictive that it interrupts patient care. The medical devices being built today with high security standards will still be used in the future and by then, new vulnerabilities may have been discovered. It is therefore imperative that we design a security framework that remains relevant for the foreseeable future. A healthcare organisation might be hesitant to move to a zero-trust architecture in light of the financial aspects, the implementation efforts and the time that might be required, as they cannot risk interrupting patient care for too long. So it is always better to make a start on transitioning towards a zero-trust architecture than to not move forward at all.