WHITEPAPER

CISO Guide:

5 STEPS TO IMPROVE ENTERPRISE CYBER-RESILIENCE

Balbix®

CISO Guide:

5 STEPS TO IMPROVE ENTERPRISE CYBER-RESILIENCE

Introduction

Cyber-resilience is becoming a popular term in the cyber security business and is defined as *the ability of an enterprise to limit the impact of security attacks*. Focusing on cyber-resilience is part of a broader approach to digital risk management that aims to not only defend against cyber-attacks, but also ensure that the enterprise is able to survive and recover quickly following an attack.

The attack surface of a typical organization is massive, growing rapidly, and the threat landscape is constantly evolving. Enterprise networks contain large amounts of insecure software (and hardware) which fail regularly from a cybersecurity standpoint. People in the organization also make routine cybersecurity mistakes, like clicking on the wrong link. No matter how good your security program is, it is impossible to avoid 100% of such failures due to reasons of scale. The ultimate objective of security teams is therefore to implement risk mitigations that result in a cyber-resilient enterprise on top of insecure components. Incorporating an understanding of cyber-resilience in strategic planning is a key to implementing and operating an effective cybersecurity program.

How is cyber-resilience defined and measured? How are breach risk and cyber-resilience related, and what is the best way to improve cyber-resilience for your enterprise? This paper will answer these questions and describe specific steps you can take to improve the cyber-resilience of your enterprise.

Cyber-resilience is the ability of an enterprise to limit the impact of security incidents.

Defining and Measuring Cyber-resilience

Measuring cyber-resilience means measuring the ability of an enterprise to limit the impact of security incidents. In this context, we are quantifying the effectiveness of preparations that an organization has made with regard to threats and vulnerabilities, the defenses that have been developed, controls implemented, and the resources available for mitigating a security failure after it happens. In essence, we want to enumerate all possible breach scenarios, evaluate the likelihood of every type of initial compromise, how the attack could propagate to other systems, and the response of the enterprise defenses.

Cyber-resilience is not an opaque score which may be derived from some simple scoring of some “n” properties of a network. It is also not a number that you arrive at by answering a checklist of questions about your network. It is quite a bit more complicated than that but can be calculated from continuous observations of the state of your enterprise and a series of probabilistic mathematical calculations.

Software is fragile. People make mistakes. The ultimate objective of your security program is to implement a cyber-resilient enterprise on top of insecure components.

To understand how we might quantify cyber-resilience, let's first double click into the structure of the enterprise attack surface, i.e., all the ways in which an adversary can attack an enterprise, as shown in Fig 1.

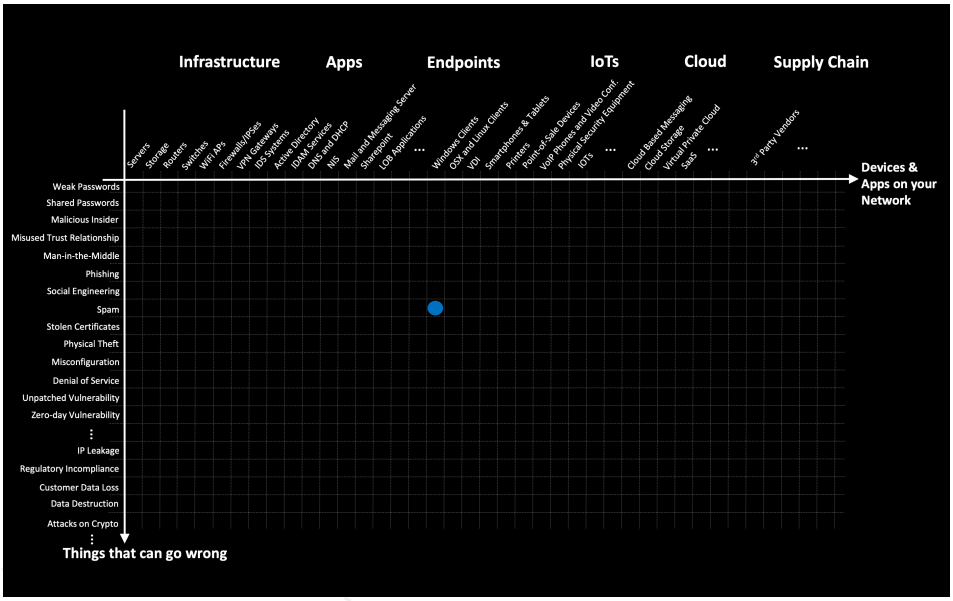

Fig 1: The enterprise attack surface

In this picture, the x-axis has the different parts of your enterprise– your traditional infrastructure, servers, databases, switches, routers, etc., apps, endpoints—managed, unmanaged, fixed and mobile, IoTs. We also have your cloud apps – personal apps of employees (e.g., Gmail, LinkedIn), official SaaS apps, public facing web sites, etc. At the extreme right, we have your 3rd party vendors who bring risk into your network because of certain trust relationships. Note that today, it is generally quite difficult for most organizations to even enumerate their x-axis for an accurate and up to date asset inventory.

On the y-axis, we have different attack methods—starting from simple things like weak, default and reused passwords, passwords stored unencrypted, or transmitted in the clear, and then more complex things like phishing, social engineering and unpatched software. Further down the y-axis, we have zero-day vulnerabilities, security bugs that are "unknown" until they are first used by the adversary. There are 100s of items on the y-axis in dozens of categories.

This gigantic x-y plot is your attack surface. In a typical breach, the adversary uses some point on this attack surface to compromise an (Internet facing) asset. Other points are then used to move laterally across the enterprise, compromise some valuable asset, and then to exfiltrate data or do some damage. Fig 2 shows how the Equifax breach in 2017 unfolded.

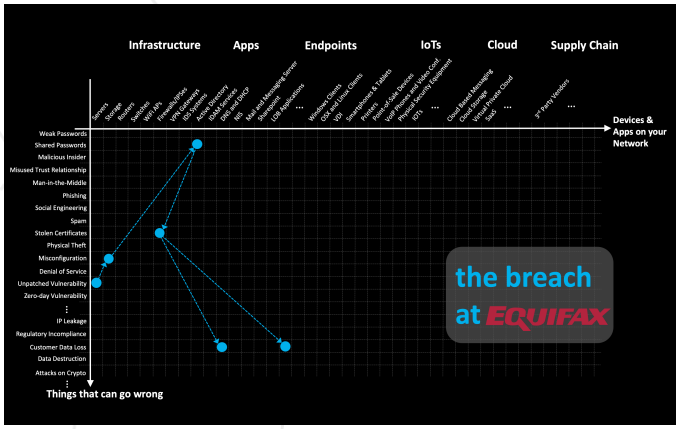

Fig 2: How the Equifax breach unfolded

It is instructive to think of the stages of a data breach as follows (Fig 3):

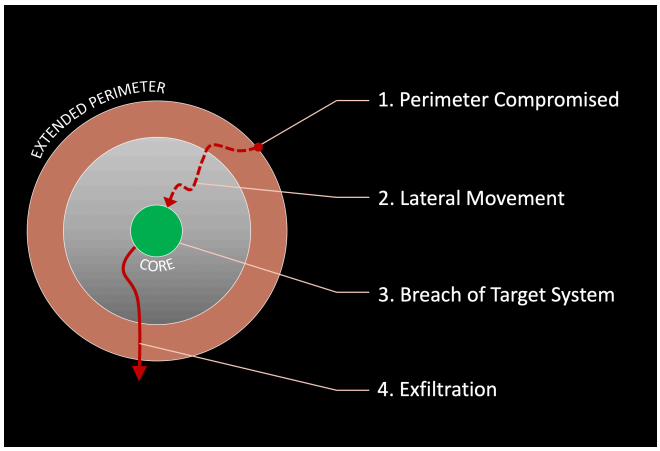

Fig 3: Stages of a breach

1. **Perimeter Compromised:** The adversary compromises some asset in the extended perimeter of the enterprise using some method from Fig 1's y-axis. Your organization's extended perimeter includes all enterprise assets from the x-axis of Fig 1 which are in direct contact with the Internet. This includes, for example, your DMZ and firewall complex, public web servers, VPN gateways, mobile devices, IoTs, every device that receives external email/messaging or is used to browse the Web, etc.
2. **Lateral Movement:** From this initial foothold, attackers use various techniques (from Fig 1's y-axis) to move inwards and sideways inside your intranet to gain access to additional systems. Lateral movement is generally easier than breaching the perimeter because we tend to optimize our internal access control policies and protocols for ease of collaboration and not for cybersecurity.
3. **Breach of Target System(s):** Ultimately the adversary is able to reach their target asset or find some system that contains valuable data or is otherwise important. At this time the adversary can attempt to steal data, destroy data, impair the availability of the system or just lie in wait.
4. **Exfiltration:** Most major cyber-attacks involve the adversary transferring valuable enterprise information to some system outside your network.

Now take a look at Fig 4 which plots the breach likelihood of an enterprise vs effort by the adversary.

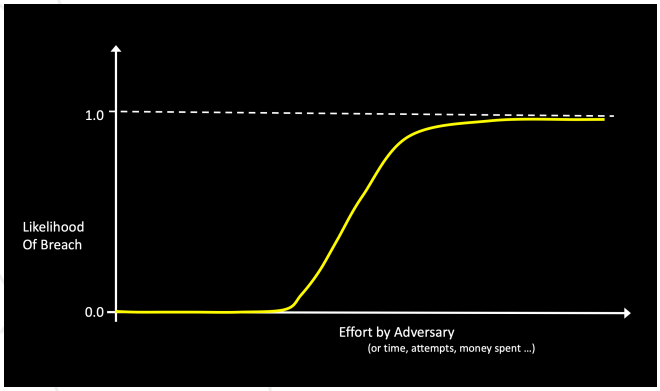

Fig 4: Breach likelihood vs effort

The y-axis in this picture is the likelihood of breach, a number between 0 and 1. The units on the x-axis are arbitrary and could be measured in a number of ways, for example, in terms of time spent by the adversary, or number of attempts, or money spent.

The shape of this likelihood vs effort curve is consistent with the intuitive notion that all security practitioners will agree on, i.e., *given enough effort, anything can be breached*. Every enterprise on this planet has a likelihood vs effort curve shaped like the one shown in Fig 4. The only difference in the curves of different organizations is where the knee of the curve is on the x-axis, and the slope of the rise from 0 to 1. This is shown in Fig 5.

An simple axiom of security: Given enough effort, *anything* can be breached.

For networks of security-mature companies, like a Fortune 50 bank, the knee would be expected to be well towards the right of the axis. For a smaller, less security mature company, we would expect the knee to be more towards the left (see Fig 5). This is largely true, but there is also a natural entropy in play that tends to shift larger and more complex networks to the left in Fig 5. Everything else being the same, it is easier to break into a network with 10K moving parts than it is to break into a network with 10 moving parts.

As you might imagine, this curve is not fixed for your network for all time. As you make changes to your network, the curve changes. The deployment of a new mitigation might push the curve significantly to the right, decreasing the slope. The discovery of a new vulnerability which ends up getting used by cyber-criminals, will move the curve to the left and perhaps make it steeper, until the vulnerability is patched.

If you are wondering about the steep rise from 0-to-1, recall the discussion earlier about the stages of a breach (Fig 3). A breach usually begins with one or a handful of vulnerable systems on the extended perimeter being compromised. This might happen because of an employee being phished, or via a weak password, or some unpatched vulnerability. In rare cases, the adversary may use a zero-day exploit to establish this initial beachhead. These initial compromise methods require varying degrees of effort by the adversary, as shown in Figure 6.

After the initial compromise, the adversary is able to jump quickly from system to system, which is why we see the sharp increase in the slope of the graph. Propagation is rapid because a fundamental design goal of our intranet is to enable ease-of-use and collaboration. Once the adversary has a beachhead on an enterprise device, they can use the fast pathways that exist inside the enterprise network to enable high productivity for the legitimate users of the compromised device.

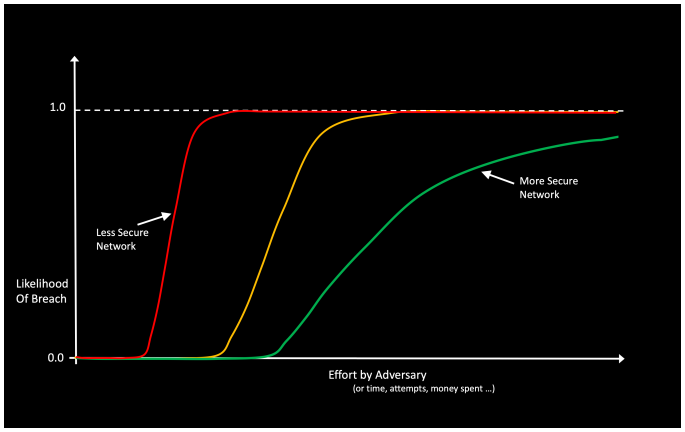

Fig 5: Breach likelihood vs effort for networks of different levels of security maturity

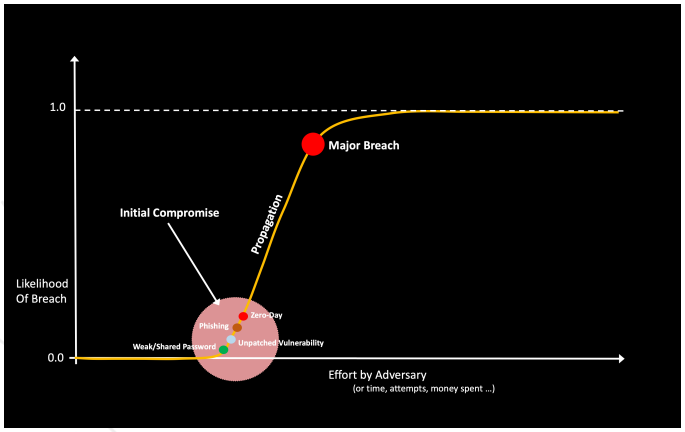

Fig 6: Initial compromise, propagation and major breaches

Once attackers have a beachhead on some enterprise device, they can use the optimized-for-collaboration "*flat network*" that exists inside the enterprise to quickly move laterally and cause a major breach.

So what is cyber-resilience in these graphs? Let's take a look at Fig 7 which is like Fig 4, but with *Expected Breach Impact* on the y-axis. The x-axis is *Effort by Adversary*, or *Threat Level*, both of which have congruent units.

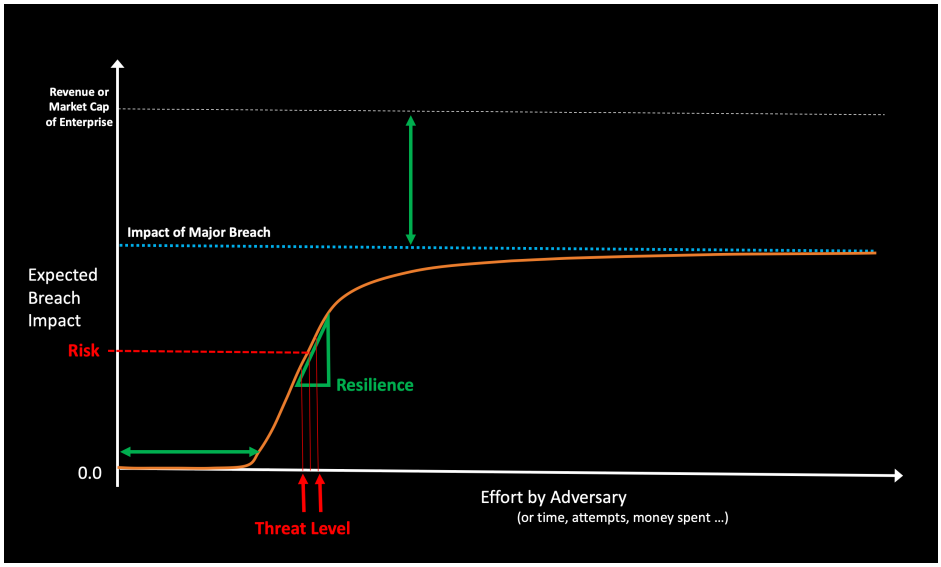

Fig 7: Breach risk and cyber-resilience

Breach risk is the expected breach impact at the current threat level. As you can see, *Breach Risk* varies greatly with small changes in both the level of efforts by the adversary for targeted attacks or in the global threat level for untargeted attacks. For this reason, *Breach Risk* is not a very useful metric for quantifying and comparing your overall cybersecurity posture (particularly if your organization's curve is towards the left edge of Fig 7).

A much more useful metric is cyber-resilience– whose factors are *the slope* of this Fig 7 curve, *the position of the knee* and *how high the curve can go*. If your defenses offer resistance to the adversary all three green metrics in Fig 7 are better.

What you ultimately want is a Breach Risk vs Effort curve for your organization that looks like Fig 8– *this is high cyber-resilience.* Notice how with this shape of curve, Breach Risk stays small and does not change much even though the threat level or targeted effort by the adversary changes a lot.

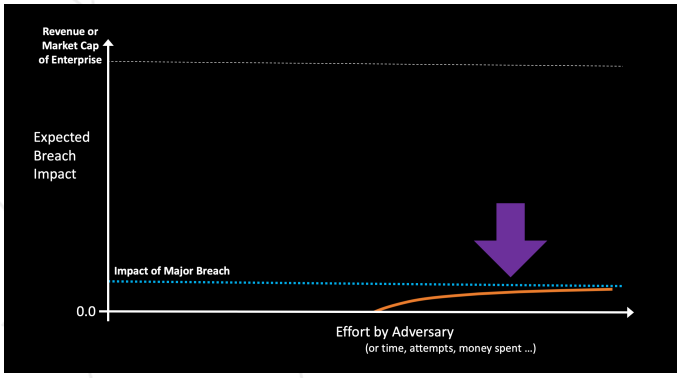

Fig 8: Improving cyber-resilience

Improving Cyber-resilience

Here is a systematic series of 5 steps you can take to enhance your organization's cyber-resilience and decrease breach risk.

1. **Gain 100x real-time visibility.** There is an adage which you are probably familiar with– *you cannot improve what you cannot measure.* Step 1 to improving your enterprise's cyber-resilience is to put a system in place that will discover and analyze your enterprise attack surface continuously and comprehensively and calculate resilience. Given the size and complexity of the enterprise attack surface described in Fig 1, this task requires some serious work.

For an organization with a thousand employees, there are over 10 million time-varying signals that must be analyzed to accurately predict breach risk. For organizations with 100K employees, we need to incorporate 100 billion+ signals in the risk calculation. Running a Nessus scan or getting a penetration test will not give you the visibility you need. Legacy techniques cover less than 5% of the attack surface and do not calculate mathematically sound measures of risk or resilience. You also cannot throw an army of people at this task– they won't be able to keep up. You need a specialized, self-learning system to do this scale of observation, analysis and calculation for breach risk and cyber-resilience.

2. **Strong identity and cyber-hygiene for your extended perimeter.** In Step 2, we focus on two things: robust mechanisms for identifying people, devices and applications in the enterprise (strong identity) and cybersecurity posture issues on your Internet-facing systems.

   Strong Identity: The objective of ***strong user identity*** is exactly what the name indicates– the ability to identify users who are trying to access some enterprise resource or application in a robust manner. The idea of robustness here is to make the identity system hard to subvert even when an adversary has guessed or stolen user passwords or lost devices, or been able to intercept/modify communications. Your strong user identity project must also take into account practical considerations such as the fact that your users will access managed and unmanaged applications, and some applications/services do not support mechanisms such as 2-factor authentication. You can't pretend that these "hard to control" legacy or shadow IT applications do not exist, or fit within your clean canonical managed IT architecture.

   Practical strong user identity can be established using an enterprise Identity and Access Management (IAM) product like *Okta*, combined with a password manager like *1password*. This will give you robust multi-factor authentication and policy control where possible and enable good password hygiene across managed and unmanaged applications.

   Strong device identity can be established by using client-side certificates, and ***strong application identity*** can be established by using server-side certificates. Good certificate management does requires some work and it is quite

IMPROVED CYBER-RESILIENCE IN 5 STEPS

1. **Gain 100x real-time visibility**
2. **Implement strong identity and secure extended perimeter**
3. **Secure the core and implement risk-aware security operations**
4. **Segment your network for dynamic risk-based access**
5. **Advanced cyber-resilience**

unfortunate that most organizations do not give it the attention that it requires. Mechanisms must be put into place to refresh expiring certificates and identify and educate users who click past certificate warnings.

Just because a system is behind a firewall does not mean that we can be sloppy about securing it.

Extended perimeter cyber-hygiene: Your Internet-facing assets are "ground-zero" where cyber-attacks start. In Step 1 above, you gained knowledge about cybersecurity posture issues involving devices, users and applications on your extended perimeter from the various vectors of Fig 1's y-axis. Typical examples include unpatched systems, missing or default passwords, port exposure, security configuration issues, easy-to-phish users, broken certificates, privileged users or high-value users who have poor cyber-hygiene, etc. You will want to fix all of these issues using appropriate mitigations including re-configuration, patching, rolling out controls for endpoint protection and anti-phishing, user training for phishing, etc. You should also think about how your fixes will stay in place, i.e., how you will maintain good cybersecurity posture for your extended perimeter as time passes. More on this in the next step.

3. **Improve Core systems cyber-hygiene; implement risk-aware security operations.** In Step 3, we will focus on three areas, improving cybersecurity posture for your non-Internet facing devices (core-hygiene), implementing risk-based vulnerability management, and automated, risk-aware incidence response.

Core systems cyber-hygiene: The premise behind improving the cyber-hygiene of core systems (which by definition are not Internet-facing) is two-fold. First, even though these systems are behind the firewall, it is quite likely that some device on your extended perimeter will get compromised at some point of time. Then, we don't want the adversary to to be able quickly attack and compromise a valuable core system that is not properly secured. This is an important reality to appreciate: just because a system is behind a firewall does not mean that we can be sloppy about securing it. Core cyber-hygiene can be improved using mechanisms similar to what you used to improve your perimeter cybersecurity posture.

Risk-based vulnerability management: It would have been great if once you fix your cyber-hygiene it would stay fixed. However the threat landscape shifts continuously with new vulnerabilities being discovered, exploits being created, and constant change in the enterprise in terms of devices, apps, users, configuration and behavior which results in new vulnerabilities. Risk-based vulnerability management is the process of continuously discovering and mitigating vulnerabilities across all your assets (Fig 1, x-axis) and all attack vectors (Fig 1, y-axis) in priority order based on risk.

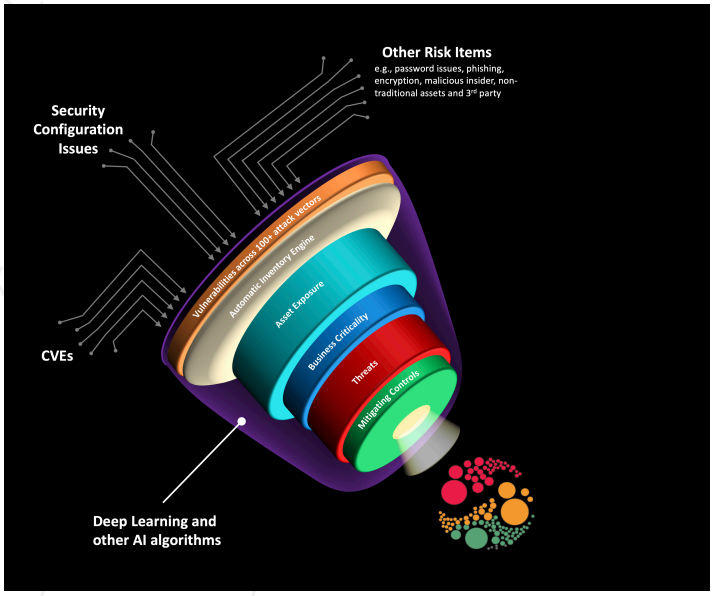

Fig 9: Risk-based vulnerability management

Risk-based prioritization of vulnerabilities is very important because otherwise the number of issues becomes too large to remediate efficiently. Today, a state-of-the-art implementation of risk-based vulnerability management (Fig 9) will take into account threats, exposure, compensating controls and

business criticality to prioritize vulnerabilities.

As you know, it is not enough to simply surface vulnerabilities. They need to be fixed. You need a system that identifies the correct responsible owners for each vulnerability and then generates prioritized tickets containing all relevant context and assigns them to these owners. Of course, some tickets will call for tactical mitigating actions, while other required actions correspond to strategic projects. Wherever possible, we want to automate these required actions, particularly for issues that will come up repeatedly, e.g., installing Windows patches.

It's not enough to just surface vulnerabilities… You need a system that automatically identifies the correct owners for each vulnerability and then tickets them, and tracks and nags them as required.

We also need to track progress and feed it back to relevant stakeholders. This is critical to get stakeholders outside the security team to contribute to improving cyber-resilience. State-of-the-art systems like Balbix take this to the next level, essentially gamifying the process of continuous vulnerability management and cybersecurity posture transformation for the enterprise, with the CISO and CIO as gamemasters.

Automated, risk-aware incidence response: Your SOC is the center of your real-time monitoring and response capabilities for security events. The top challenge dealing with the volume of security event data from all your enterprise assets that enters the analysis funnel efficiently, but without missing important events. Tools like SIEMs are you first filtering/collating mechanism, but are severely limited to what you can do with expensive consultants writing straight-line rules to describe context.

Deep Learning and other advanced AI techniques can continuously analyze 10s to 100s of millions of data points to mimic expert analyst knowledge and detective techniques. This can supercharge the effectiveness of your SOC by prioritizing activities based on risk, and providing deep context for analysts and automation tools. Of course, you must automate as much you can, and SOAR tools make this possible.

4. **Dynamic network segmentation.** After Steps 1-3 of improving cyber-resilience have been completed, you are ready to take on the task of modulating network-level access control for your extended network.

   The theory behind network segmentation is that an attacker cannot (practically) attack a system which they cannot route packets to, even if that system has some vulnerabilities. Network segmentation is equivalent to introducing a virtual air-gap between your more exposed perimeter systems and high-value critical systems. For example, in theory, you can set up your DHCP configuration so that your privileged users' laptops and desktops are given IP addresses in a certain VLAN while all other client endpoints get DHCP addresses in other VLANs. You can also configure your servers to be in a different VLAN and arrange your firewall/forwarding rules in your network infrastructure so that the administrative ports of your servers, such as RDP and SSH, are only routable from the privileged users' VLAN. All other VLANs cannot send packets to these ports of any device in the server segment. You can also prevent devices in the server VLAN from making any outbound connections. These are examples of north-south segmentation.

   Similarly, you can also modulate east-west traffic within a data center to prevent the adversary from moving laterally between your core systems.

   In practice, network segmentation is quite hard to implement because of the overhead of keeping up with changing

configuration and workloads in the enterprise. Second generation network-segmentation systems, e.g., Illumio, do a better job of this by giving you better tools to visualize workloads and manage rules to modulate east-west traffic. However, it is quite hard to use such systems to modulate north-south traffic. Setting up a bastion host complex is another way to modulate administrative north-south traffic.

Emerging third generation network segmentation systems bring the self-learning capabilities of modern AI algorithms to the challenges of dynamic network segmentation. With such systems it is possible to arrange the forwarding of traffic such that a) only authorized strongly-identified users and devices can route packets to specific systems, and b) only if the risk of such traffic is below acceptable levels. This is shown in Fig 10.

5. **Advanced cyber-resilience.** In Step 5, we can take the concept of dynamic network segmentation to the next level. For example, you can configure 3rd-gen network segmentation solutions such that your servers in AWS or in your data center are accessible from a client device only if the device has a specific class of 802.1x certificate, its user has gone through a full Okta verify within the last 2 minutes, the device is physically located within a certain geographical area (as measured by 2-3 methods), the device is fully patched and compliant with certain security configuration requirements, the user has good cybersecurity behavior– has not gone to poor reputation websites or clicked past broken certificates in recent time, and so on.

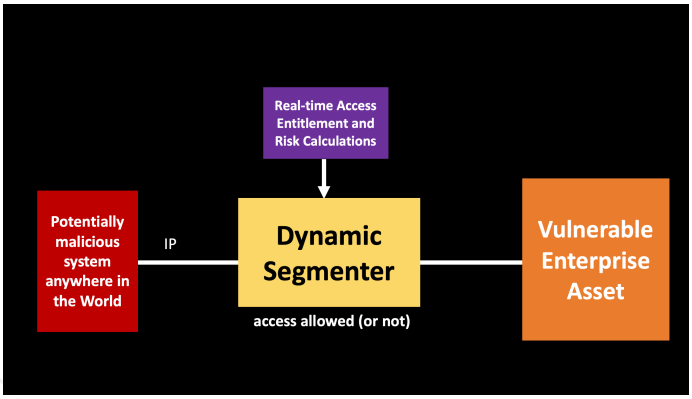

Fig 10: Risk-aware dynamic network segmentation

The general form of this calculation is shown in Fig 11. Before any IP packet is forwarded to a host, we can do a risk calculation based and then based on the results of the calculation and policy, allow the packet through, drop it, hold the packet while we do a out-of-band reauthenticate, or forward the packet to a honeypot or a subterfuge network.

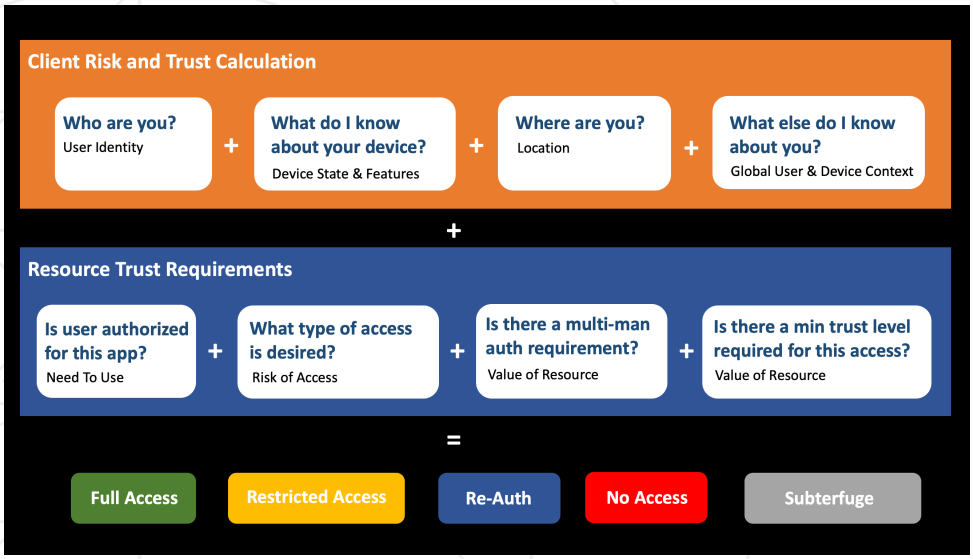

Fig 11: Risk-based packet forwarding and session management

Making cyber-resilience possible

The computation involved in supporting robust cyber-resilience is quite complex. The observations that are needed to feed into this math need to be gathered from all the entities in the extended enterprise network — users, assets, applications, their interactions, system and application configuration, human behavior, ACLs, mitigations, the state of all security products and intimate knowledge of the global threat model.

We created Balbix for this purpose. The Balbix platform makes continuous observations of your extended enterprise network from host, network and external vantage points funneled automatically through an AI backend to learn risk-related aspects of the network. Balbix builds a comprehensive and predictive, bottoms-up cyber-risk model which presents a clear picture of which users, apps and devices drive the enterprise's overall breach risk. The model incorporates information about vulnerabilities across 100+ attack vectors, threats, exposure, mitigating security controls, and business criticality.

Figure 12: The Balbix dashboard

You can see some examples of Balbix's dashboards in Fig 12 and 13. You can drill-down to the device level in these heatmaps, or use google-like natural language search.

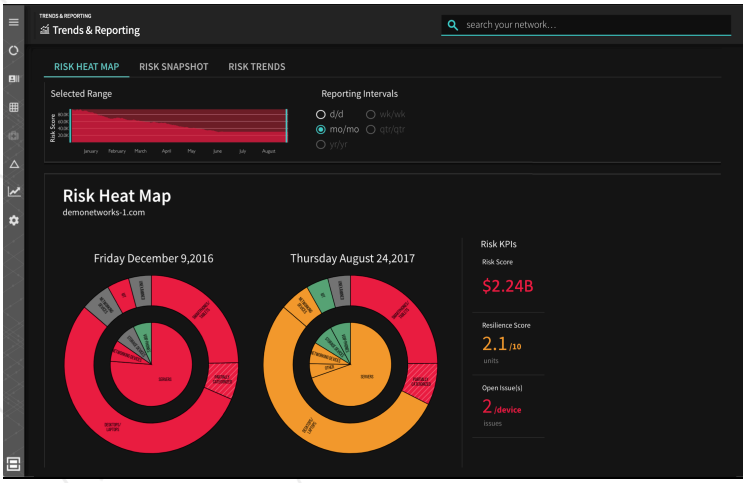

Figure 13: Balbix trends & reporting

Balbix also provides a prioritized set of actions that you need to take in order to improve your organization's cyber-resilience. The system has all the necessary integrations with your ticketing and other systems, with support for notifications, nudges, comparative reports, trends, etc., to help you gamify the process of cybersecurity posture transformation.

Please contact us at info@balbix.com to see how we might be able to help with your cyber-resilience objectives.

3031 Tisch Way, Ste 800
San Jose, CA 95128
866.936.3180
info@balbix.com
www.balbix.com