A Practical Guide to Implementing Zero Trust Architecture

How to extend the zero trust umbrella over all your assets in 8 steps.

GoodAccess

Table of contents

Executive summary

Today's society depends on information technology so much that it's hard to imagine we could function without it at all. Power industry, water supply management, banking, transportation, and a number of other fields and industries all rely heavily on IT for their functioning.

Small and medium enterprises (SME) face a growing number of threats in today's digital world. One of the most effective approaches to protecting critical resources and information is implementing the principles of zero trust in the company's security architecture.

Besides cyber threats, companies are also enduring the pressure of ensuring compliance with cybersecurity regulations, such as NIS2 and industry standards like SOC2, ISO/IEC 27001, and similar, whether this requirement comes from national governments or fellow companies seeking to strengthen their vendor supply chain.

Though zero trust architecture (ZTA) is a global trend, few know how to bring it to practice in a usable way. The biggest challenge is implementing zero trust architecture that encompasses the entire company; i.e., securing all its critical systems and networks. Especially nowadays, when it is common for organizations to have multi-site and multi-cloud infrastructures, where employees connect from any location and often using their own devices.

This whitepaper provides practical guidance to implementing zero trust architecture in a small or medium enterprise.

Why implement zero trust architecture

Threat protection

There's a good reason why zero trust architecture is a global cybersecurity trend. It allows organizations to create effective protection around their systems and data regardless of the location. It also addresses modern security challenges like insider threats while allowing organizations to build their networks more simply and efficiently, which is the key advantage that makes this an effective and practical approach in the real-world cybersecurity.

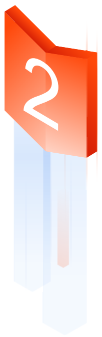

Full control over the organization

One of the motivators for implementing zero trust architecture is gaining full control over the organization's assets, no matter how complex its infrastructure. Zero trust forms an "umbrella" that extends over the entire organization and enables centralized access control and monitoring of all network activities. This approach ensures that the organization can manage and protect all its data and resources even in highly complex and diverse architectures.

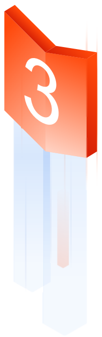

Compliance with legislation and security standards

By implementing zero trust architecture organizations can achieve compliance with security standards and regulations like ISO/IEC 27001, SOC2, NIS2, HIPAA and more. Zero trust not only protects against cyber threats, but also prepares the groundwork for obtaining certification and passing security audits.

When organizations implement zero trust architecture they automatically meet a large number of cybersecurity requirements that regulations and industry standards require, which minimizes the risk of legal and financial sanctions.

How to deploy zero trust architecture in your organization

The core idea of zero trust is that no network traffic can be implicitly trusted unless proven otherwise. Implementing a zero trust architecture in an organization is always a combination of technical measures (tools) and processes.

Below are a few steps that can be taken to reach this goal.

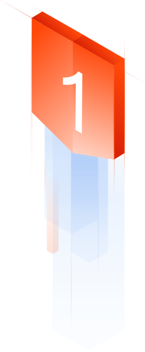

Analyze security risks and gaps

The first step is to take inventory of the organization's security, identify key resources, critical systems, and relevant risks. If the goal is to deploy a zero trust architecture, the next step would be to carry out a gap analysis that would identify vulnerabilities in the organization's protective matrix as opposed to the end-goal (ZTA), and subsequently compile an action plan and implement the necessary measures. The main steps include:

- Identify assets (data, applications, systems, networks, devices, users, etc.).
- Map the relationships between the assets, including data flows and network architecture.
- Understand the organization's context and business.
- Rank the assets by criticality (more below)
- Take inventory of existing security policies and toolsets.
- Identify relevant risks.
- Analyze identified risks.
- Rank the risks by relevance.
- Decide on measures to counter the risks.
- Compile a risk management plan.
- If needed, coordinate with the gap analysis.

Both risk and gap analyses can be carried in-house or outsourced to an MSSP provider.

Define critical assets and their networks

The goal of this step is to create an isolated network perimeter (Software Defined Perimeter, SDP). That means separating the critical systems from the rest of the organization’s IT and ensuring that only authenticated network traffic can access them.

But the question is, what constitutes a critical system? Best practice is to view the organization’s IT systems in terms of the criticality of the service they provide; i.e. to ask whether it is necessary for the system to function even during an attack (either to provide a service to customers or an internal service).

When building a zero trust architecture, a critical system can be any system that is network-accessible (locally or remotely) and identified by a URL or IP address, network protocol, and port, or a combination of protocols and ports (since web services are defined by a URL, TCP protocol, and ports 443 and 80).

Another example of a critical system is Microsoft Active Directory stored locally. This is a highly critical system that should be separated from the rest of the local network, hidden from the rest of the world, and only be accessible to authorized traffic. No company wants this system to be accessed by someone who visits the local office and connects to the Wi-Fi there.

PRO
TIP

How does GoodAccess do it? Using the user-friendly web interface you can name and define critical resources in the Systems section. The upside is that by defining systems there you also segment your network at maximum granularity (i.e. at the level of network services). This is because the way you define systems in GoodAccess also defines how they can be accessed. The goal is to close all access pathways to the systems except the GoodAccess gateway that handles authenticated user connections. This can be done in several ways:

SaaS systems

Restrict access to all IP addresses except the GoodAccess gateway(s) used by your organization (these addresses are dedicated to your organization). This ensures that the only way employees can access the service is via the gateway, and all the traffic generated from this will be properly authenticated. But for the rest of the world, the system is invisible and inaccessible from any network.

Cloud and data center systems

Though the above-mentioned IP whitelisting would work in this case as well, best practice is not to expose these systems and the cloud networks to the public internet at all. Ideally, there should be no public IP address and the whole network should be completely private and invisible to the internet at large.

Inside this network the organization must establish persistent tunnels between the GoodAccess gateway (or gateways in high-availability configuration) and edge routers/IPsec gateways or Linux instances within the private network. This is an easy way to create an IPsec tunnel that forms a permanent encrypted connection with the gateway. Such a network will become accessible and routable only via GoodAccess.

GoodAccess has a built-in Cloud & Branch Connector to create these connections.

Next, the organization can define systems within the private cloud network. This ensures secure zero trust authenticated access for the employees.

Branch systems
(in local private networks)

It is best practice to move all local systems to the cloud. Unless the system is truly needed on a local network, it is safest to remove it from the local network, where there are many network devices, users, and IoT equipment that is difficult to control and increases the number of attack vectors.

However, if any of the local systems are considered critical (e.g. production lines, CCTV, building management systems), these systems must be isolated in a separate network (e.g. a VLAN) where they'll be safe and access to them will be provided via a local GoodAccess gateway. This gateway will be deployed as a virtual instance within the local network.

An employee entering the network establishes a secure encrypted tunnel with the local zero trust gateway, which will grant them access according to the least-privilege principle to local critical systems. It is important to configure the VLAN so that the GoodAccess gateway is the sole access point to the critical system.

Remote zero trust access is then facilitated as described above. An IPsec tunnel between an edge router and a GoodAccess gateway in the cloud, enabled by the Cloud & Branch Connector. Here, the local network is perfectly concealed from the public internet, as the edge router does not have a public IP address.

The most complex situation occurs when the critical systems deployed on a local network cannot be moved to the cloud or a data center. Such cases require creating a VLAN that will house the critical systems, which would be separate from networks that contain network devices that do not need to conform to the strict zero-trust access policies (such as printers, servers, fax machines, and other non-critical IoT devices).

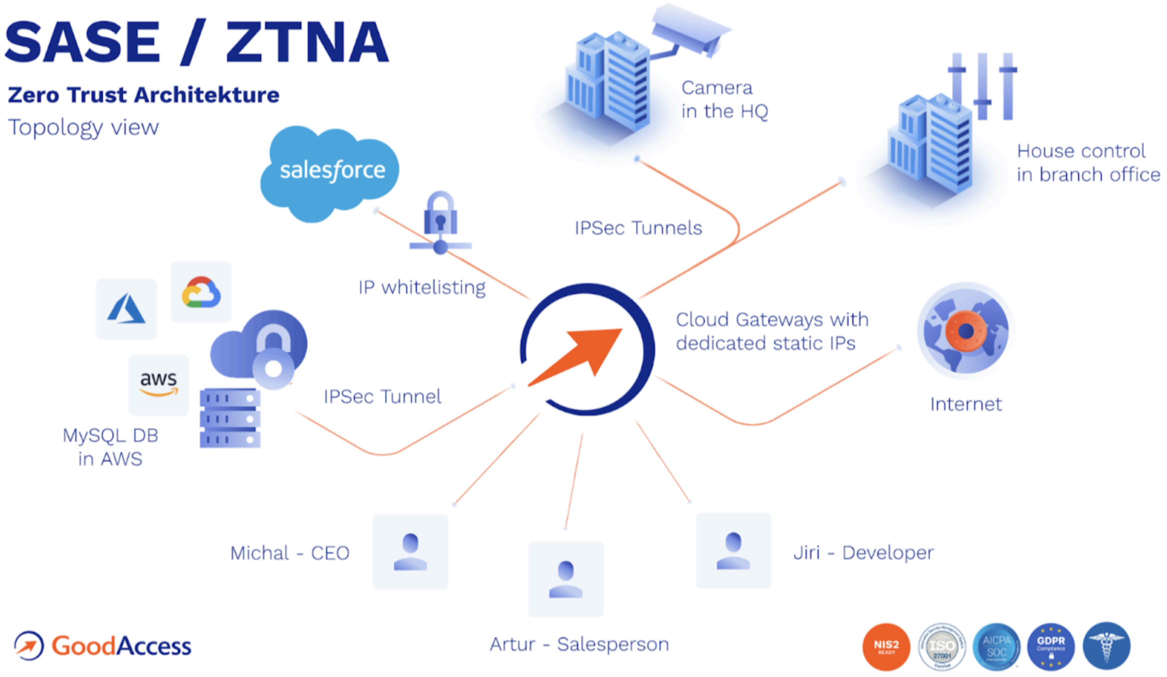

Fig. 2: Zero trust architecture – Topology overview

]

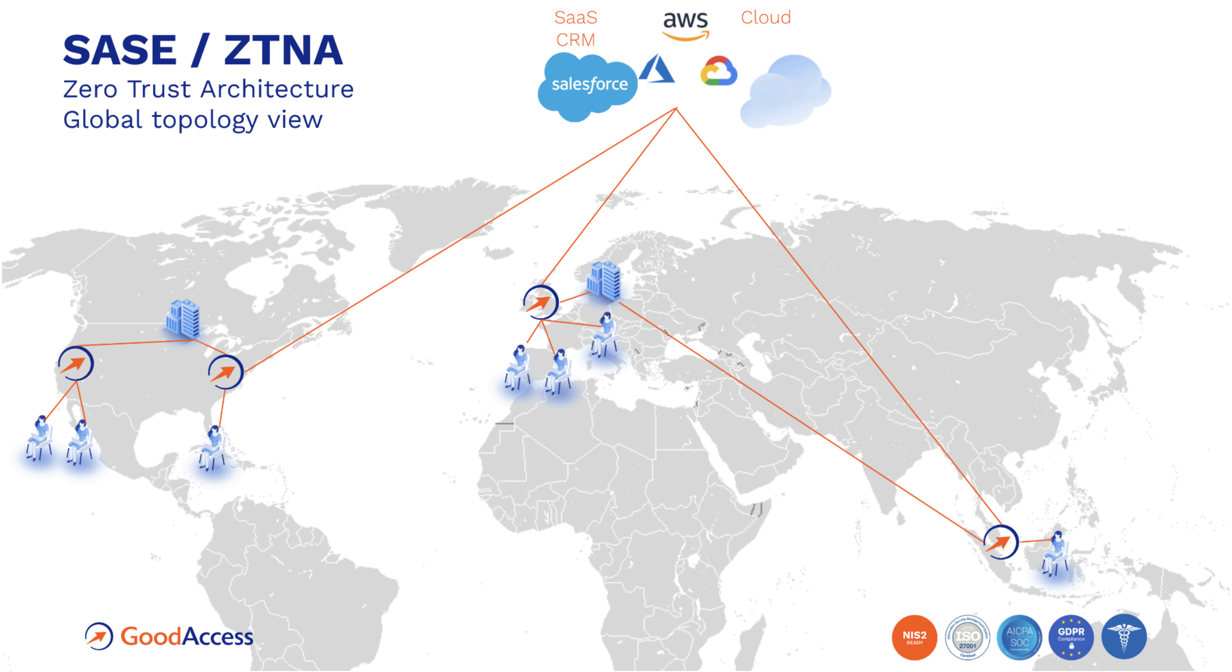

Fig. 2: Zero trust architecture – Global Topology View

Implement an identity management system (IAM)

An integral part of zero trust architecture is an identity and access management (IAM) system, such as Microsoft Entra ID, Google Workspace, Okta, and others. These systems ensure:

- **Central management of employee identities,** which eliminates the labor-intensive and error-prone manual access management, password rotation, and MFA on each system. In addition, it enables departments like HR to create employee identities or remove them during offboarding.
- **Multi-factor authentication (MFA)** on the application and network layer (VPN, ZTNA).
- **Grouping employee identities** and managing access privileges to each application per group.

GoodAccess can be easily integrated with all major identity providers or any provider supporting the SAML and SCIM protocols. We also collaborate with ProID, an identity provider that can combine physical identity with electronic identity and thus offer multiple options of multifactor authentication, such as hardware tokens, physical access cards, and more.

IAM integration enables full control over system logins and facilitates MFA on the network layer (ZTNA). It is best practice to configure the following additional measures:

- Multi-factor authentication for all users.
- Enabling single sign-on (SSO).
- Synchronizing users and groups of users between the IAM and GoodAccess using the SCIM protocol, which eliminates manual user setup in both systems. The SCIM protocol automatically carries all the account settings over to GoodAccess, allowing the administrator to manage all the identities and their access privileges in one place (IAM).
- Setting up identity-based access control that allows users to access only those systems they need for their work (see section 5).

We often encounter a scenario where a critical system does not support SSO or MFA. That prevents central access control and application-level MFA on this system, and it typically concerns older legacy systems.

In a zero trust architecture, and when ensuring regulatory compliance, we always combine access control and MFA on both the network and application layer. This ensures that security policies are observed throughout the organization. In other words, if application-layer MFA is not available, it can be implemented on the network layer, and vice versa. But, it's best to have both at the same time.

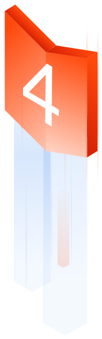

Define policy for employee devices

Even in mid-size organizations it's becoming virtually impossible to manually check every user device that connects to the company's assets. A typical employee uses several devices at the same time (3 on average), and the final number of devices becomes very high indeed.

The goal of this step is to ensure that all devices, whether company-issued or employee-owned (BYOD), meet the organization's security policy, and regularly and automatically check their security posture.

A majority of cyberattacks today start in employee devices, and when that happens, the attacker's activity can often be detected and immediately stopped.

The organization must determine what criteria describe a safe device. And based on that, the organization must set up rules for devices connecting to the network and enforce them automatically.

To do this, GoodAccess uses a function called Device Posture Check. Every employee device has a GoodAccess client app installed, i.e. a software agent. It not only establishes secure connections with the zero trust network, but also handles a number of other functions, including regular device posture checks and comparing the posture with the administrator's centrally defined security policy. The administrator can choose from a long list of parameters that they can require user devices to meet. For example, in Windows devices, they can check:

- Operating system version and latest update,
- Running and up-to-date antivirus software,
- Disk encryption,
- Windows Firewall,
- Screen lock protection,
- Domain,
- Registry settings,
- A specified file on the hard disk,
- A specified process running, and more.

PRO TIP

The policy settings take effect immediately. Employees whose devices do not meet the security requirements will be automatically disconnected from the zero trust network and must remedy the security posture of their devices before they can continue using the device.

Another case is when the device is compromised by an attacker who disables a security feature, e.g. firewall, to allow them to carry out their malicious activities. Device Posture Check will detect this change and disconnect the device from the zero trust network, which will cut off the attacker from the critical systems and instantly protect them from damage.

By allowing access only via the GoodAccess zero trust network ensures that critical systems won't be accessed by unauthorized personnel using unsecure devices.

5

Set up access privileges – identity based management

This step involves configuring access privileges for user identities that determine what critical systems they can access; i.e. identity-based access control. The goal is to render the systems invisible from the outside and prevent users from accessing them any other way but through the zero trust network.

In practice, this means defining what user identity (role) can access which specific system. Typically, users who work with the CRM will not have access to the development environment, while developers will not be able to access the CRM system unless their work gives them a legitimate reason.
Generally, best practice is to set up access privileges on both the application layer (IAM, SSO, MFA), and on the network layer (ZTNA). This ensures all systems are protected equally, including legacy systems that do not support SSO and MFA on the application layer.

GoodAccess handles this setup in the Access Control module using virtual access cards. They work just like physical access cards that we use to open doors, only virtual access cards grant access to selected critical systems and network segments. On the backend there is the FWaaS (firewall as a service) technology, which grants or denies authenticated network communication through the GoodAccess Gateway (cloud or local) depending on the configuration of the access cards.

6

Log network activity and access

Just like in a physical building all points of entry and movement within are watched and recorded, so it is in a zero trust architecture. The administrator needs to have records of access attempts to the perimeter as well as individual systems. Ideally, access logs will be sent to a SIEM, where they will be correlated with logs from the other parts of the infrastructure, e.g. endpoints, applications, and network appliances. This enables the detection of security events and anomalies.

GoodAccess offers several types of logs from the entire zero trust network:

- **Gateway-level access logs** – access records on the perimeter. Whenever a device connects to the perimeter, a log entry is created that shows who is connecting and when, what device they are using, from what location, and how much data was transferred.
- **System-level access logs** – records of access to individual critical systems on the network level. Here the zero trust architecture enables the centralized monitoring of all access to the critical applications regardless of whether the application natively supports log collection. These logs show who connects and when, what device they are using, from what location, and how much data was exchanged with the system.
- **Threat Blocker logs** – records of network communication with blacklisted domains, e.g. malware, phishing, C&C, etc.
- **Device posture check logs** – records of the security posture of devices and records of each regular check. These logs are automatically updated with any improvement/deterioration of the security status of all devices that employees use.
- **Admin logs** – records of all the changes made to the GoodAccess zero trust network configuration by the administrators via the Control Panel. Logs in GoodAccess cannot be changed or deleted, which ensures that no adversary can secretly make changes and cover their tracks.

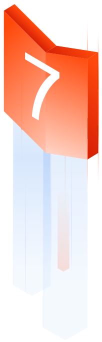

Automate anomaly detection and response

Response to anomalies is an important security component under a zero trust architecture. The general goal is to automatically assess network activity as accurately as possible and respond whenever a potentially dangerous behavior is detected. This can be achieved by API calls from a SIEM to GoodAccess, which can:

- **Block the user**, disconnect them from the perimeter, and immediately prevent their access to critical systems.
- **Block the device** that is behaving suspiciously (for example, a malware communication has been detected or its security posture decreases).
- **Isolate the critical system**, for example, when an unpatched vulnerability is found. In such a case, all access cards to the system can be invalidated, rendering it inaccessible.

8

Appoint a compliance officer and conduct regular security awareness trainings

Throughout an organization’s life cycle there are many changes to IT infrastructure, staff, devices, new technologies, and systems. All of this must be continually maintained as part of zero trust architecture, and it will not happen by itself.

Organizations must appoint a cybersecurity and regulatory compliance officer, whose task is to implement and oversee processes that ensure a stable level of cybersecurity and conduct regular cybersecurity training.

In an ideal world, users (employees) should not even notice that any zero trust architecture has been implemented, but in practice, they are likely to lose some of their access privileges, the way they log in to the network may change, and they will be required to adhere to certain processes. It is therefore necessary to explain the changes being made and train them to achieve compliance.

Conclusion

Zero trust is not one specific piece of technology, but rather a comprehensive approach to cybersecurity. In practice it involves the use of several tools at once, such as multi-factor authentication, device posture check, network segmentation, access control, traffic monitoring, logging, threat protection, etc. A complete zero trust ecosystem could look like fig. 3.

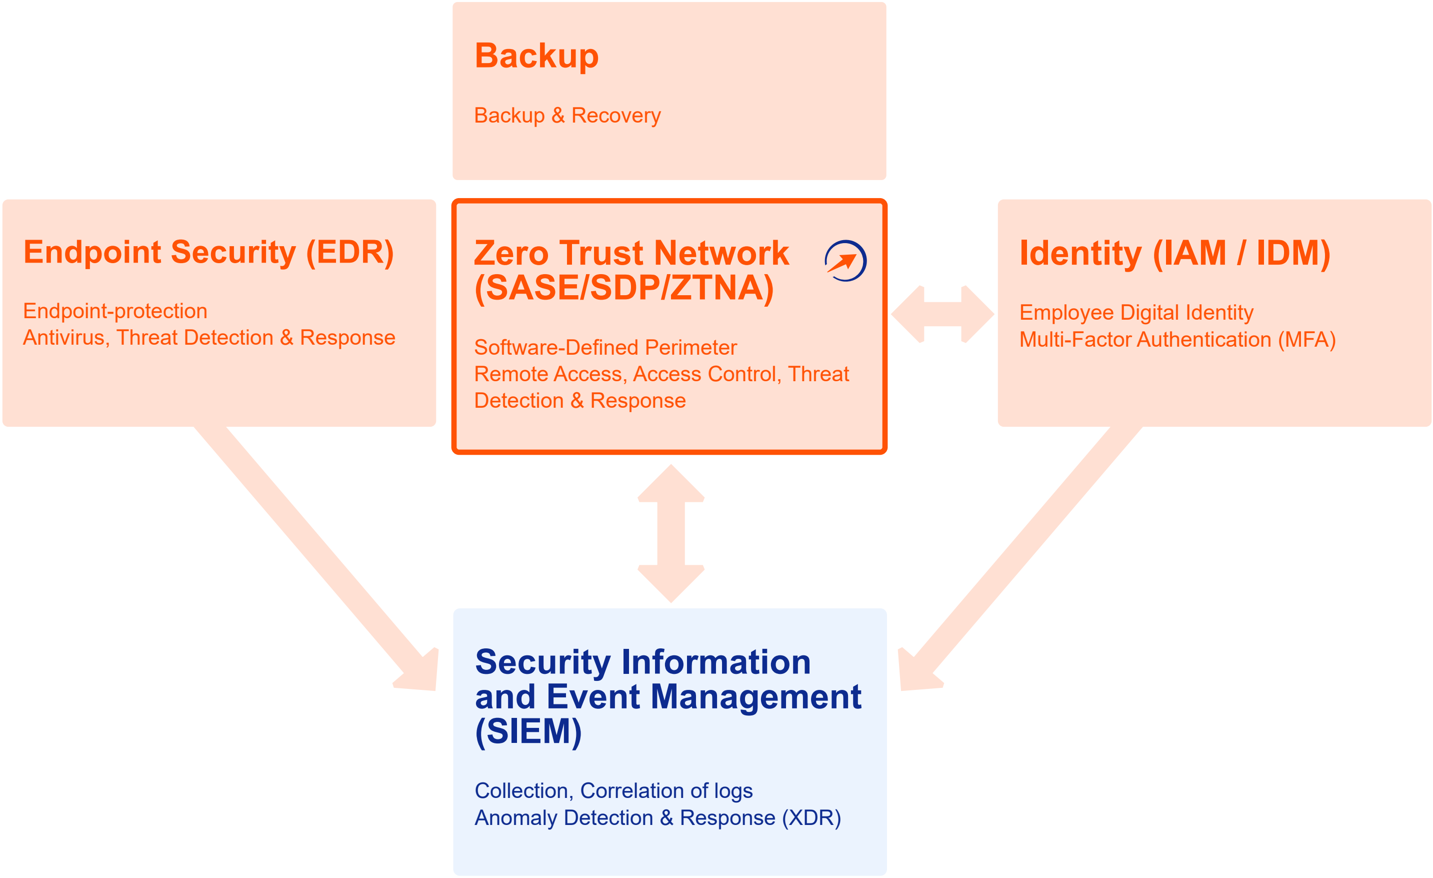

Backup

Backup & Recovery

Endpoint Security (EDR)

Endpoint-protection
Antivirus, Threat Detection & Response

Zero Trust Network (SASE/SDP/ZTNA)

Software-Defined Perimeter
Remote Access, Access Control, Threat Detection & Response

Identity (IAM / IDM)

Employee Digital Identity
Multi-Factor Authentication (MFA)

Security Information and Event Management (SIEM)

Collection, Correlation of logs
Anomaly Detection & Response (XDR)

Fig. 3: An example of basic security architecture

The eight points discussed above are a basic template for the successful deployment of a zero trust architecture. But the eventual composition will depend on the size of the organization, complexity of its infrastructure, existing security measures, and the abilities of the key personnel to implement the changes.

However, modern cloud-based solutions like GoodAccess make the deployment of these tools much simpler. They can significantly increase the cyber resilience of the organization, ensure regulatory compliance, while remaining user-friendly and able to automate many routine tasks. In addition, cloud-based solutions are infinitely scalable, which is equally important for the long-term usability of the solution.

If you are interested in a non-binding consultation on NIS2, or you wish to know more about the GoodAccess solution, please, contact us at sales@goodaccess.com.

Michal Čížek

E: ceo@goodaccess.com
W: www.goodaccess.com