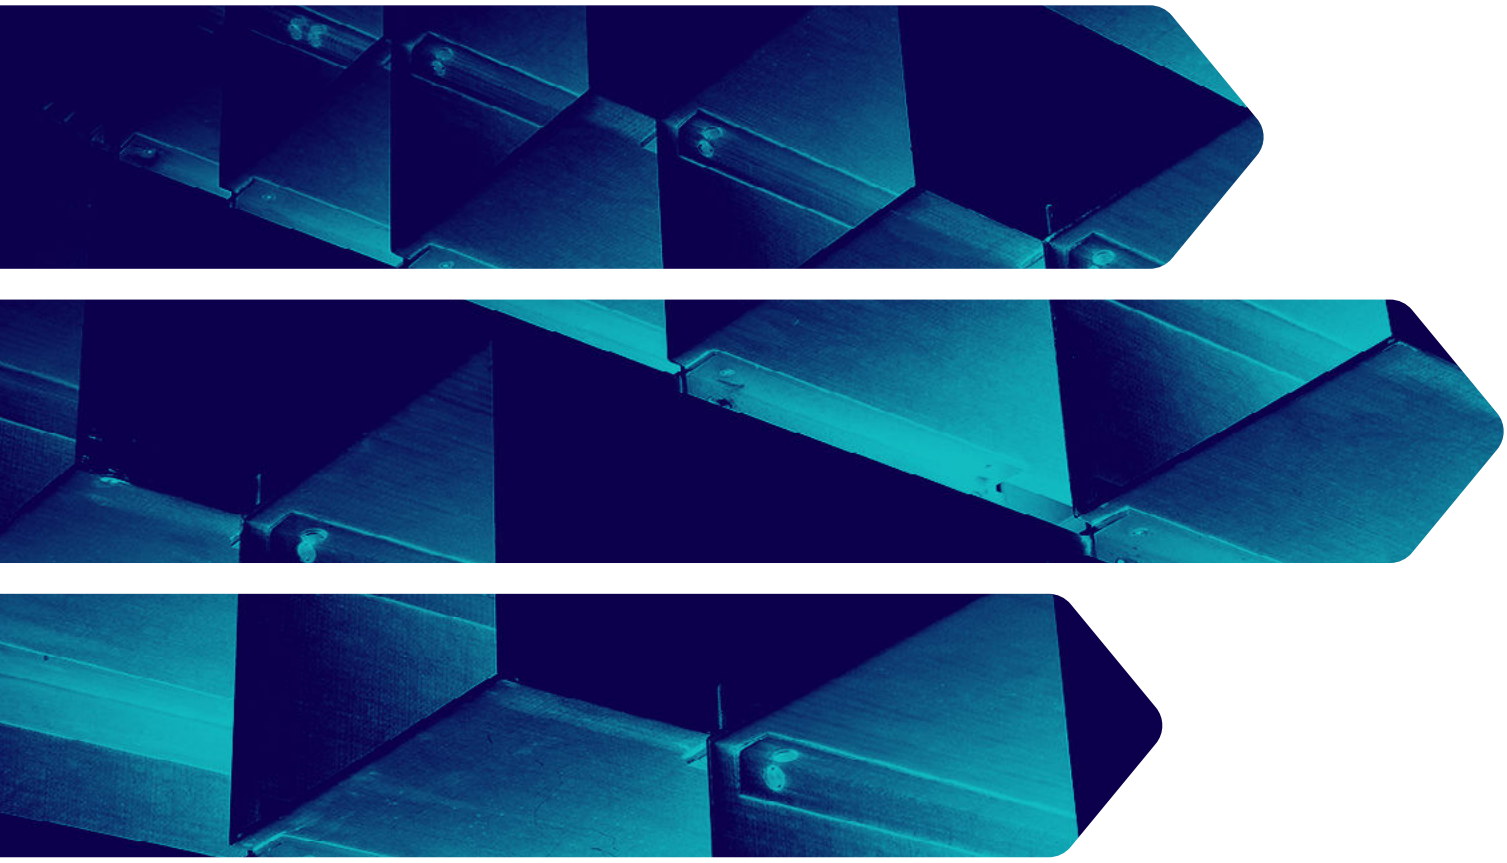

Buyer's Guide to

Active Application Segmentation for Ransomware Protection

TRUEFORT™

ActualTech MEDIA

Contents

2 Unprecedented Threats

5 **Discover:** What discovery capabilities are offered by the platform?

8 **Understand:** How do the platform's proactive capabilities improve your security?

10 **Enforce:** How does the system enforce security to enable real-time detection and response?

12 **Analysis:** How does the solution enable recovery and improve post-action data gathering for root cause analysis?

Unprecedented Threats

It's getting harder to keep your organization secure. Recent trends have created an unprecedented threat landscape.

Practical elimination of the defined network edge

The past decade has seen the eradication of neatly defined network perimeters. Today's computing environments are inherently hybrid - distributed across multiple clouds, data centers, and edge environments. In addition, in many cases, the rapid rise of remote work has further eroded the perimeter. Removing this once well-defined barrier has created new opportunities for bad actors to penetrate networks that were once considered secure.

The ongoing rise in coordinated attacks

Security incidents, once newsworthy, have become routine occurrences. It's so common that many CISOs now discuss them as inevitable and create preparedness plans assuming that an attack will eventually be successful. Today, while prevention remains a critical component of a security plan, preparedness around recovery has taken center stage. Today's attacks are far more common but more damaging to an organization, as ransomware attacks threaten target organizations with the public release of their data. This double extortion attack gets worse as attackers seek additional compensation through so-called *triple extortion* attacks that seek to extract payment from other potential targets by mining the stolen data and carrying out threats to release damaging information, which could harm the original victim's customers, partners, and thus reputation.

Increased attacker sophistication

Although organizations take steps to protect what remains of the network edge, attackers are finding ways to penetrate these once-praised bastions of security. Once an attacker has made their way through the first line of defense, it often doesn't take long before they find a soft internal target and secure this as a launching point for further breaches. Attackers begin lateral movements from this internal island to seek additional vulnerabilities to accomplish their objectives. With today's ransomware so prevalent, that objective is invariably the discovery of data hoards that the attacker can then ransom back to the victim. The goal of many security tools is to reduce the ability of an attacker to move laterally and to contain the impact of any attacks they may launch. However, due to the inability of many tools to identify a latent breach, it's often too late to contain the attack when other indicators of compromise appear.

The exponential increase in the number of potentially vulnerable endpoints

Clearly defined network and application endpoints that could previously be managed via static security rules have evolved into chaotic application environments with rapidly shifting characteristics that make it impossible to define "set it and forget it" rules. These dynamic environments are often infused with components from third-party software supply chains to enable critical functionality. However, these services also bring potentially unknown security vulnerabilities into the environment. There are numerous news stories detailing the severe impact that can result from software supply chain vulnerabilities that went undetected until it was too late. The modern application environment is a veritable stew of vulnerabilities that requires a solution that takes an application-centric and contextual view of activities occurring within and between applications, servers, workloads, and users.

Your application-centric security platform needs context from across the application environment to allow your organization to benefit from the defense it can provide.

These are just some trends that define modern computing and application environments. It's not possible to overstate the complexity of these environments. These challenges require a modern application security platform that legacy products can't adequately address. In this buyer's guide, you'll discover the critical questions you need to ask as you assess application security platforms for your organization.

Why focus on the application environment? Because that's where the most significant security gaps in many organizations exist, leading to devastating attacks.

The security of the application environment is directly impacted by other areas of vulnerability, including:

- Network vulnerabilities, including unpatched or outdated software, malware, and firewall misconfigurations
- Operating system vulnerabilities which enable attackers to exploit and gain access thanks to outdated operating systems, exploits such as buffer overflows, and denial of service attacks
- Human vulnerabilities include insider threats and insiders having the ability to access applications or systems that they shouldn't be able to
- Process vulnerabilities can include systems interacting with external systems, such as unauthorized FTP servers or payment processing systems

Your application-centric security platform needs context from across the application environment to allow your organization to benefit from the defense it can provide. Only a zero trust segmentation and workload protection that adapts to application behavior can keep pace with the dynamic characteristics of modern application environments.

This guide is broken down into four sections:

- **Discover:** The application and endpoint discovery capabilities of a platform
- **Understand:** The level to which a platform can proactively understand "who, what, where and when" of security incidents
- **Enforce:** How the platform takes automated steps to prevent the spread of an incident and how it guides administrators to high-value response activities
- **Analysis:** The ways by which the platform actively monitors and enables quick recovery and root cause analysis of an incident

Discover: What discovery capabilities are offered by the platform?

As mentioned earlier, today's computing environments span everything from edge environments to data centers to multiple public cloud environments and, increasingly, employee homes as remote work has taken hold. In addition, today's applications are a far removed from their ancient, monolithic cousins. Whereas the dinosaurs of the application world once stayed relatively consistent between patches and version upgrades, today's applications consist of thousands of ephemeral microservices and containers. This makes it impossible to manually map all the potential endpoints created and destroyed during an application's existence.

Application Architecture

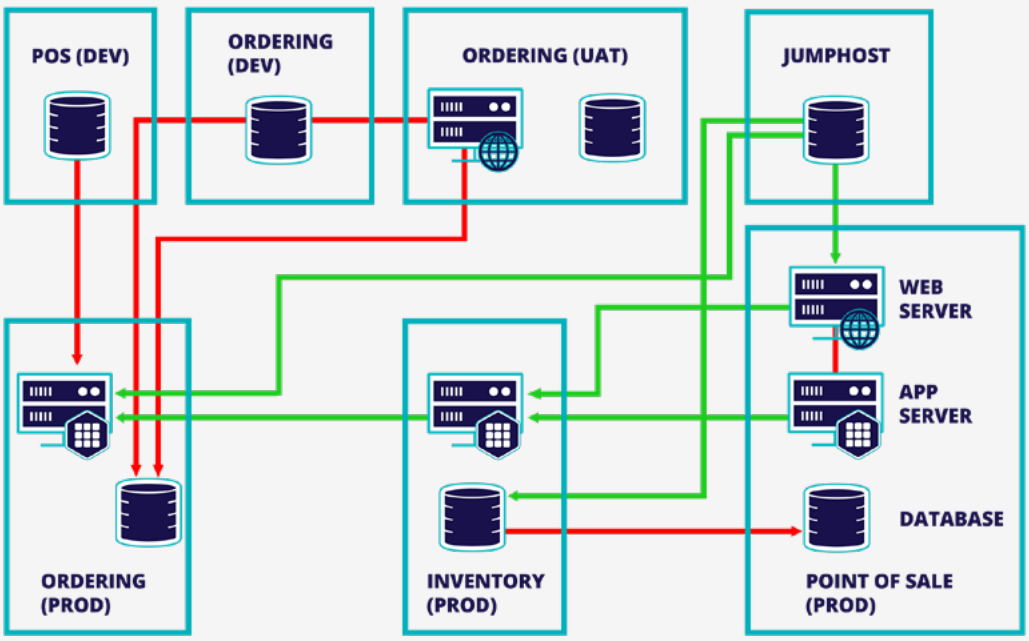

Figure 1: Note how core services, infrastructure, and business applications interact.

As you assess security solutions for these chaotic application environments, your focus should be on identifying a solution that can discover all the endpoints it needs to protect and what a normal operating state for that environment looks like. The standard operating state for an application includes understanding routine, expected behavior, and what anticipated movement between endpoints might look like. Most importantly, the solution must consider the importance of uptime for mission-critical applications and balance that need against security considerations.

If you have an environment of any significant complexity, likely, you have already deployed other agent-based security tools, such as CrowdStrike or SentinelOne. As you add more layers to already complex configurations, the "yet another agent" paradigm can become problematic as stacked agents reduce application performance, increase overall security complexity, and create silos in incident response. These silos can lead to wasted efforts when time is at a premium.

As you assess security solutions for these chaotic application environments, your focus should be on identifying a solution that can discover all the workloads and what a normal operating state for that environment looks like.

You may be tempted by the promises of many extended detection and response (XDR) tools available on the market, but make sure you do a thorough analysis. In many cases, XDR tools report what they see—they do not always perform in-depth assessments to determine what is normal behavior. So, while they may technically be doing their job, they are missing critical context, increasing the potential for false positives, or worse, not being able to identify an infiltration because the XDR tool does not know what "normal" actually looks like.

There are several questions you should consider as you compare and contrast security platforms.

QUESTION TO ASK	WHY IT'S IMPORTANT
Does the platform protect applications wherever they exist and in any form?	Applications operate across a range of environments, from the public cloud to the data center to the edge. They span physical servers, virtual servers, and containers. The platform you're considering should dynamically adapt to unusual activity to maintain protection across cloud, hybrid, virtual, containers, and traditional on-premises environments.
Can the solution detect and map live application interactions and relationships with network actions, servers, users, other applications to verify proper business policies and processes?	The ability to profile a contextual unified view of all application behavior including workload roles, network connections made, processes executed, and service accounts used for execution. This comprehensive view of the environment makes it far easier to understand normal behavior and to identify abnormal behavior that may require further investigation.

QUESTION TO ASK	WHY IT'S IMPORTANT
Is it possible to identify high-risk traffic to prevent unauthorized communications and prevent vulnerable interactions?	Even in environments in which principles of least privilege are strictly enforced, there are still privileges available—and at times, these may enable significant access depending on the needs of the application. A solution that can identify high-risk traffic can help protect you from an increasingly common attack vector and limit the potential damage from an attacker that begins to achieve success in lateral movement around the environment.
Can you verify required business-critical application and workload communications to prevent business disruptions?	It's important to ensure that business-critical applications aren't needlessly impacted by unending false positive alerts by forcing administrators to investigate alerts that lead to nowhere. Avoid false positives by using a positive security model to protect the behavior needed for optimal application performance and cease resource-draining activities.
Are you able to peer into the behavior of supply chain-derived software to reduce the potential for a vulnerability from impacting your environment?	Software procured through the supply chain is common in application stacks. However, since it's often pre-packaged, local developers may not always fully understand what constitutes normal behavior. Your security tool should have the ability to monitor such tools for their own baseline behavior and alert you if it begins to act in an abnormal fashion.
Does the tool enable you to leverage existing agents, such as CrowdStrike and SentinelOne, for discovery and behavior mapping?	Agent sprawl is a real challenge. Every agent you deploy requires resources to be diverted from serving the needs of the application. Reduce agent sprawl by leveraging existing agents for information gathering rather than continually adding new agents for every layer in the stack

Understand: How do the platform's proactive capabilities improve your security?

Mapping and behavioral analytics are important elements of a security solution, but they are only the first step in a long journey toward reducing the threat potential in a highly dynamic environment. While organizations may attempt to proactively gain visibility into application behavior before production deployment, as the saying goes, "no plan survives contact with the enemy," which is as true for deploying applications into production as it is for meeting enemy combatants on the battlefield.

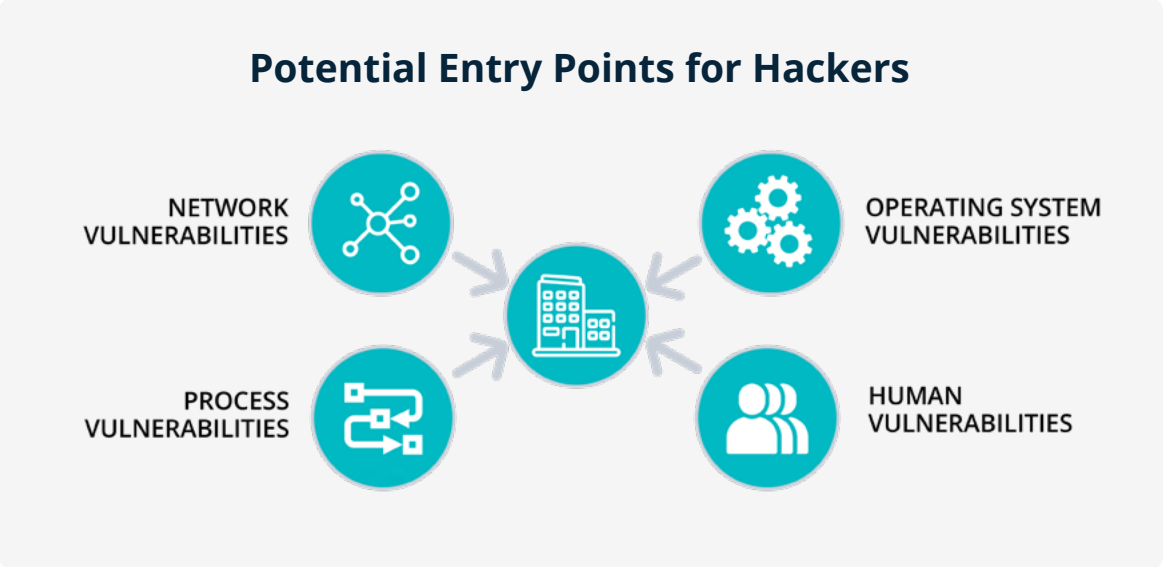

Figure 2: These are some of the most common vulnerabilities for an organization.

Application characteristics in test, development, and staging environments change dramatically in production as real-world users navigate applications in more ways than developers, and QA teams can possibly predict. As such, insights gleaned from pre-production environments are not reliable indicators for how production applications will behave. This is also true when it comes to application security. Remember—bad actors do not follow rulebooks. They seek weaknesses that may not even have been considered in pre-production. As such, it is vital to select a tool that can adapt to this real-world usage to help security teams proactively respond to potential attacks.

Mapping and behavioral analytics are important elements of a security solution, but they are only the first step in a long journey toward reducing the threat potential in a highly dynamic environment.

Historically, it has been common to use various segmentation techniques to reduce the overall attack surface of an application proactively. As applications have become more complex, and as environments have become increasingly distributed, many existing tools and network segmentation solutions have proven to have limitations when it comes to the applications they support. They simply were not designed to see or understand production application and their interaction within the infrastructure—the key context that is enabled only by a comprehensive understanding of applications talking to servers, other applications, and more. Many of these network-centered, IP-based tools do not understand applications and their workloads and how users interact with them, leaving gaping holes in the security layer.

Here are some questions you should consider around proactive defense as you evaluate new security tools.

QUESTION TO ASK	WHY IT'S IMPORTANT
Does the solution provide comprehensive workload support?	You can't afford a security solution that can't provide proactive capabilities for all your workloads—both what you run today and what you may add tomorrow. Make sure your solution supports whatever workload environment you do or could run, including Windows, CentOS, Ubuntu, Red Hat, SUSE, Kubernetes, and Docker services. Get ahead of tomorrow in a proactive way.
Is it possible to reduce attack surfaces for ransomware and other techniques targeting misconfigurations via advanced microsegmentation?	Dynamic microsegmentation with contextual understanding of how an environment is supposed to look makes it possible for a proactive security solution to reduce the potential blast radius from an attack, or prevent an attack altogether.
Does the solution notify security teams when there have been unanticipated file and configuration changes taking place in the environment?	A solution that understands the environment can more quickly identify unexpected file and configuration changes that are indicators of compromise. This behavior is unusual in nature, and the platform's ability to understand normal behavior makes it more likely that you will catch and stop ransomware and other attacks more quickly than otherwise possible.

Enforce: How does the system enforce security to enable real-time detection and response?

Traditional approaches to security often involved simple IP- or VLAN-based network segmentation. These now-archaic approaches were fine in a world where a well-defined edge existed, and where static routes and virtual machines ruled the day. In the current era in which tens of thousands of containers are created and destroyed daily, and the network perimeter is now a mirage, these static constructs no longer work, requiring a dynamic security platform.

Microsegmentation reduces an organization's potential attack surface, but even basic microsegmentation is not sufficient. The microsegmentation platform you deploy must support segmentation by application, account, and action. Essentially, you need to be able to establish policies to automate the creation and maintenance of a perimeter around your applications. This policy can block unnecessary network connections and disable privileged accounts being used incorrectly, even as the application environment changes.

The goal is to reduce the attack surface as much as possible. This means exposing fewer holes for network and application penetration from the start and reducing the potential for lateral movement around the network from an in-progress attack.

Most importantly, your platform should allow you to kill anomalous or rogue processes in their tracks even as they execute so that you can eliminate—or, at a minimum, reduce—the spread of an attack that may be underway.

QUESTION TO ASK	WHY IT'S IMPORTANT
Does the platform reduce the attack surface for ransomware and other techniques targeting misconfigurations via advanced microsegmentation?	Robust microsegmentation services enable real-time reduction in attack surface that limits the spread of ransomware and other attacks.
Can it harden systems to prevent unauthorized changes to meet compliance, auditor, and cyber insurance requirements?	Cyber insurance policies are increasingly strict and prescriptive, requiring organizations to adopt specific services and policies to even be eligible for such policies. A robust platform can monitor misuse of credentials and prevent potential human issues from reducing the overall security posture of the application environment.

QUESTION TO ASK	WHY IT'S IMPORTANT
Can the platform prevent anomalous or malicious activity from expanding, and is that prevention adapted to each application and workload environment?	A solution that only raises alerts while an attack spreads will allow far more damage than one that can take automated action against abnormal behavior in real time by compartmentalizing the attack. Make sure your chosen solution can block unnecessary network connections between applications, disable privileged accounts used on new workloads, and kill anomalous processes as they execute.
How does the platform enforce policy based on baseline behavior and alert administrators when it deviates?	The creation of a known-normal baseline should be the foundation of the platform. As services deviate from that baseline, administrators should be notified and, more importantly, proactive steps taken to reduce potential impact.
Can the platform immediately kill network connections, application interactions, workload processes, or user access to prevent threats?	Automated action in real time is the only way to truly prevent the potential for significant damage to the environment. The platform, based on its knowledge of the application environment, should be able to limit network connections and other interactions and user activity to prevent the spread of an attack.
Does the platform prevent false positive indicators from minimizing application availability by thoroughly understanding the environment and the context in which each service operates?	False positives lead to productivity and uptime-killing wild-goose chases that go nowhere. Choose a platform that thoroughly understands the application environment and surrounding context to prevent false positives from wasting time. More importantly, when you do get a hit on a more robust platform, you know that something is truly amiss, which leads to a different kind of urgency than one spurred by "yet another goose to chase" mentality.

Analysis: How does the solution enable recovery and improve post-action data gathering for root cause analysis?

Information security professionals need to understand what has happened so that appropriate recovery steps can be taken, and mitigation activities taken to prevent a future occurrence.

In the event of a successful attack, quick recovery is of paramount importance. For most organizations, downtime translates to adverse bottom-line outcomes, so there is generally a solid desire to get back in operation, but in a way that does not create further problems or erase the "paper trail" that can help shine a light into how an attack took place.

As you consider recovery, only solutions that provide deep forensic timeline insight and event correlation will maximize your efforts to understand how an attack took place. Discovering the correlation between services makes it possible to understand actions that led up to an incident, then provide insight into what triggered the incident. Very often, the culprit is not the most obvious trigger. Comprehensive correlation reporting helps you understand what was happening simultaneously, enabling you to hone in on the root cause of a breach quickly.

As you consider recovery, only solutions that provide deep forensic timeline insight and event correlation will maximize your efforts to understand how an attack took place.

When an incident occurs, the importance of understanding the reach of the incident cannot be overstated. It is the only way you can take the proper steps at the right time to mitigate damage to the organization and prevent drastic action, such as taking down applications not impacted by the breach and jeopardizing mission-critical operations.

QUESTION TO ASK	WHY IT'S IMPORTANT
When an incident occurs, does the solution provide the appropriate level of visibility and observability to understand what took place?	For recreating the "scene of the crime" for forensic analysis and understand the reach of a successful attack, deep insight into the application environment is a must-have feature.
Does the solution understand application context?	A solution that understands the context in which your applications operate makes it much easier to correlate incidents across all aspects of the environment, and more quickly determine the root cause of a breach.

QUESTION TO ASK	WHY IT'S IMPORTANT
Can the solution help you discover what applications and workloads were impacted by that incident?	Without knowing what applications and workloads may have been affected by an incident, it's impossible to not take down all workloads to make sure that they're not all impacted. No one wants to take down every application just because one was compromised.
How long does discovery take?	From the moment you deploy your new platform, it should be learning and adapting and able to provide baseline analytics within days. Once an incident occurs, it should be easy for the platform to report exactly what it knows so that remedial actions can be both timely and limited to just the affected parts of the environment.
Can the solution prioritize the downstream impact to other applications?	This is the entire point of a platform that offers real-time response to incidents based on a contextual and complete understanding of the environment. A platform that can take automated action to block abnormal activity can prevent downstream impact on other applications and ensure that damage and downtime is limited to just the affected applications.
Does the solution help you determine whether there was unauthorized access of your critical data stores or other sensitive information?	In an era in which it's increasingly common to need to report on the impact of an incident, understanding whether critical data stores were impacted is important. Imagine if you knew there had been an incident but had no idea what was impacted. You would likely need to assume the worst, even if the worst-case scenario didn't occur—you would have no proof otherwise. The results are additional cost, additional downtime, and additional reputational damage.

To learn more about putting yourself in control of lateral movement against ransomware, visit TrueFort at www.truefort.com.