REDSEAL

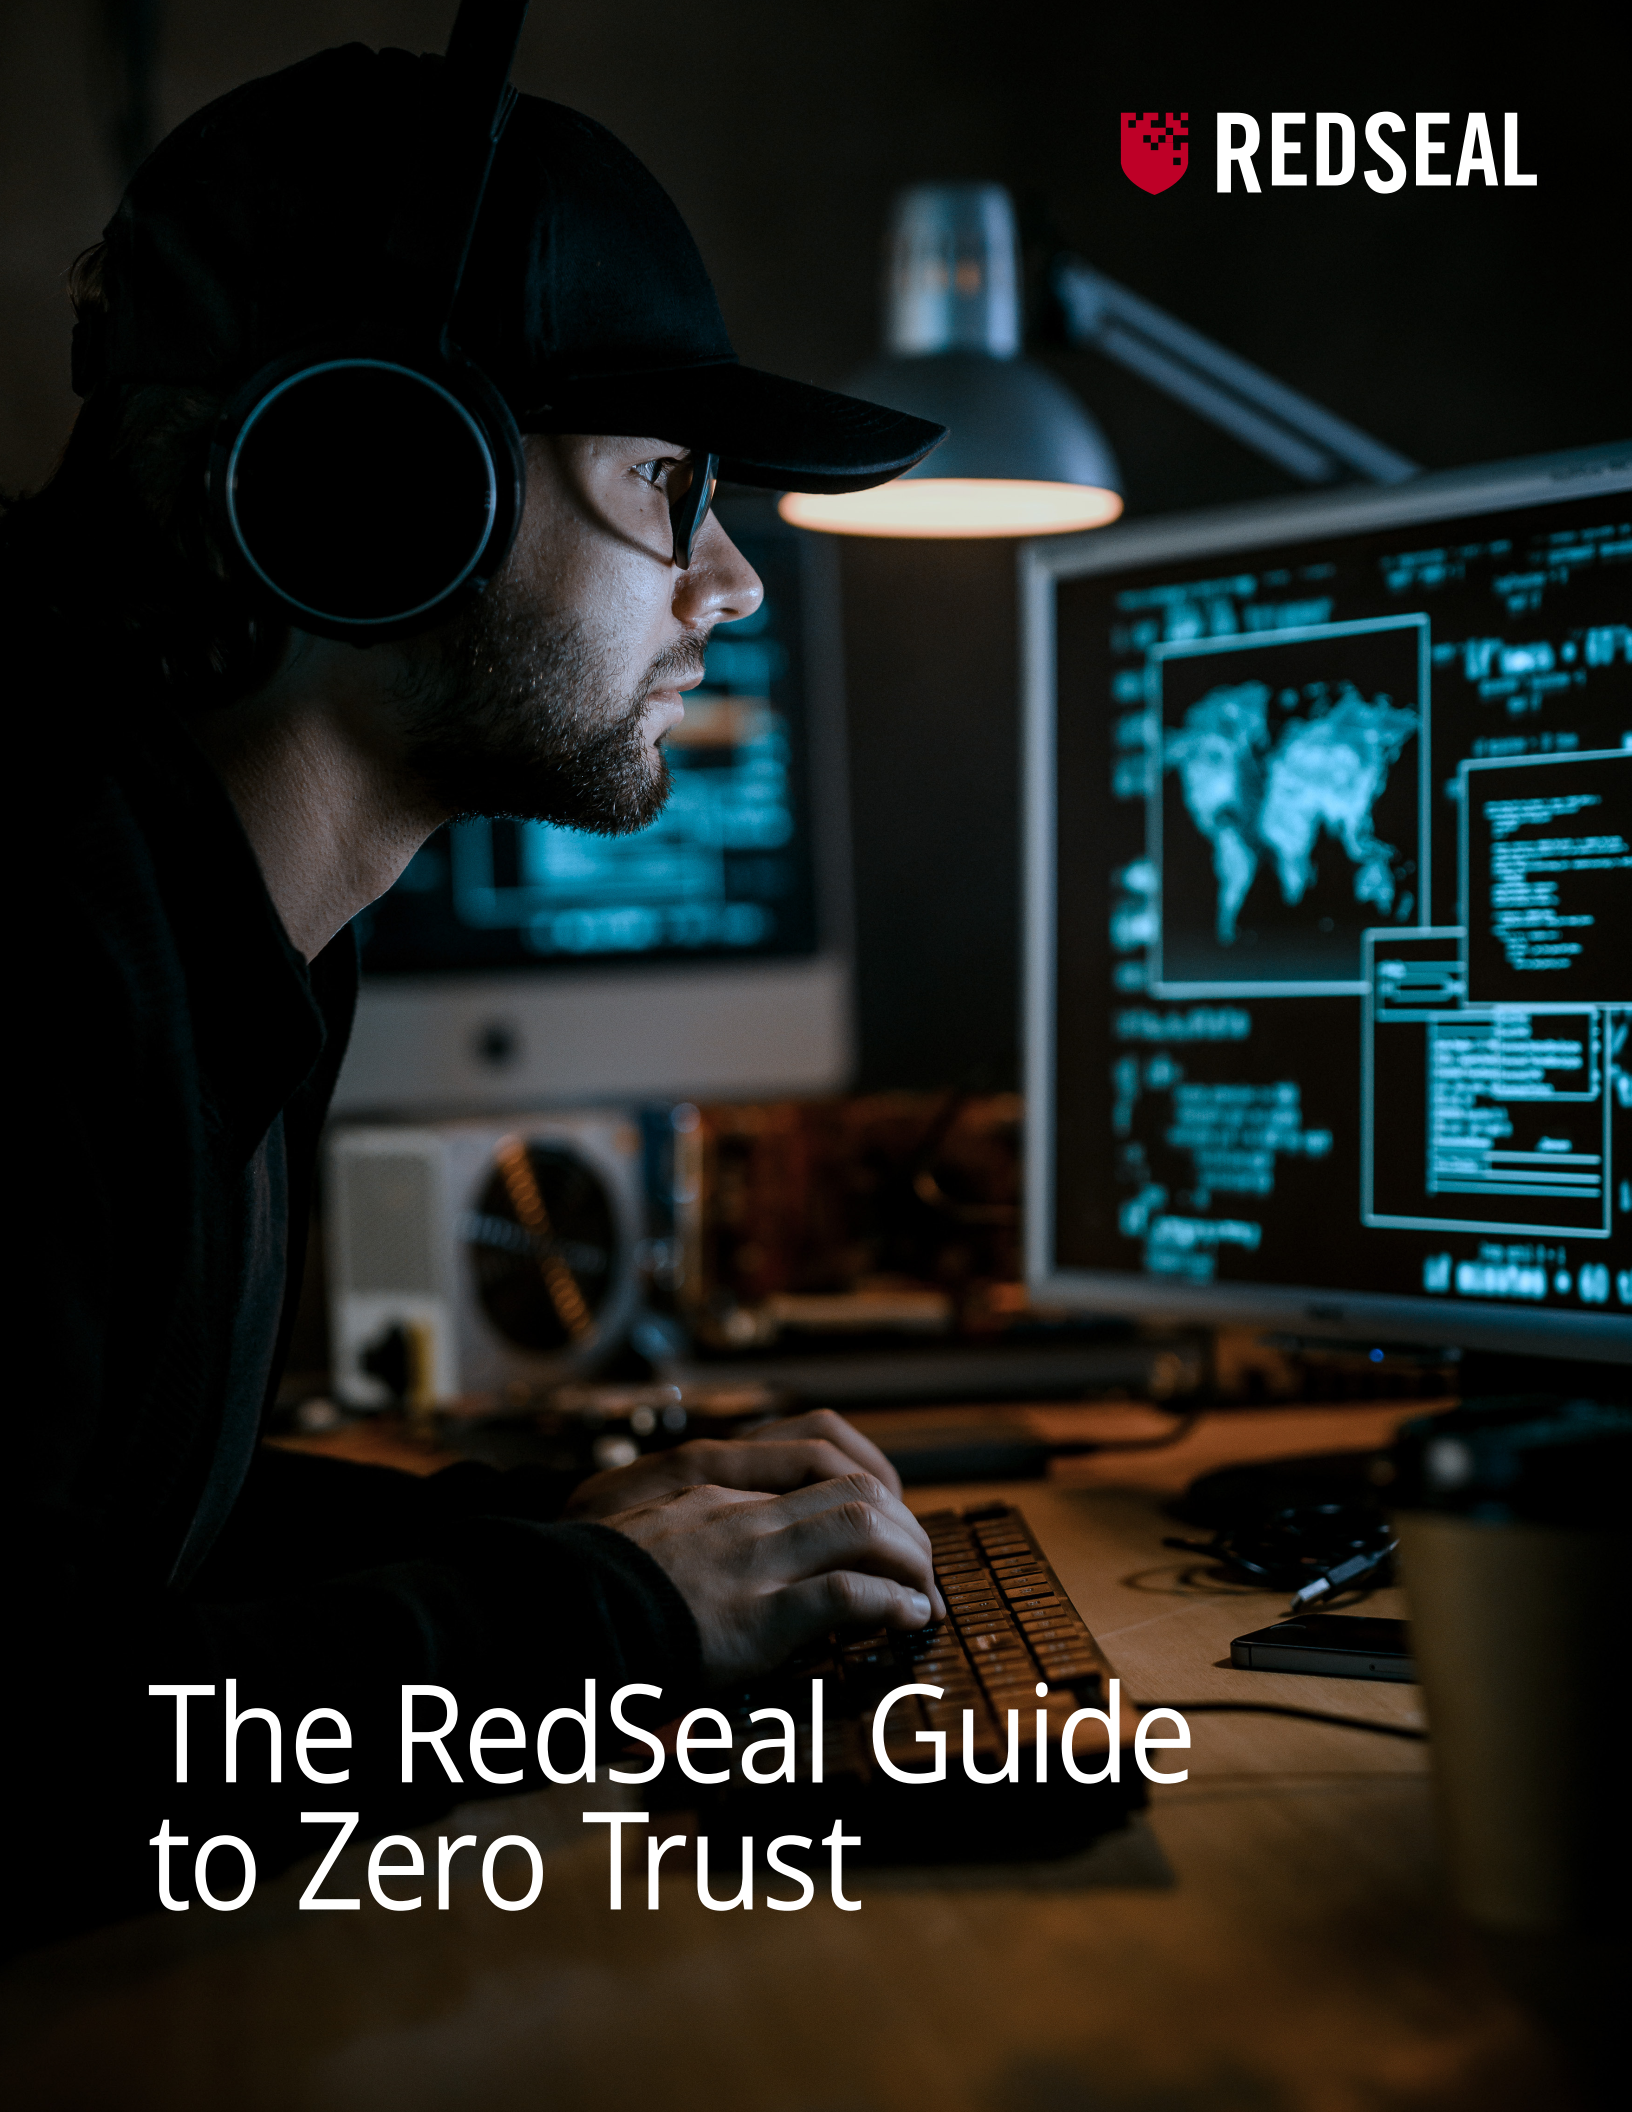

The RedSeal Guide to Zero Trust

Foreword by Wayne Lloyd

2021 will be remembered as a tipping point in how senior leadership views cybersecurity. In boardrooms and government CISO offices across the globe, the realization dawned that what happened in 2020 wasn't a one-off. In fact, it was the start of a new era of rapid cloud migration and digital transformation, mass hybrid working and an explosion in connected devices.

These trends opened up new fronts for threat actors which they were only too ready and willing to exploit. Recorded US data breaches were on track to reach record highs in 2021. Global ransomware attacks hit 470 million, another all-time high.

In 2021 global ransomware attacks hit another all-time high of

470 million

The days of paying lip service to cyber are over. Across the public and private sector, and from the President down, security is finally being treated as a strategic imperative. And zero trust looms increasingly large over these high-level discussions, thanks to a Presidential Executive Order issued in May 2021. Its requirements may strictly speaking apply only to federal agencies and suppliers. But the message is clear: all organizations should be moving in the same direction of travel to protect the nation, its people, its businesses and our way of life.

Yet getting to a place of zero trust won't be easy, especially given the complexity of modern IT environments. Today, 92% of enterprises have a multi-cloud strategy, while 80% have a hybrid cloud strategy. Legacy and digital, on-premises and cloud networks, sit side-by-side. In-house staff are struggling under the burden of securely managing this complexity amidst rapidly advancing innovation and iteration across multiple providers. Where should they begin?

It starts with visibility and control. Organizations must understand what devices, apps and services they have running on-premises and across cloud networks. This kind of inventory is the essential foundation of a zero trust approach which demands that organizations continuously identify who and what can access their most sensitive IT assets.

They must also be able to see all the various ways data flows through these networks. And they need to know how existing segmentation policies support compliance with regulations and best practice standards.

The following guide will describe this path in more detail, following the lead of the Office of Management and Budget (OMB)'s Moving the U.S. Government Towards Zero Trust Cybersecurity Principles memorandum. We all need to remember that zero trust, like any best practice cybersecurity, is not a destination but a continuous journey.

That journey starts here.

A View from the Government

President Biden's Executive Order gave agency heads 60 days to develop plans for implementing a zero-trust architecture, to help secure the wider push towards accelerated cloud adoption. It describes zero trust as follows:

"The zero-trust security model eliminates implicit trust in any one element, node, or service and instead requires continuous verification of the operational picture via real-time information from multiple sources to determine access and other system responses. In essence, a zero-trust architecture allows users full access but only to the bare minimum they need to perform their jobs.

If a device is compromised, zero trust can ensure that the damage is contained. The zero-trust architecture security model assumes that a breach is inevitable or has likely already occurred, so it constantly limits access to only what is needed and looks for anomalous or malicious activity. Zero Trust Architecture embeds comprehensive security monitoring; granular risk-based access controls; and system security automation in a coordinated manner throughout all aspects of the infrastructure in order to focus on protecting data in real-time within a dynamic threat environment. This data-centric security model allows the concept of least-privileged access to be applied for every access decision, where the answers to the questions of who, what, when, where, and how are critical for appropriately allowing or denying access to resources based on the combination of sever."

The OMB memorandum goes further, citing five key pillars to zero trust as described in the Cybersecurity and Infrastructure Security Agency (CISA)'s Zero Trust Maturity Model. These are:

1) **Identity:** Agency staff use an enterprise-wide identity to access the applications they use in their work. Phishing-resistant multi-factor authentication (MFA) protects those personnel from sophisticated online attacks.

2) **Devices:** The federal government has a complete inventory of every device it operates and authorizes for government use, and can detect and respond to incidents on those devices.

3) **Networks:** Agencies encrypt all DNS requests and HTTP traffic within their environment, and begin segmenting networks around their applications. The federal government identifies a workable path to encrypting email in transit.

4) **Applications:** Agencies treat all applications as internet-connected, routinely subject their applications to rigorous testing, and welcome external vulnerability reports.

5) **Data:** Agencies are on a clear, shared path to deploy protections that make use of thorough data categorization. Agencies are taking advantage of cloud security services to monitor access to their sensitive data, and have implemented enterprise-wide logging and information sharing.

Fortunately, RedSeal can help customers across three of these key areas.

Ensuring Complete Device Inventory

The conversation about zero trust is still too myopically focused on identity. While it is a cornerstone, there are many other bricks in the zero trust wall that need addressing. One other foundational block is device inventory. If zero trust is about allowing only certain authorized individuals to do certain things on certain devices, you need to first know about those devices—all of them.

This used to be easier in the days when inventory meant recording the serial number on the bottom of large machines. Today it's much harder because networks are distributed across on-premises and cloud infrastructure. And those "devices," or IT assets, may be hidden in walled gardens within clouds like Kubernetes. They're also moving around, all the time, in the cloud—dynamically provisioned into different locations.

So where do you start? You could try:

- Reviewing finance department licensing agreements
- Running network scans
- Deploying traffic monitoring solutions
- Talking to network "greybeards"

These approaches will get you part of the way. But shadow licensing agreements are far from comprehensive, tribal knowledge can get lost over time, and scanning tools will only identify devices on pre-defined IP space. As for monitoring tools, what about the traffic that doesn't pass through those monitors?

RedSeal works differently. We gather comprehensive data on the configuration state of all network- and cloud-based assets. These files will tell us if there's any missing IP space in our database, enabling us to present a comprehensive picture to federal security teams. Not only that, but we can detect if any of those on-premises or cloud infrastructure assets are misconfigured. As IT networks have grown in complexity, with execs able to spin up new cloud accounts in minutes, configuration errors of this sort are all too common.

Segmenting Networks Around Applications

If you're serious about zero trust, you need to anticipate the prospect of a network breach. The third "I" to consider after Identity and Inventory, is therefore Isolation, or segmentation. When done properly, network micro-segmentation can help to limit the blast radius of attacks, fencing in the threat actors so they can't reach high-value data and assets. Why is this important if you've already got identity controls watertight? Because there's no such thing as 100% secure. Bad actors are finding increasingly sophisticated ways to impersonate legitimate users and bypass authentication: just look at the SolarWinds attack.

Zero trust is ultimately not about protecting people or assets, it's about securing data. To that end, it makes most sense to segment at the application layer. It's often said to be easier to do this in the because with the cloud you can isolate workloads at the virtual machine or even application level using containers and pods. However, this becomes hard to manage at scale without automation of some sort. And in reality, few organizations are 100% cloud based. Take an aircraft carrier. Or an industrial control system. These kinds of IT environments will always need on-premises networks.

But as they're complex and growing fast, the chances of things slipping out of policy are high, with potentially serious repercussions. Eight or nine such incidents per week is not uncommon on some networks.

In federal government, there are various requirements for checking configuration settings at various intervals, based on the criticality of the information system. RedSeal can help you do that on a continuous basis, no matter how broad or granular your segmentation policies. We identify the "thing" you want to segment (server, workstation, app, container, pod etc). Then we continuously check the network is configured the way you want it, flagging to third-party tools of any compliance drift so they can be quickly and easily remediated. It works across on-premises and even the most dynamic of cloud environments. RedSeal makes sure that as workloads come and go, they only travel into the segments they're meant to, and nowhere else.

Discovering Internet-Accessible Applications

The zero trust mantra is that everything is compromised until proven otherwise. That's akin to assuming all the doors and windows of your house are open. By the same rationale, wouldn't it also be helpful to know specifically which ones are open, so you can close them? When it comes to your IT infrastructure, anything connected to the public internet should raise immediate red flags.

The challenge for modern security and network functions is that IT systems are increasingly siloed, with knowledge divided up between different members of the team. Do you know the difference between a Google VPC and an AWS VPC? How about an Azure Vnet? What about configuring a Juniper router versus one made by Cisco? These skills gaps are growing, and will continue to do so as vendors innovate, organizations invest and digital transformation continues. So how do you get a clear, holistic picture of your attack surface?

The answer is RedSeal. We consolidate all the information you need to know in a single place, offering a lingua franca across platforms to show which apps are exposed and why. Then your teams can investigate whether that's the result of a misconfiguration and if it needs to be remediated. It's work that would otherwise take many times longer if you have to consult different teams and knowledge silos and come to consensus. That's if you trust those individuals in the first place. Technology is moving so fast, skills gaps can appear out of nowhere and grow rapidly. RedSeal makes it simple.

Conclusion

In the private sector, marketing hype has if anything confused the zero trust message and the kind of approach end-user organizations need to take. Vendor X may claim its solution can get you all of the way there. But in reality, no silver bullet exists. There's much more clarity in federal government, where the zero trust imperative is clear, and the guidelines from NIST and the OMB are a great place to start.

Yet challenges exist in the public and private sector. There's still too great a focus on identity as the key to unlocking zero trust. While important, it's only one piece of the puzzle. Organizations need invested partners they can trust, to guide them on this journey. That means vendors capable of providing unified visibility across on-premises and cloud networks, into every single asset and how it's connected. Then comes continuous monitoring to ensure that when zero trust segmentation policies slip, you know immediately what to do, without needing to consult siloed teams.

This is the value of RedSeal. As cloud complexity builds, skills challenges persist and the bad guys get smarter, we provide a critical single source of truth—helping you to take those first zero trust steps with confidence.