



How to Apply Zero Trust at the Edge

govloop

Red Hat





Introduction

Zero trust security assumes that all traffic on your internal network is potentially malicious. In the edge computing environment, this becomes both problematic and especially necessary.

In security, a key concept in zero-trust architecture (ZTA) is the notion of the protect surface, the catalog of all assets requiring protection. With the rise of remote work and the explosion of Internet of Things (IoT) devices across government, that protect surface has expanded dramatically. In this environment, "it is critical to move from the castle-and-moat architecture of the past to address the needs of a hybrid cloud, and that includes the edge," said Michael Epley, Chief Architect and Security Strategist at Red Hat.

There's an urgent need to bring ZTA to the edge, where systems are subject to physical and network compromise to a much higher degree than non-edge-based systems. Applied at the edge, ZTA offers "a holistic security framework that addresses many pillars and capabilities of security," said Kenny Peeples, Principal Cybersecurity Architect at Red Hat. In this framework, government is able to:

- Authenticate all connections
- Identify all devices, users, applications and services
- Ensure that traffic goes only where it is needed, not just at the network level but also at the application level

In this report, developed in collaboration with Red Hat, we'll look at some of the key challenges to achieving ZTA at the edge and lay out a strategy for federal agencies looking to elevate their security practices.

Need to Know

If ZTA is to be implemented effectively at the edge, it's important for government leaders to understand two key concepts:

What makes zero trust different from the classic security model?

Zero trust is an approach to designing security architectures based on the premise that every interaction begins in an untrusted state. "This contrasts with traditional architectures, which may determine trustworthiness based on whether communication starts inside a firewall," Epley said.

More specifically, zero trust attempts to close gaps in security architectures that rely on implicit trust models and one-time authentication.

Zero trust has gained popularity because the global threat landscape has evolved, challenging long-held assumptions about the inherent trustworthiness of activities inside a network. Well-organized cybercriminals continue to recruit insiders and otherwise find new ways past the outer shell of traditional security architectures.

At the same time, "sophisticated hacking tools and commercialized ransomware-as-a-service platforms have also become more widely available, making it easier for new kinds of financially motivated criminals and cyber terrorists to operate," said Epley. "All of these threats have the potential to exfiltrate valuable data, disrupt business and commerce, and impact human life."

Given this new threat landscape, the federal government is shifting toward zero trust, with ZTA being mandated both by a May 2021 executive order and the March 2023 National Cybersecurity Strategy. Many other enterprises are weighing the costs and benefits of adopting this approach.

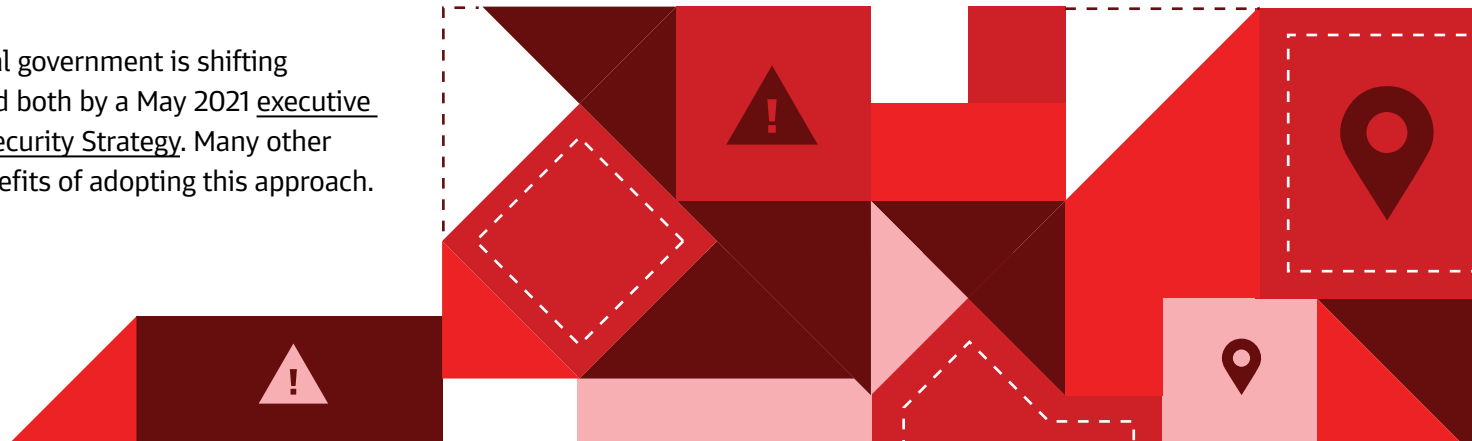
What makes edge different from the classic computing model?

Edge computing is computing that takes place at or near the physical location of either the user or the source of the data. "By placing computing services closer to these locations, users benefit from faster, more reliable services while companies benefit from the flexibility of hybrid cloud computing," said Peeples.

Edge computing offers the means for a government agency to use and distribute a common pool of resources across a large number of locations.

All this represents a change from the classic computing model, in which on-location users were tied to a central system and secured within a cyber perimeter. That model no longer applies across government: Remote work and the rise of IoT in support of myriad mission functions, among other things, have effectively dissolved that perimeter.

"The enterprise is now remote. Devices are deployed at the edge, and yet they still need access to enterprise services and resources," said Peeples. "This is a change from the classic model where all the dependent services and resources are in the same location."





How to Bring Zero Trust to the Edge

For government IT teams and business-line leaders looking to implement ZTA in support of both remote workers and IoT devices, four key issues need to be considered.



Distributed ID Management

Identity and access management are key to ZTA, and yet these present special challenges at the edge.

A core principle of zero trust is that we don't trust identity assertions that are provided by the actor. We assume breach. In this case, that means assuming that the edge system's self-asserted identity is potentially invalid or incorrect — that the assertion does not match the reality.

But that's a complicated proposition in practice. "Edge systems are composed of lots of independent instances of systems that are deployed in hostile environments," Epley said. "At the same time, transactions between actors and resources in a zero-trust architecture must be individually ascertained, which requires we know the actor and resource involved."

Further, most identity systems rely on an actor holding a credential that can be misused. That becomes particularly problematic in the edge environment, where there is less centralized control available to manage those credentials.

How do you close the gap between the architectural reality and the security imperative? "We have to devise methods to identify an actor at the edge that don't rely on the actors themselves, or at least not exclusively," said Peeples.

Distributed identity management solutions offer a way forward. Such systems rely on multi-part approaches: They might leverage blockchain (an immutable record of digital events) as well as "merkle trees," data structures that enhance the efficiency in blockchain.

"There are a number of distributed identity systems being developed today, and they represent one approach to managing this problem," Epley said.

Automation of IoT Fleet Management

With hundreds of IoT devices to inform everything from inventory management to military strategy, agencies face a significant management challenge. It's not only the vast numbers of devices, but the variety.

"There's the problem of scale, with just a very large number of devices to manage. And because it takes a long time to build out these networks of IoT systems, there's often a large diversity in the implementation," said Peeples. "Each individual device may have different software, different hardware, different sensors, different generations of what they can and can't support."

That makes it hard to implement zero trust effectively, particularly as these devices operate according to automated routines. "It's no longer about whether you can trust the human; it's whether the hardware can trust the software, or the software can trust the hardware," said Epley. "Can the systems trust the cloud, or other systems that they're interacting with?"

To ensure trusted connections, agencies need to be looking to automate their identity solutions across the IoT fleet. Fleet-wide automation of a distributed identity and access management solution would help define allowable interactions and reduce risk across the entire system.

"You need an identity and access management system that allows for flexibility, giving the diversity of that IoT fleet, and the fact that these interacting systems will change over time," Epley said.

Such an approach would support the requirements of the OMB Memorandum on Zero Trust and the Defense Department's Execution Roadmap, both of which call for organizations to inventory users and devices, and to proactively manage access to devices.

"With automation, we can do the configuration-as-code across all those diverse devices, and be able to track the configuration drift," Peeples said. "That's what will allow agencies to effectively manage across such a large and changing landscape of devices."

Business Process and Policy

From a business-process perspective, zero trust demands the alignment of access controls to specific business needs. In the case of edge systems, those needs are highly dynamic.

"Mission requirements may change, or the edge device might literally be moving: Think of satellites orbiting the Earth. The business function that it's supporting might depend on literally where it is at a given time," said Epley.

Suppose that a satellite needs to upload new commands or download new data, based on its position. That's a fluid mission set, and it will require business processes and policies that are more fluid than those that govern more conventional computing capabilities.

"We need to think about how we express those policies in a consistent fashion, and in terms of zero trust, we need to match those policies to the associated access controls," Peeples said. In the satellite use case, for example, there will be policies governing when the device is in "receive" mode, and access controls associated with those policies. "It should not be able to receive a command when it's over some random part of the Earth, such as an area that we don't control," he said.

Traditionally, there hasn't been much need to define fluid or dynamic access controls based on changing mission or device parameters. But that's just what will be needed in order to make zero trust effective in the edge environment.

"It's a business challenge, and a technological challenge as well: figuring out how to make the access control systems understand and be responsive to those policies," said Epley.

Agencies will need to do some risk analysis to understand what policies are actually needed, and they will need to focus on the engineering aspects to ensure they are putting in place systems that can effectively support those dynamic access control policies.

Integration With Cloud Infrastructure

As the government looks to support zero trust at the edge, planners will need to be thinking hard about the role of cloud in supporting those widely deployed endpoints. That means they'll need to pivot toward a more holistic view of the situation.

"We have to stop thinking of edge and cloud and IoT as separate," Peeples said. "The practical reality is that these systems are going to be interoperating with each other. That means we will be making the zero-trust access-control decisions across our cloud infrastructure, as well as our IoT and edge systems."

Without effective security controls across that entire chain of interactions, he said, agencies risk losing the availability of mission-critical edge devices.

With edge data increasingly tied to the cloud, security efforts also will help ensure the validity of that vital information. "With locations being geographically dispersed, we have to think about the integrity of the data," said Peeples. "There already is guidance that has come out from the government that requires agencies to deal with encryption, for example. They'll need to be validating the source and the destination of that data."

Zero trust will play a key role. "We will need security mechanisms that enable these edge devices to talk securely to the cloud," Epley said.

Cloud exists outside the traditional IT perimeter, and conventional firewalls are of minimal use in that ecosystem, "and yet that is going to be a central place where they are sending information," Peeples said. "So we need to ensure the security around that. Our zero-trust policy has to extend into that arena."

There is no perimeter in these cloud-dependent edge systems, "and that's exactly why zero trust is so important here," said Epley. "Given the scale and diversity of these edge systems, you need the continuous validation that zero trust requires."

Zero Trust at the Edge: Solutions Spotlight

No one product will make ZTA a reality in support of federal edge computing efforts. Rather, leaders need to examine their inventory devices as well as their current security posture across the enterprise in order to build out a roadmap of capability priorities.

Red Hat works with the Defense Information Systems Agency to apply DISA Security Technical Implementation Guides to Red Hat offerings. Red Hat, in turn, has developed Compliance as Code playbooks and/or scripts to help remediate findings.

Red Hat Enterprise Linux

Red Hat Enterprise Linux supports and powers software and technologies for automation, cloud, containers, middleware, storage, application development, microservices, virtualization, management and more. The latest release, RHEL 9, is designed to meet the needs of the hybrid cloud environment, a necessary step as cloud is integral to the support of edge devices. In terms of zero trust, RHEL 9 supports identity management tools and services to help you centralize identity management, enforce security controls and comply with security standards across your entire environment.

OpenShift

The OpenShift Container Platform is a platform for developing and deploying containerized applications. OpenShift Container Platform documentation helps users understand OpenShift Container Platform development, from simple containers to advanced Kubernetes deployments and operators. Integrated with the tools needed to implement ZTA in less time and with less risk, it enables agencies to easily create projects in order to organize and share the solutions they develop, as well as applications.

Ansible Controller

The Ansible automation controller provides dashboards and centralized logging to manage and track job execution from a single location, enabling collaboration across teams. Users can define, operate, scale and delegate automation across the enterprise with a centralized web user interface (webUI) and application programming interface (API). In support of zero trust, Ansible allows administrators to enforce access controls to assign permissions, privileges and roles to users.

To improve their security posture, agencies need to address CISA's ZTA Maturity Model, and may also need to follow the DoD CIO execution roadmap.

Red Hat products help them to get there, with RHEL supporting infrastructure security, and Openshift and Ansible enabling automation and orchestration. In support of greater compliance, agencies can leverage specialized tools in Openshift such as Compliance Operator, Advanced Cluster Security, GitOps and Pipelines. A single hybrid cloud platform used to build, run and manage applications at scale, OpenShift integrates with the security measures needed to implement ZTA faster and with less risk.





Zero Trust at the Edge: Use Cases

Three potential use cases help to demonstrate the potential power of ZTA in government edge deployments.

1. Mission Infrastructure

The military increasingly relies on edge devices as a key part of its mission infrastructure. It's looking to field sensors, for example, to help build a more complete and accurate battlefield view. The difficulty is one of scale: how to manage hundreds or thousands of devices over a broad geographic area.

Edge management strategies "can help to securely manage and scale these deployments, providing control and peace of mind at every stage of an edge system's life," said Peeples. With a modernized ZTA implementation, it's possible to achieve zero-touch provisioning, with system health visibility and quick security remediations available from a single interface.

Robust tools help ensure the secure deployments and operation of these mission-critical devices. "Using Image Builder, for example, IT teams can quickly create, deploy and easily maintain custom, edge-optimized OS images over the life of the system," he said.

Image Builder, provided by Red Hat Enterprise Linux, contains everything needed to run edge workloads in specific places. These images can be customized for unique edge workloads, helping keep edge deployments consistent, scalable, compliant and more secure.

With proactive management of device security, the military can make use of these devices to get a fuller view of the battlefield without introducing unnecessary risk.

2. Operations Automation

Federal agencies are looking to edge devices to move processing power closer to data sources. The edge allows them to deploy latency-sensitive applications, gather IoT data and support facilities operations even when a connection to the data center or cloud is lost.

Without local IT support staff, however, it's hard to ensure the secure operation of these devices. "Limited staff are often forced to deal with incompatible technology at sites that vary in size, network connectivity and energy use," Epley said.

To support zero trust across these key operational-IoT deployments, agencies can leverage the power of an automation mesh, an overlay network intended to ease the distribution of automation across a collection of execution nodes, using existing connectivity.

With zero trust as the backbone securing the connections, the automation mesh creates peer-to-peer links that make it possible to deliver timely, scalable support for operations-critical IoT deployments, even in the face of latency and network disruptions. With automation across environments, operations teams can standardize configuration and deployment from data centers, across clouds and all the way to edge locations.



3. Radio Access Networks

Agencies are looking to telecommunications service providers to implement cloud-native radio access networks (RAN) as part of a 5G mobile network. That 5G workload may play a critical role in support communications with mission-critical edge deployments.

“One crucial detail of this use case is ensuring that this workload can be ‘autonomous,’ so that it can continue operating with its existing configuration, even when any centralized management functionality is unavailable,” said Epley. Zero trust is critical to supporting these disparate workloads.

When edge workloads span myriad locations, technical requirements and physical footprints, government needs a consistent workload lifecycle approach — one that includes on-premises, hybrid/multi-cloud and edge architectures, as well as continuous operations and management of sites that might have intermittent or very low-bandwidth connectivity.





Conclusion

Zero-trust principles start with making security a foundational component of all projects, whether developing new products or implementing new infrastructure. Instead of building security around network access, ZTA is applied to every interaction as a practice across the organization.

By using ZTA, along with a zero-trust network and zero-trust edge as the core security principles, the government can address the challenges in support of a range of mission-critical use cases.

For more information on ways Red Hat can help your agency innovate while meeting regulatory requirements and reducing risk, visit https://www.redhat.com/en/solutions/public-sector





About Red Hat

Red Hat helps customers integrate new and existing IT applications, develop cloudnative applications, standardize on our industry-leading operating system, and automate, secure, and manage complex environments. As a strategic partner to cloud providers, systems integrators, application vendors, customers, and open source communities, Red Hat can help organizations prepare for the digital future.

To learn more, please visit www.redhat.com/government

About GovLoop

GovLoop's mission is to inspire public sector professionals by serving as the knowledge network for government. GovLoop connects more than 300,000 members, fostering cross-government collaboration, solving common problems and advancing government careers. GovLoop is headquartered in Washington, D.C., with a team of dedicated professionals who share a commitment to the public sector.

For more information about this report, please reach out to info@govloop.com.

www.govloop.com | @GovLoop