tufin The Security Policy Company.

Achieving a Zero Trust Network Security Model with Tufin

Executive Summary

Zero Trust (ZT) requires two core capabilities across your entire infrastructure — visibility and automation. The role visibility and automation play in a ZT model is captured by the Forrester Zero Trust framework, depicted in Figure 1 (right). Visibility and automation are the outer "rings" that enable the segmentation, validation, control, speed, and accuracy required to implement a comprehensive Zero Trust model among all the elements of a modern environment.

This white paper presents an overview of the challenges organizations face implementing a Zero Trust threat prevention model, the six components of a Zero Trust model, and capabilities needed to implement these components. It concludes by presenting how the Tufin Orchestration Suite (TOS) can help organizations establish and maintain effective Zero Trust network access.

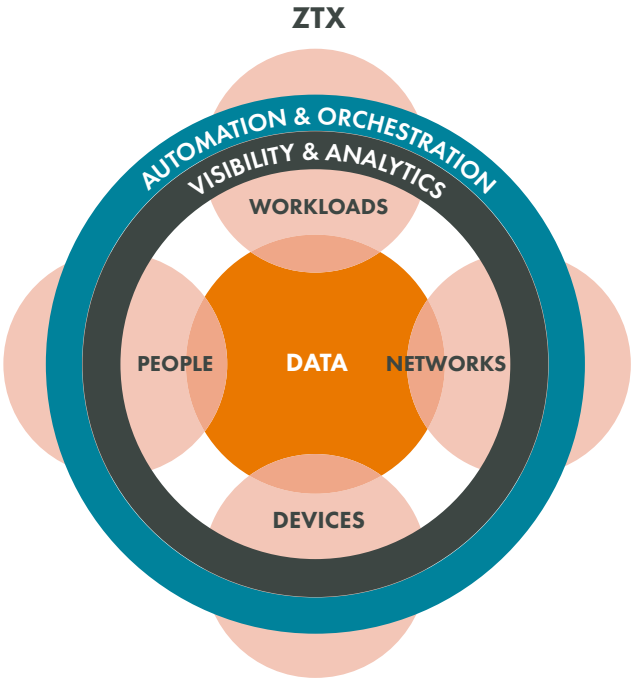

Figure 1: Forrester's Zero Trust model illustrating how a Zero Trust program must include all the elements of an organization's operating environment: workloads, networks, devices, people and data, whereby the Zero Trust program is enabled with visibility and automation across all of these elements.

What is Zero Trust and Why it is Important

"Never trust, always verify" is the core tenet of global analyst firm Forrester Research's Zero Trust philosophy. According to Forrester, it's Zero Trust that dictates how organizations handle their overall network security. Perimeters now have so many ingress points that malicious actors are assumed inside the network. Once inside, 'trust' is an easily exploitable vulnerability if a Zero Trust model is not in place.

Based on Forrester, the traditional approach to network security was within an organization's network perimeter; services and users were trusted, and therefore, could "talk" or connect. It was, for the most part, an open, almost flat network, where traffic or users were trusted by virtue of being inside the perimeter.

The Concept of "Inside vs. Outside" is Obsolete

The problem with this approach, is that attackers are often already inside the network. It could be that they have gained credentials, be it maliciously or lawfully; if they are employees with bad intentions, or for example, or exploit a known vulnerability to access the perimeter. But the moment they are inside the perimeter, without Zero Trust principles in place, they have gained virtually unlimited access across the entire network.

As such, according to Forrester, Zero Trust suggests we discard the idea of the traditional model, that inside the network, it's a trusted zone, and shifting security from a perimeter-based model to a securitymodel that's based on continuous verification of trust. This requires that organizations create micro-perimeter or micro-segments to control access to sensitive assets, limit user privileges to the necessary minimum, and enhance risk detection and response with analytics and automation.

One thing we know to be true – Zero Trust fundamentally enhances an organization's network security posture by removing its absolute reliance on perimeter-based protection.

Zero Trust fundamentally enhances an organization's network security posture by removing its absolute reliance on perimeter-based protection.

Why Zero Trust is Infinitely Challenging to Implement

One of the most effective ways to achieve Zero Trust and protect sensitive assets is through segmentation. Segmentation or micro-segmentation can limit the impact of attacks by making it more difficult to traverse the network. If segmentation is implemented and managed automatically based on policy, it will help organizations to achieve self-securing environments and maintain a least-privileged principle. This is done by controlling inbound/outbound communication flows of zones/services, isolate risky assets, and prevent security incidents from spreading.

Zero Trust and segmentation are the goal for many organizations, but they can present their own challenges. One of the biggest obstacles with any type of segmentation is that it's complex to set, implement and maintain, and it can potentially impact business continuity. Today, the largest problem for security and network admins is how to minimize the attack surface, and write segmentation policy that will be granular enough and easy to maintain, without impeding business continuity, and avoid, for example, shutting down business-critical applications.

Why Segmentation/Micro-segmentation is so Complex to Implement and Maintain

With rapidly-changing environments becoming more and more heterogeneous and dynamic, organizations are now turning their attention to solutions that will help them achieve a Zero Trust network.

The traditional network is fast expanding to cloud, and includes more endpoints to monitor due to IoT, apps, and different types of workloads (containers, VMs, servers, etc.) which are deployed across multiples clouds and on-premise. In addition, the LAN is becoming significantly larger with the deployment of SD-WAN, coupled with the surge in the number of users needing network access from everywhere. This translates into complex networks that provide an enormous attack surface.

i/o Segmentation

When it comes to the hybrid environment, there are two dimensions to segmentation – one is inside the cloud (i.e. workload-level segmentation/micro-segmentation) where IP-based perimeters are no longer relevant as IP addresses are ephemeral, and the other, is segmentation across the hybrid environment.

Managing these two dimensions of segmentation becomes especially tricky and complex, requiring a lot of overhead and expertise. What's more, segmentation for these two environments (on-premise datacenters and clouds) is controlled and managed by different management solutions and teams. In the cloud, more often than not, developers and CloudOps are the ones to create and manage workload-level access rules using security groups, where frequently, they assign overly permissive access to workloads. We have seen many instances of the default cloud setting of any-any-any. Access permission misconfigurations in the cloud can scale in seconds, resulting in a magnitude of problems in production.

The problem becomes even more complex in the cloud as the pace of the changes is growing, and as more organizations are adopting DevOps methodologies, security cannot play catchup with development.

Multi-vendor Platform

With on-premise, however, segmentation is managed by network security teams, but is often managed manually, as there are at least two network security management solutions from different vendors deployed throughout the network (e.g. Palo Alto Networks Panorama, FortiManager, and Check Point SmartDashboard). The challenge here, is that provisioning segmentation policies with two different tools, no matter how hard you try, requires too much time and effort to manually configure and manage.

Further, there are no two policies alike when they are configured manually on different tools. It's impossible to manually apply consistent policy control across all network devices. This leads to inconsistent security policy enforcement and policy gaps, re-work due to errors and excessive time spent on routine tasks (e.g. decommissioning rules, servers or applications, etc.).

Manual Processes Cannot Keep Up with the Pace of Change

Many organizations manage security policies and network access changes manually using ticketing systems for approvals, Excel files (as in... an 18,000 row Excel file!) for tracking, and manual provisioning or custom scripts. In this case, routine firewall changes, such as adding new servers, updating firewall rules or decommissioning objects, can be an arduous task. That's because when implemented manually, crawling the firewalls one by one, or (hopefully) using a searchable index of rules, can take days if not weeks. Keeping up with all change requests, clean-up, recertification, and audit is impossible to manage without automation.

The problem becomes even more complex in the cloud as the pace of the changes is growing, and as more organizations are adopting DevOps methodologies; security cannot play catchup with development.

In addition, there's always the concern that uncontrolled changes may lead to security risks and compliance violations – changes that enable more access than required or approved, or admins that are changing rules on an incorrect firewall can all expand the attack surface.

Solution: Decouple Policy from Infrastructure

To effectively manage segmentation across the hybrid environment, organizations should decouple segmentation policies from the underlying network. This way, it can be applied to any app or workload across the datacenter into the cloud. Centrally managed rules (based on IP, security groups, namespaces, IAM profiles, Network Access Control Lists (NACLs), etc.) can utilize policy automation and provide full visibility across the environment.

Deploying policy automation can also help you implement changes quickly and automatically based on accurate topology path calculations, and policy analysis to ensure fast and precise provisioning of new or changed access.

As a security best practice, you want to ensure that every change is vetted by the relevant task owners and meets security and compliance mandates. It's recommended to create pre-defined workflows for common change handling to ensure full user accountability and control with comprehensive audit trails of all changes. These workflows should be integrated into an organization's ITSM, Jenkins, Slack, or other collaboration processes.

Only by automating changes through an auditable, secure and accurate process, will you be able to enable new access, fix what's no longer compliant, and also ensure that all policy access changes are designed based on the most optimized path.

Automated provisioning will ensure an accurate and error-free process. The last mile of the Zero Touch change process is to provision the change automatically to all firewalls, routers, and Kubernetes for example.

Finally, using automation to create and manage the segmentation policy, you ensure each app/workload is provided the right level of access for its daily operation, and that every policy change is vetted, designed and implemented accurately to achieve Zero Trust security. No human intervention, and no changes are allowed outside of the federated process. This creates self-secure networks and simplifies operation and security.

How Tufin Enables the Zero Trust Security Model

Tufin Orchestration Suite (TOS) offers a segmentation policy management solution that helps network, security, and cloud teams to embed and maintain Zero Trust practices across their hybrid network. Security automation drives optimized, least-privilege segmentation, and trusted and secure changes, while ensuring policy adherence and business continuity.

Tufin specifically maps to the Zero Trust principles in 7 key areas:

1. Complete Visibility & Accurate Topology Modeling

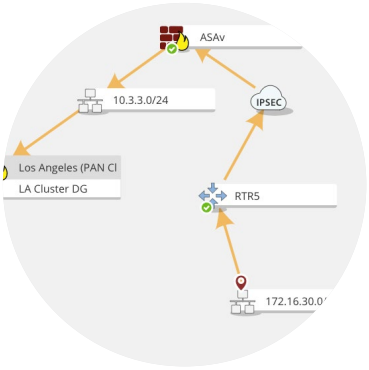

Tufin provides full visibility and control of which apps, workloads, and network security devices (e.g. SDNs, firewalls, NGFW, routers, switches, or security groups) are currently deployed and how they're connected, as well as what can talk to what, and who can talk to who (e.g. traffic flows) across the multi-cloud, hybrid environment.

Tufin topology map is created by connecting to all of the hundreds of firewalls, thousands of routers, switches, and cloud services, and retrieving all routing tables, as well as taking into account all Network Address Translation (NAT) and port number translation, VPNs, Multiprotocol Label Switching (MPLS), IPV6, security groups, IPAM data, and more.

This results in a precise and highly-accurate model of your hybrid environment, so you can immediately start to monitor actual traffic across environments, and identify anomalies, such as misconfigurations, and potential threats.

2. Security Segmentation

With Tufin, you can consistently and automatically apply any level of segmentation to microservices, network zones, user IDs, or App-IDs (even if, for example, the database tier resides in your datacenter, and the web tier is deployed in AWS public cloud). Using the Tufin solution as an orchestrator between zones, tags, and namespaces, and from security group rules to firewall rules, ensures segmentation is enforced across cloud environments and datacenters, following the workload anywhere it's deployed.

By mapping apps, workloads, business units and subnet connections (east-west and north-south traffic), you can start modeling security policies and segmentation options. For example, you can locate and prioritize low- and high-value assets, view which assets are most connected, and apply an appropriate segmentation strategy and granular security policy. By deploying a network topology map, you essentially create a shared understanding of security concerns and requirements between the various stakeholders – app owners, developers, and network and security personnel.

Tufin provides a path to segmentation. In the cloud, it starts by monitoring traffic and automatically learning the app/workload communication flows, and creates a whitelist policy which you can then edit, using natural high-level language to define segmentation policies. The result is a policy baseline, including rule properties and flow restrictions. Further, the policy can be generated as a YAML file, so it can be easily embedded in the SDLC for shift-left security.

For IP-based segmentation, you can use the Tufin Unified Security Policy (USP) matrix to create security policies to control what traffic is allowed between the zones. To quick start IP-based segmentation, you can use one of Tufin's predefined compliance segmentation policies whereby all rules can be compared to these policies. You can define exceptions to policies, if needed, and once defined, policies are automatically distributed and enforced.

3. Secure Kubernetes Deployments and Automate Segmentation Policies Across the App Lifecycle

With Tufin, micro-segmentation can be set and applied automatically across hybrid and multi-cloud environments. By assigning tags (e.g. PCI-sensitive) to workloads, identical policies can be applied to all workloads with the same tags. Once a workload with the same tag is instantiated, the corresponding policy can be automatically applied to it. This policy can then be enforced to all network security devices as needed, to enable secure communication between cloud environments and the datacenter.

In addition, Tufin enables the removal of reliance on developers to configure network parameters manually. Tufin learns the app/workload's intended behavior, and suggests the right policy (security groups, rules, etc.) which can then be configured into Kubernetes. This behavioral profile will then be automatically applied to heterogeneous workloads, irrespective of their location in the network, cloud or on-premise.

To do this effectively for cloud-native apps, Tufin security checks can be added as a step in the CI/CD pipeline, to ensure secure-by-default setting. This will help CloudOps teams detect misconfigured rules early in the process, and apply remediation before deployment, without halting development efforts.

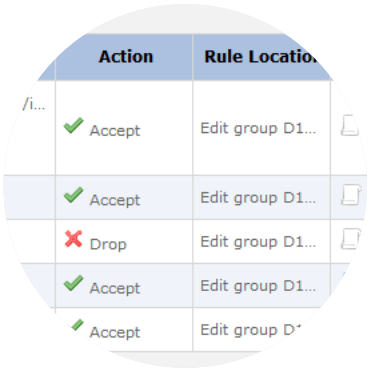

4. Policy Optimizations vis Change Automation

Tufin provides the ability to visualize network connectivity, assess whether that connectivity is risky, and understand where changes need to be applied. With Tufin, you can implement changes quickly and automatically based on accurate topology path calculations, and policy analysis to ensure fast and precise provisioning of new or changed access. To achieve this, Tufin not only locates the exact applicable rules in all relevant network security devices and infrastructure components, such as firewalls, SDNs, routers, etc., across the hybrid environment, but also pinpoints which changes in these rules need to be made to enable the required change.

Tufin ensures that every change is vetted by the relevant task owners and meets security and compliance mandates. Tufin also offers pre-defined workflows for common change handling to ensure full user accountability and control with comprehensive audit trails of all changes. These workflows can be integrated into an organization's ITSM, Jenkins, Slack, or other collaboration process.

In addition, Tufin ensures that all policy access changes are designed based on the optimized path, and via a predefined change window, automatically implements the changes to all relevant firewalls, routers, switches, and Kubernetes, to ensure an accurate, error-free process.

5. Monitor and Control Network Changes

Tufin identifies unauthorized changes made directly to network devices outside of the authorized change management and configuration orchestration processes as controlled by Tufin, and alerts via email when policy violations occur. Tufin can also identify when unauthorized personnel attempt unpermitted actions or actions that take place outside of the organizationally defined change processes. The authorized and unauthorized change tracking and reporting from the Tufin dashboard identifies where network security devices were modified outside of the approved workflow and operations windows. Similarly, Tufin can be tuned to monitor unusual network activities according to policy based on hit count, zone changes, or anomalous port-specific network behavior.

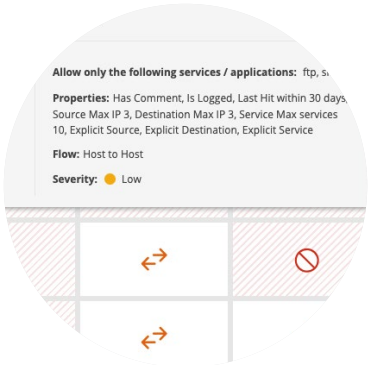

6. Assess, Prioritize and Mitigate Risk

The challenge with risk has always been that too many critical vulnerabilities are discovered and not enough resources are available to patch them. Moreover, the reality is that vulnerabilities with high CVSS scores aren't necessarily the ones exploitable in your network. For example, a medium-level vulnerability associated with a business asset having multiple access points may be used by an attacker to gain a foothold in your network and move laterally to pivot to other high-value, sensitive assets. As a result, this vulnerability is more prone to exploitation by attackers, and consequently, should be considered a high-level priority for remediation or mitigation.

Organizations need a method of prioritizing the vulnerabilities that should be patched first and find a way to mitigate the risk of exploitation until they can be remediated.

Tufin integrates with leading vulnerability management solutions, including Tenable.io, Tenable.sc, Qualys VMDR, Rapid7 Nexpose, and Rapid7 InsightVM to provide risk-based network insights that help organizations efficiently prioritize remediation and mitigation efforts by correlating vulnerability data with network insights. By combining vulnerability scanner output with network access data, organizations can understand how these vulnerabilities are contextually exploitable today, enabling security admins to identify and address vulnerabilities that pose the greatest threat to critical business assets.

7. Maintain Continuous Compliance with Automation

If done right, automation provides two-fold benefit when it comes to achieving continuous compliance and Zero Trust.

With policy automation you can generate a global policy that is optimized for apps/workloads based on a least privilege principle, is compliant with regulations and other security mandates, and is automatically applied across your hybrid network.

To help you jump start compliance across your Zero Trust network, you can use Tufin pre-defined PCI-DSS, HIPAA, NIST, CIS policies where we translate access-related requirements into segmentation policy. Any violation is automatically alerted and blocked, and because you define these policies using Tufin, you can apply them, irrespective of the underlying security network infrastructure, SDN, or cloud. If you plan to switch vendors, or migrate to a different environment, you can do this without the need to redefine your policies.

The second benefit is in maintaining compliance, where, with automation you can easily maintain your policy globally. By default, any automated change will be vetted, tested against the policy, designed and implemented in minutes, with no human errors or misconfigurations.

When it comes to compliance and Zero Trust, automation is key in helping you integrate security best practices at scale, while dramatically shortening the time it takes to continuously manage them—all while staying aligned with security and compliance policies.

By adopting a Zero Trust approach, you are essentiality adding a heightened level of predictability to your network access practices. By whitelisting what can talk to what and who can talk to who, you will be able to define your intentions in parallel to responding to anomalies. Simply put, this ability is far greater than attempting to guess where the next attacks may come from, or which vulnerabilities will be used by attackers to penetrate your network, and move laterally across your environment.

About Tufin

Tufin (NYSE: TUFN) simplifies management of some of the largest, most complex networks in the world, consisting of thousands of firewall and network devices and emerging hybrid cloud infrastructures. Enterprises select the company's Tufin Orchestration Suite™ to increase agility in the face of ever-changing business demands while maintaining a robust security posture. The Suite reduces the attack surface and meets the need for greater visibility into secure and reliable application connectivity.
With over 2,000 customers since its inception, Tufin's network security automation enables enterprises to implement changes in minutes instead of days, while improving their security posture and business agility.

tufin The Security Policy Company.

DP20210415