Identity-Based Improvement to Zero Trust with BlackRidge Technology Case Study

Edward Amoroso, TAG Cyber
eamoroso@tag-cyber.com

August 30 ,2019

Abstract

An improvement to zero trust security is introduced by John Hayes of BlackRidge Technology[1] based on the notion of separating identity and confidence. A measurement scale is explained with boundary conditions of maximal and zero trust. Practical use of the scale is demonstrated in a case study analysis of the BlackRidge TAC platform.

Introduction

An improvement to a well-known technical concept emerged recently from John Hayes, CTO of BlackRidge. That is, he has introduced a refreshingly helpful view of *zero trust security*. The enhanced view that results should help to make zero trust security more useful as a design tool for analyzing systems, protocols, and other areas of interest for security engineers, developers, and researchers.

This article explains the conventional model of zero trust and then outlines the enhanced view introduced by Hayes. The BlackRidge Transport Access Control (TAC) commercial platform is used as a case study to demonstrate the utility of the model for security analysis. The work reported here will hopefully spur other ideas on how this enhancement to zero trust might be useful in the cyber security community.

Conventional Model

As most cyber security experts attest, the conventional view of zero trust security that has evolved in our community is described in the context of an enterprise firewall. The familiar definition is that zero trust removes the edge perimeter from the access validation equation. Applications, services, or workloads must therefore authenticate devices and users, regardless of the network context. LAN presence is no longer sufficient to demonstrate trust.

This model has been popular for a couple of reasons: First, it is consistent with the obvious dissolution of the enterprise perimeter that is occurring in organizations, large and small. So, the definition connects with the observations of IT professionals. But second, the model helps

[1] The work described in the paper was communicated to the author during private conversations with John Hayes.

drive computing to the cloud, which increases its support by vendors of any cloud-based service. Commercial motivation is thus a strong driver for zero trust.

The challenge from a research, development, and analytic perspective is that zero trust is a deeper concept, one that should be explained in a non-architectural context. Least privilege, for example, is a similar security concept that is poorly defined in terms of specific applications, but much better defined in terms of the intent of minimizing access when not needed. As such, zero trust would seem to deserve a more elegant definition.

Identity Model

As suggested above, an enhancement to this definition of zero trust was offered recently by John Hayes, who serves as the CTO for cyber security vendor BlackRidge. His company provides a creative TCP-based gateway and endpoint solution called TAC (Transport Access Control) that uses injected session tokens to validate the identity of incoming packets using their First Packet Authentication™ (FPA) approach. Hayes' enhancement to zero trust security is as follows:

Zero trust is improved when identity is separated from confidence.

This observation stems from the notion that identity information is always associated with confidence. For example, when a username is entered to an application, then low (but perhaps non-zero) confidence exists that the true owner of that name is really the one doing the entry. Upon successfully meeting an authentication challenge, perhaps biometric-based, that confidence is quickly driven upward along some scale.

Hayes envisioned a confidence scale that can be moved up and down in a way that is measured separately from how an identity is established. In this way, one can examine the boundary conditions of that scale. At one end of the confidence scale, two entities have high, or maximum, trust and can freely interact without concern for any threat. At the other end of the scale, however, two entities have low, or *zero*, trust.

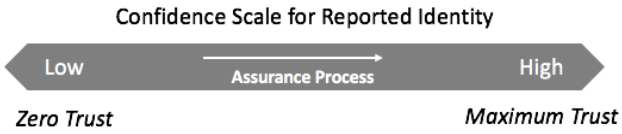

Figure 1. Defining Zero Trust in the Context of a Confidence Scale

This improved view of zero trust as an edge case in a confidence scale is attractive because it correctly models trust as an attribute with value that can be increased or decreased. Security architects can then determine where on the confidence scale a given system resides in order to determine the correct security controls that must be imposed. The scale thus helps provide a more threat-driven allocation of security investment to reduce risk.

BlackRidge Use-Case

The BlackRidge team has apparently been using this confidence scale to drive security policy decision-making in the TAC solution referenced above. Briefly, the TAC solution involves a variety of identity-related functions including insertion of an identity token by a user or device directly into a TCP session. Once received by a BlackRidge gateway, which is managed by an enterprise, the session identity is resolved, and a policy decision can be made.

In contrast, most enterprise teams today allow inbound TCP sessions to proceed across their perimeter toward some premise-hosted applications. And this is no different for cloud-hosted workloads: Packet access is still granted through a defined cloud hosting perimeter. Once the session reaches the desired premise-hosted workload, identity and access management (IAM) functions decide on whether the workload access should be granted.

Viewed in a stepwise manner, this familiar enterprise access scenario helps to illustrate the zero trust concept. First, a user is assigned a source IP address (SIP) with little assurance other than an operating system has been used to populate a packet header with a well-formed address. This is a zero trust situation, because spoofing SIP is a simple process. The first step in the process is thus SIP assignment to the first packet from a user:

Figure 2. First Step: SIP Assignment to the First Packet from a User

Once the first packet reaches the perimeter gateway for the premise-hosted app, it will presumably be directed to a router with the ability to perform packet filtering operations. This gateway will inspect the SIP and any other header-resident information, but it cannot do much more – simply because the session has yet to be established. Because SIP is so easily forged, this step doesn't change the zero trust nature of the session.

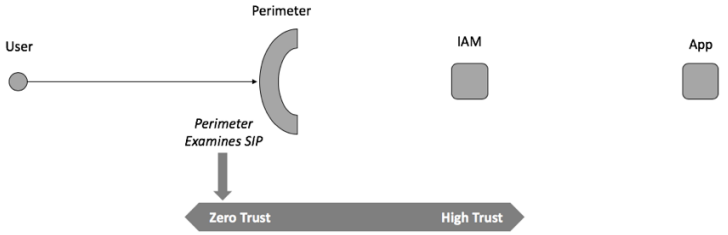

Figure 3. Second Step: Perimeter Examines SIP of First Packet

The highest risk aspect of this process comes next in the process: That is, because the SIP provides essentially zero additional confidence in the veracity of the reported source, the perimeter cannot make an informed access decision. It is forced to accept the packet and

forward it to the IAM function. This step creates an unsafe access condition, because it allows spoofed traffic to traverse the perimeter with impunity.

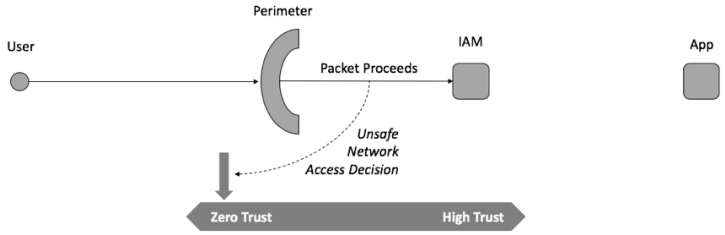

Figure 4. Third Step: Perimeter Gateway Allows Packet to Proceed

Once the packet finds its way to the IAM, a TCP session can then be established across the perimeter. Within this session, the user and IAM would follow whatever application-level identification and authentication protocol has been selected. Assuming that this protocol has been properly implemented with multi-factor proof, we can assume that trust in the veracity of the source has been increased from zero to high.

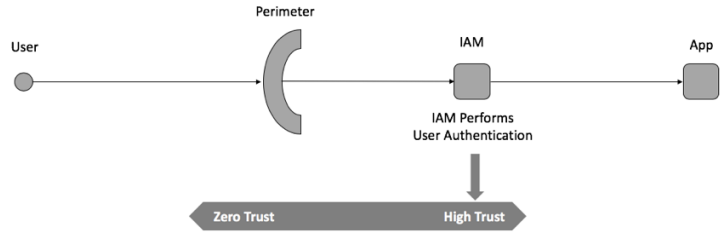

Figure 5. Fourth Step: The IAM Protocol Establishes High Trust

Now that trust has been increased from zero to high, allowance of the user to proceed with a new first packet to the app is no longer unsafe. Instead, the high trust allows the app to be confident in the veracity of the source. This might not preclude use of an additional location or app authentication steps, but the session is no longer highly suspect. Regardless of this last step, the overall process is unsafe between the perimeter and IAM-related functions.

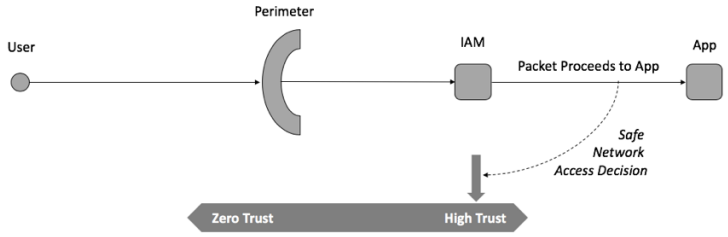

Figure 6. Fifth Step: New First Packet Proceeds to the App

The BlackRidge solution is designed to remove the unsafe network access shown in the stepwise illustration above. It is specifically designed to acknowledge the zero trust nature of

SIP authentication by using TCP session-injection of identity-based tokens for use at a perimeter or other access device, including cloud-hosted infrastructure. The idea is that users or devices begin the process by injecting a cryptographically-strong identity token into the first packet.

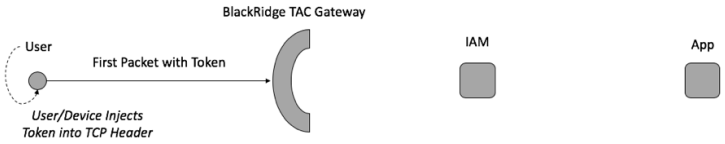

Figure 7. BlackRidge First Packet Token Injection

By installing a BlackRidge authentication gateway or agent adjacent to the perimeter or other access gateway, the accepted packet can be used to establish higher levels of trust for authorized users. It essentially assures that before any session is initiated to an IAM, app, or other hosted workload, that the trust level associated with the first packet has been appropriately vetted. In this way, zero trust security is achieved in an elegant manner for enterprise security teams.

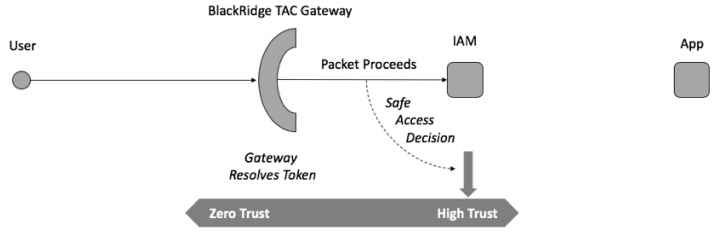

Figure 8: BlackRidge Gateway Token Processing

The BlackRidge case study illustrates how easily the confidence-scale definition of zero trust security can be used to highlight and explain the risk of access decisions on a network. Interested readers might examine how the model might be used to show interim values on the scale such as moderate trust (somewhere in the middle of the scale). The model easily supports such security risk analysis.

Concluding Remark

An enhanced means for using zero trust security in enterprise security design and architecture work has been long overdue. Once the perimeter-based enterprise is officially dead, the use of zero trust security as a useful design concept will remain. So, it only stands to reason that our community begin to employ views of zero trust that are design independent. BlackRidge Technology's enhancement to zero trust security as separating identity from confidence is a great start toward deeper understanding and better application of an important security principle.