The Three Practical Use Cases for Network Layer Identification; Interview with BlackRidge Technology CTO John Hayes Part I

Followers of my previous blog entries should recognize the next company in DCIG's Executive Interview series. I have previously discussed both the technical and operational impact of BlackRidge Technology's patented breakthrough technology known as Transport Access Control (TAC). Today, BlackRidge announces their first product, Eclipse, based on their TAC technology. I begin a discussion of this release, in the form of a multi-part interview series, with BlackRidge Technology's CTO John Hayes.

In this first entry, Hayes and I discuss the need for network layer identification and some of the more popular use cases of the technology.

Ben: *John, thank you for agreeing to talk with me today. I've written about TAC, now known as Eclipse, a number of times but I think it would be good for us to review what your company does. What is network identity and why should I care?*

John: Very simply, BlackRidge applies identity to networks. Identity itself can be something as simple as a user name and password. It can be more secure smart card credentials. It can be how you access a single sign on system.

But whatever form of identity you are using today, identity is interpreted and understood only at the application layer. By that I mean identity is used only by applications, whether I

am logging into a service on the web, or an email server or what have you, the application layer is aware of identity.

But all of the network below it, for instance TCP/IP and all the protocols that underpin the internet, are completely unaware of identity. This results in networks being managed based on topologies, and it's very hard for organizations to do strong policy enforcement at the network layer.

For instance, you start off a company and the company grows. Like any organization it grows in fairly random ways. Ideally my company would grow along lines that would be easy to implement with a network topology, VLAN, subnets and things of that sort. In reality, there's no resemblance of the organization to the network, and therefore I need to compromise my policies in order to make them work well within the network.

If you can apply identity to your networks, it gives you the ability to say, *"OK, I can implement the policy I want, and I no longer have to compromise."*

What BlackRidge does is we can take identity, compress it down to a small enough element that it can be communicated and interpreted at the network layer to allow you to do this.

Ben: *In real world terms, what does having this additional layer of identification get me? What problems can I solve?*

John: There are three use cases that we see most often. Their applicability obviously depends on the customer needs.

The first one is what we call the absolute security use case. That is where we use identity to effectively remove resources from the network. Those resources only appear when they are being accessed by a recognized requestor, which has both the identity and the authority to access those resources.

To all unidentified and unauthorized requestors, there are

essentially no resources on the network. And when I say the requestors cannot see them, I mean they cannot scan, they cannot look around, they cannot access, they cannot coerce the resources to give up their user name and password. They cannot see the machine or resources there. This is of primary importance to people where security is the overriding concern.

The second use case is what we call the ROI use case. If I am an enterprise type organization, I have a public facing internet presence, and I am trying to provide a service to known users, whether they are employees or customers.

Today, we are talking to customers and what we are finding is those public facing interfaces range anywhere from 50 to 90 percent of the traffic that they are seeing is unwanted. Now unwanted traffic can range from network scans to network reconnaissance. However it can also be attacks such as denial of service attacks or it can be spam; in short, it can be all sorts of these things.

But what is significant is that in order to get the information I want from the people I want to communicate with, I need to over provision my front end pipes and security resources, anywhere from a factor of two to a factor of 10, just to get the communications I want out. That means my firewalls, my deep packet inspections, my intrusion protection system, all need to be scaled up correspondingly.

For instance, if I have a gigabit connection to the Internet, that's a pretty big pipe. But if I've got a gigabit connection, and I see 90 percent garbage, that means I have have to put a gigabit's worth of firewall, intrusion detection, deep packet inspection, all of this other stuff out there, in order to get 100 megabits of good data.

There's two different ways we would approach this problem. If I believe that all the traffic coming at me should be known, in other words it is going to be only my own employees and my

customers, and we can provision and assign identities, we will just simply filter out and drop all unauthenticated traffic. Now you take your resources and you focus them on suspected good traffic instead of all of the traffic coming at you.

The second one is I've got a mix of anonymous traffic coming at me. I still need to service it but I want to have different levels of service. What I mean by that is, my employees and my customers are known and I have an established relationship with them so they are going to get a certain quality of service.

This anonymous traffic, I can now separate and give them best effort service, as opposed to a predefined level of service that otherwise I would have to do. So in both cases customers can realize an ROI in the provisioning of front-end security services.

Now the last use case actually has to do with how I manage my networks. As we discussed earlier, my networks grow in certain ways as I add resources, and resources tend to get added organically. I also add clients. And then I also merge with organizations or all sorts of other things happen that cause me to have networks that are not ideal to the way I'm doing things. So I might have a finance guy sitting next to an engineer sitting next to somebody in sales.

Now I sure do not want my sales guy and my engineer to be accessing my finance server. But if they're all on the same VLAN or subnet because of the way the building is wired, I am stuck with that situation. Identity gives me the ability to segment those resources based on the policies I want, based on identity, instead of being tied down to traditional network policies.

This occurred very recently. We were talking to a customer that had just acquired a company. But part of the facilities they acquired with this new company were actually still with

another company.
They needed to carve out this piece of the facility and isolate it without moving everybody around and actually doing some fairly unnatural organizational acts. In other words, they could not do this with current technologies.

With BlackRidge, they are able to very clearly say, "*OK, we are going to isolate these sets of users, and we are going to have them only access the resources they are supposed to see which is going to be outside the company. Everything else inside is now protected.*"

It's just a much easier way to do it and it saves a lot of what are relatively non-technical things we're working around. You can actually apply the policy you want without having to compromise based on other, in many cases non-tech
nical, requirements. So those are really the three major use cases that we see.

In Part II of this executive interview series I will discuss the deployment and technical aspects of Eclipse and how BlackRidge continues to find new use cases for the product.

In Part III of this interview series I will discuss the management of Eclipse and how BlackRidge continues to find new use cases for this product.