Blitzkrieg and VDI Edge Protection.

Don MacVittie, 2012-26-01

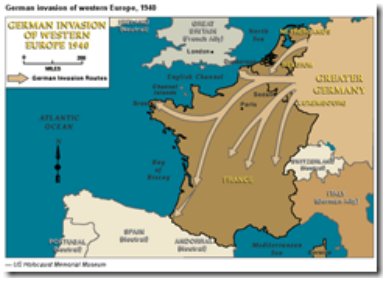

By now, everyone even vaguely familiar with information security knows the military maxim of blitzkrieg – burst through the hardened defense at a single point and then rush pell-mell to the rear where the soft underbelly of any static army lies. It is a good military strategy, provided you have the resources to break through the defenses and follow up with a rapid advance into the rear areas. While there are variants of this plan, and a lot of discussion about how/when it is strategically worth the risk, historically speaking it has been a smashing success. Germany did it to France and the Low Countries in 1940, to Russia in 1941, Russia returned the favor in 1943, and the western allies joined used it successfully at Normandy in late 1944. Sherman's March to the Sea in the American Civil War was just such a ploy (though Sherman was more willing to hit civilian targets than a 20th century general would have been, it was still a rush to the soft rear), and the first Gulf War had the coalition forces doing much the same. These are just the large-scale instances of this theory in operation, but you have to admit it works. The risk is high though, as the Germans found out at Prokhorovka, and that alone makes generals cautious that they have the resources and intelligence reports to burst through in the first place.

The difference between the military maxim and the theory that information security should follow it is an important one. In military theory, you only harden behind the lines if there is a high likelihood that the enemy forces will find a weak spot in your lines and exploit it to get at the rear areas. The conundrum for the defensive leader finding themselves in such a situation is that every combat soldier placed to the rear is one less combat soldier on the front, increasing the likelihood that there will be a breakthrough. In information security, the problem is that the resources of the attacker are theoretically unlimited. Unless they are apprehended by the authorities in their home country, there is no penalty for attacking over and over and over. The limiting factor for the attacker – that they might smash themselves upon their opponent – does not exist at this time in Internet parlance. An attack fails, that merely means the attacker marshals the same exact set of resources and tries again.

The defense, on the other hand, still has a limited number of resources (dollars and staff hours) to defend themselves with. And they must make the most of them. Defense in depth is an absolute necessity, simply because the attacker can continue ad-infinitum to try attacking, and the number of attackers is unknown but large. That leaves a heavy burden on information security staff, who have settled into the glum belief that it is "not if, but when" they will be defeated. While the ultimate solution to this problem rests outside the purview of corporate security, in the interim, it is necessary to do what can be done to simplify and strengthen the fortifications that are between ne'er do wells and corporate resources.

Just to add fuel to the fire, this is all happening at the same time that organizations are facing increasing pressure to expose more and more of their internal architecture to the Internet so that users can access their applications from essentially anywhere. So to put it into military terms, there are numerous hostile entities, an ever increasing front length, and a static number of defenders and resources. That is ***not*** a recipe for success in most scenarios.

So what is the serious information security professional to do? Well the first steps have already been taken. Defense in depth is just a fact that most organizations live with, down to firewalls between departments for some organizations. Anti-virus tools and encryption are the norm, not the exception, and external access is generally protected by a VPN. But new technologies bring new challenges, or more frequently make old but low likelihood challenges into higher priority issues.

As we deploy VDI – and we are deploying VDI at a faster rate than I'd expected – the issue of edge security becomes more and more of an issue. If you expose VDI desktops to the world so that your workers can log in at any hour and get some work done, or an employee who's sick can stay home to avoid infecting others but is well enough to work can do so, you will have to find a way to lock that interface to the world down so that users can get in, but hackers cannot. This is more important than most interfaces because the interface sits in front of user desktops, and they generally have more access than a server.

While there are a variety of ways to attack such an inlet, DDoS – to keep employees from working remotely – and Trojans are the two most likely to be successful. What you'll want on this inlet is a way to check that the client – be it PC or iPad or whatever – complies with security policy that includes at least rudimentary virus checking (since the client device is outside your network and possibly not even a corporate resource), and a way to resist DDoS attacks. A network level tool that shunts detected DDoS attacks off to neverland, like F5's own BIG-IP is going to be the best solution, since traditional firewalls are aimed at detecting more traditional attacks and can become victims of a DDoS. Regardless of what you choose to protect against this type of attack, it should be something you can guarantee will stay standing when hit with thousands of dropped connections a second.

And you'll want to be able to apply more generally corporate security policies. That's a tough call in a VDI environment. While a product like BIG-IP can be set up to use your corporate security policies for access and authentication purposes, it is difficult – both legally and technologically - to force corporate security policy on employee-owned devices. Legally you can limit access based upon the status of the machine requesting it, the user name, and the geographic location, but you can't insure that the device meets with the same stringent policies you would require on your internal network. And that's a problem, because VDI *is* your internal network. Time will tell how large this threat looms, but I wouldn't ignore it, since we know it's a threat. Legally you can ask employees to agree to be bound by corporate security policy when accessing the corporate network from a home machine, but I honestly don't know of anyone doing that today – and I am not a lawyer, so maybe there's a good legal reason I haven't heard of anyone doing just that.

In the end, the benefits of allowing some or all users to access their desktop remotely is a huge benefit, but be careful out there, the number of attackers isn't going down, and while we're working all of this out is their opportunity to take advantage of weaknesses. So protect yourself. I'd recommend F5 products, but there are other ways to try and resist the hoards should they come knocking at your public VDI interface. Whatever you choose, just make certain it is implemented well.

F5 Networks, Inc. | 401 Elliot Avenue West, Seattle, WA 98119 | 888-882-4447 | f5.com

F5 Networks, Inc.
Corporate Headquarters
info@f5.com

F5 Networks
Asia-Pacific
apacinfo@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.com