Article

By Leif Nixon,
Security Expert at Sectra

The mythical air gap

That morning when I arrived at the office, I was met with shouts of, "There he is!" I'd like to be able to say that they were happy shouts, but unfortunately this was not the case. I was quickly informed that during the night, the big uninterruptible power supplies (UPSs) that powered the data center had turned out to be, in fact, interruptible. And that was my fault.

To provide some background to this, I was working at the time as an IT security officer at a large data center, and one of the many things it was my job to worry about was whether any unexpected devices or services appeared on our network. Security people don't like unexpected things. And so, I had developed a monitoring service that scanned our whole external network range every night and reported any changes in our network posture. After having run it successfully for some time, the previous day I had arranged for it to also scan our internal networks. This was a mistake.

We had a separate, supposedly isolated network for all the scary stuff: the supervisory control and data acquisition (SCADA) devices and management interfaces that we didn't dare expose to anybody but our most trusted staff. However, the computer that ran the scan actually had a connection to the SCADA network, and I had failed to exclude it from the scan. As it happens, the UPS controllers were on that network, and they proved to be so fragile that a simple port scan caused them to keel over and die.

In fact, there are many devices that ought to be confined to separate networks. Devices that might be too frail, or might rely on inherently insecure communication protocols like Modbus or IPMI. Or devices that might not have any particular vulnerabilities at all, but contain sensitive data that warrants an extra layer of separation. For the ultimate separation, conventional wisdom says you put your network behind an "air gap."

An air gap is just what its name suggests. There should be a gap of air around the network. No physical connections. Nothing goes in, nothing comes out. If the attacker can't reach your network, they can't hack it, right?

The human factor

Unfortunately, there are very few computers that don't need to communicate with the outside world somehow. Even if you are not aware of it, there is probably a network cable across the gap somewhere. Even if not, information is probably being passed back and forth via USB sticks, or via people plugging their laptops into the isolated network.

In practice, many air gaps can simply be viewed as a particularly slow network connection. Communication over it may be slow and somewhat random, but information still flows through it.

There are many, many real-world examples of this. The uranium enrichment systems in Natanz, Iran, were of course air gapped, but were still infected by the Stuxnet malware, introduced via USB sticks. I have seen critical control systems infected by malware inadvertently brought in via a service technician's laptop. And of course, there was the time I killed those UPSs myself.

Is an air gap the ultimate security solution?

There is a superstition that air gapping networks is the ultimate solution that solves all security challenges within a network. However, how many networks actually have no connection at all to the outside world? It could be as simple as a network cable, a USB stick or a connected laptop and the air gapped network is not isolated anymore. This doesn't necessarily have to be a problem, but you need to be aware of it and keep in mind that information is actually flowing, even in cases where you actually have a real air gap. There is nothing to do about this, the important thing is that you are aware of it and provide processes to manage the situation.

Tips on how to handle the mythical air gap

Air gapping is a long-established approach to separating critical networks, and was introduced and initially primarily used in the military. In today's highly digitalized society, the risk that a network is not isolated—even though it is supposed to be air gapped—is very high. You need to keep this in mind when you work with security within your networks. I have collected a few tips on how you can secure the systems behind the air gap:

Tip 1

There will almost certainly be a need to transfer data across the air gap in one way or another. Rather than inventing unofficial ad hoc solutions, make sure your organization provides approved and official processes and procedures (guidelines) for transferring data securely. This will significantly reduce the human factor of making mistakes that could lead an unauthorized person into your internal networks. Examples of ways to transfer data include:

- Opening up a tunnel in your firewall to a specific device within the protected network.
- Setting up a data diode across the air gap, allowing you to control which direction the data should flow. With a data diode, it is physically impossible to send data in the wrong direction.
- It is difficult to completely stop using USB sticks since this is an efficient way of transferring and copying data. Therefore, your organization needs to implement guidelines for how to use a USB stick in a secure way. For example, you could have a "USB laundry," which means that when you have used a USB stick to copy data between different networks, you put the device in a machine that "washes" it and erases all data, so you can reuse it.

Which data transfer method is appropriate depends on the specific situation and security requirements. As always, you need to do your own risk assessment.

Tip 2

Another important way to secure the network behind the air gap is, as usual, to always keep monitoring your networks. If you do this, you will immediately see if there is unwanted activity in your networks and have the chance to act according to the actual risk at hand instead of acting on an attack on your networks.

Tip 3

Lastly, it is also very important to verify that systems that should be isolated actually are isolated. This can be done by monitoring for unexpected connectivity and continuously testing the security features within the isolated networks. And this has to be done all the time—every day, all year round—to make sure there are no security holes for malicious actors to get into the network.

Key takeaways

This is not to say that air gaps are useless. In fact, network separation is one of the most important tools you have for protecting your most vulnerable and important systems. But it is important to acknowledge that your air gap is likely imperfect. You have to accept that there will always be some form of information flow, and instead of trying to stop it, you need to find ways to manage it.

SECTRA
Knowledge and passion