Electronic Design

By BEN WARNER, Director of Applications Engineering,
DIGISTOR, a CRU Data Systems company, https://digistor.com

Air-Gapped Networks (Part 3): Cross-Domain Solutions and Off-Site Backups

Solutions are in demand to address the need to access isolated pockets of information as well as securely back up data off-site in case disaster strikes.

In Part 1 of this three-part series, we introduced the concept of "air-gap networks"—secure computer networks that are physically isolated ("air-gapped") from their unsecured counterparts.

In Part 2, we considered the problems associated with physically moving information from one air-gapped network to another. As part of this, we noted that the removable, encrypted drive is a fundamental building block to solve information transport issues created by maintaining isolated systems with no common networks.

In this third and final article, we delve a little deeper into the issues associated with trying to implement cross-domain solutions. Also, irrespective of whether information is on an accessible network or air-gapped to keep it isolated, it still needs to be backed up. Thus, we will also delve into the secure procedures required to conduct such backups.

Implementing Cross-Domain Solutions

Cross-domain solutions are sorely needed to address how to access isolated pockets of information, or to solve problems using multiple agencies or teams. The scope of the problem transcends team, agency, and even country boundaries.

As we've previously discussed, a wide variety of agencies and programs rely on self-maintained air-gapped networks, many with multiple or forested domains. The growth of such pockets is exponential.

When a new program or gap is identified, data must be communicated to vendors and proposed solutions be communicated back. One agency may have information that's required to help solve another's problem, if only everyone knew what information everyone else had or needed.

Seeing Both the Forest and the Trees

There's an old expression in America about not being able to see the forest for the trees. This refers to someone who is so involved in the details of something that they fail to notice what's important about the thing as a whole.

Have you ever wondered where this expression came from? In fact, it originated in England as "can't see the wood for the trees." There's a city called Bath in the county of Somerset, England. This beautiful city is known for—and named after—its Roman-built baths.

In the heart of the city is a concourse of houses that were designed by the English architect John Wood the Elder sometime around the 1740s. A tree was planted directly in front of these houses, and it grew so large that people began to exclaim: "You can't see the Wood for the tree!"

This expression comes to mind when one thinks of "forested domains," which may be visualized as a domain over other domains. We might also regard a forested domain as being a parent domain that branches out and watches over child domains.

The U.S. government isn't really comfortable with the forested domain concept—for security reasons it would like to keep everything isolated and air-gapped. However, sometimes the benefits of access to information within a forested domain outweigh the arguments against this type of hierarchical network structure.

Who Knows Who Knows?

One of the greatest strengths of employing air-gapped networks inside secure facilities in the form of Sensitive Compartmented Information Facilities (SCIFs) or Special Access Program Facilities (SAPFs) is that any confidential, secret, or top secret (TS) information is as secure as it can be.

Contrariwise, some of the greatest problems associated with that are almost no one knows what anyone else knows and almost no one knows who needs to know. If you think that's confusing, "you ain't seen anything yet."

Suppose we have a forward base deployed out in the middle of nowhere. Suppose U.S. intelligence operatives somewhere in the world discover something interesting and report this nugget of knowledge back to their superiors. Suppose that, if the people in the forward base were aware of this information it would dramatically change their mission profile. Suppose that very few people in the U.S. government are even aware that this forward base exists? So how are people at the forward base going to hear the news?

While we're talking about this, it's worth noting that if you read news of some advanced new aircraft or weapons system on a news site, or hear about it on a podcast, then it's old news.

As another example, suppose our intelligence operatives discover that one of our adversaries is building something like a hypersonic cruise missile and they report this information to their superiors. Potentially, we have two capability gaps. The first is that we need to have an equivalent weapon in our own arsenal. The second is that we need to have an appropriate counter-capability to be able to take such a cruise missile down.

The problem here is that we could potentially have three different platforms under development, but none of these programs know about the others.

It's been reported that DARPA and the U.S. Air Force (USAF) have a program underway called the Hypersonic Air-breathing Weapon Concept (HAWC) program. The stated goal of this program is "to develop and demonstrate critical technologies to enable an effective and affordable air-launched hypersonic cruise missile."

I personally have never worked on or with any program like the HAWC. Of course, this doesn't mean it doesn't exist. In fact, these public announcements could all be just a show to hide the fact that we have already developed such a missile, but it's still classified as being TS (it's said that people with top secret clearances laugh at people who have only secret clearances). Similarly, it may be that we already have the technology to shoot such missiles out of the sky; we just haven't gotten around to sharing this information yet.

NASIC Knows!

Thankfully, the picture is not quite as dire as I've painted it. The National Air and Space Intelligence Center (NASIC) is the United States Air Force unit for analyzing military intelligence on foreign air and space forces, weapons, and systems. Located at the Wright-Patterson Air Force Base just east of Dayton, Ohio, NASIC has around 4,000 personnel (at least, that's what they tell us—who knows what the true number might be?).

There are several other organizations like NASIC, each with their own areas of interest and spheres of influence. NASIC receives intelligence information from multiple sources. The people at NASIC parse this information, filter it as they see fit, redact information that might compromise the source or something else, and forward it to anyone they think needs to see it.

This information is further filtered as it makes its way down the hierarchy. One of the first groups to see it may involve 20 intelligence officers who break the information into smaller pieces and send those pieces out to different targets. One of these pieces may end up at a fighter command base with five intelligence officers, for example. In turn, they will determine who needs to be made aware of the information and how much of this information they care to share.

Of course, sometimes the people at intelligence agencies distribute "straight up lies" as a way of checking for leaks, identifying bad actors, or spreading a little FUD (Fear, Doubt, and Uncertainty) amongst our foes.

Off-Site Backups

As Benjamin Franklin famously said, ""By failing to prepare, you are preparing to fail." Unfortunately, it behooves anyone who is in charge of securely maintaining data to guard against the loss of said data using the 3-2-1 Data Backup Solution, which can be summarized as follows:

3. Make **three** copies of the data.
2. Use **two** different types of media.
1. Keep **one** copy at a separate location.

Making three copies of the data

Copy the same data files to different disks. Thus, should one disk fail, it won't take the only copy of those along with it. Maintaining at least three copies of data may seem like it increases cost (in the form of purchasing more hard disks or cloud storage) and requires more effort (to keep everything organized), but the payoff in the event of a disaster is priceless.

Using two different types of media

In this context, the term "media" refers to the medium used to store the data. This could be hard-disk drives (HDDs), solid-state drives (SSDs), CDs, DVDs, Blu-ray discs, magnetic tapes, thumb drives, etc. The idea behind using different types of media for extra copies is to protect against different types of technology failures, obsolescence, and/or environmental hazards.

Keep one copy at a separate location

This is what's meant by off-site storage. The idea is to minimize the risk of losing all backups by maintaining one copy of the data at a remote location; that is, far away from the location where the primary backups are stored.

A simple rule of thumb says that the area impacted by a natural disaster can be approximated as a circle with a diameter of 100 miles. Based on this, the location of any off-site backup should be at least 200 to 300 miles from the primary facility. Of course, the environment associated with each facility must also be considered.

In the case of facilities located in California (which could potentially be impacted by a major earthquake) or Florida (which endures more than its fair share of hurricanes), the lowest-risk option is to locate the off-site backup in a completely different state.

In the case of air-gapped networks, only three good options exist for off-site backups. The first option is to create a transfer procedure based on removable media (like a removable hard drive or SDD assembly) and build a schedule around a courier service to securely relocate these backed-up copies of your data to a secondary facility that's equally secured. The second option is to transfer the data to a higher classification network if one is available to you using a cross-domain solution. The third option is to create an on-site environment that's protected from the elements (fire, water, etc.).

Each of these options has its pros and cons. For instance, couriering information is always a complex process when it comes to handling classified data (see Part 2 for more details). In addition to developing a method for transferring data onto the removable source, it's necessary to establish a process and devote human resources to execute that process. Furthermore, it's necessary to define ways to regularly test the backups to ensure they're still working.

The second option is potentially one of the better alternatives, but this option isn't always available or permitted. With the use of a cross-domain, it's possible to transfer all of the data on your air-gapped network to a higher classification WAN. It's also possible to use removable drives or CDs to transfer the data if a network is available, but there's no cross-domain on-site.

Implementing the third option requires the creation of an environmentally safe chamber for a backup system, or the use of environmentally rugged storage devices. Problems you may run into with an environmentally safe chamber are heat dissipation and some form of fire suppressant that doesn't damage the system.

An example of an environmentally rugged storage device is ioSafe, which protects data from flooding (fully submersed in fresh or salt water up to 10-feet deep for up to 72 hours) and fire (up to 1,550°F for up to 30 minutes).

If you wish to learn more about backing up your data, two free publications are available: the Best Backup Practices Booklet and the Backup, Recovery, and the Cloud eBook.

Don't Worry, Be Happy (and Secure!)

Although security is a complex business, even non-government organizations can take some steps to enhance their security footing. The first step is to adopt a zero-trust (ZT) security model, whose underlying concept is "never trust, always verify." This means devices should not be trusted by default, even if they are connected to a managed corporate network such as the corporate LAN, and even if they have been verified previously.

The next step is to recognize that computer data can be in one of three distinct states: *data in transit*, *data in use*, and *data at rest*, where the latter refers to data that's physically housed in a storage device. Organizations tend to focus on protecting their data in transit and data in use with tools like firewalls and antivirus software, but it's also necessary to ensure the security of their data at rest (DAR).

Today, protecting DAR is understood to be a critical piece of a ZT solution, but fully protecting DAR is a non-trivial matter. The solution is to use self-encrypting solid-state drives (SEDs) that contain a hardware encryption engine (EE). This EE encrypts data as it's written onto the disk and decrypts it again when it's read off the disk.

Very importantly, it's necessary for the SED to support multi-factor authentication (MFA) and pre-boot authentication (PBA). That means even the operating system (OS) is encrypted and the drive and system won't boot until the MFA has been satisfied.

Some drives, like FIPS-certified DIGISTOR's C-Series drives, can be used to enforce MFA, such as facial recognition, to enable users to even open, copy, or transmit files. These drives also maintain internal hardware logs of any file accesses that can be used for forensic purposes if required.

It's easy to become discouraged by the complexity involved. Fortunately, most of the time, making all of this work can be classed as SEP (someone else's problem). For anything that isn't SEP, we all need to diligently work to keep our nation's secrets safe and secure.

Throughout his career, Ben Warner honed his cybersecurity expertise working with the United States military. He worked on projects involving security and protection of networks, holding some of the nation's most sensitive and classified information with Applied Research Solutions at Wright-Patterson Air Force Base. He's also worked with Booz Allen, a leading cyber defense contractor, GE Aviation, and is a veteran of the U.S. Air Force.