

Mind the (Air) Gap

Published: 2021-05-13 · Archived: 2026-04-05 17:14:38 UTC

Following the ransomware incident [impacting Colonial Pipeline operations](#) in May 2021, many parties asked how such a disruption, impacting one of the main arteries delivering refined petroleum products to the Eastern and Southeastern United States, could occur. Based on information available, [the intrusion did not directly impact Industrial Control Systems \(ICS\)](#) within Colonial’s environment. Instead, the company itself initiated a controlled shutdown of operations as a precautionary matter to prevent critical ICS-related and [product tracking or billing systems](#) from being impacted by the ransomware event. Yet even though the criminal entities responsible did not or were unable to directly modify, disable, or otherwise disrupt ICS equipment, Colonial operators and defenders believed there was a non-zero chance such action could take place.




Based on the above circumstances, many persons asked a seemingly simple question: why were critical control system devices even accessible from Colonial’s IT network in the first place, resulting in the precautionary shutdown? This question brings ICS-related security back to the discussion of the “[air gap.](#)” Frequently lauded as an impenetrable standard in securing sensitive networks but nearly always misunderstood, an air gap – physical separation and lack of any connectivity between the sensitive network and all other networks – provides a deceptively easy solution for network defense. Yet further exploration of just what an air gap entails and how to actually implement one rapidly demonstrates that, although applicable in certain situations, as a general security practice (even in critical infrastructure environments) this practice rapidly becomes untenable.

First, what is an air gap? As described briefly above, an air gap represents the complete, physical isolation of a given network from all other networks. Only through physically bridging the network, such as through the introduction of removable media, can data flow in or out of the network in question. While not completely secure, as there are multiple known instances of [self-propagating](#) or [air gap-jumping](#) intrusions going back nearly two decades, such absolute segmentation will defeat the vast majority of network intrusions by virtue of simply not having an accessible network in place. Yet as tantalizing as this seems, actual, true “air gapped” networks are exceedingly rare, and frequently contain links, whether intentional via diodes or unintentional via [“shadow IT.”](#) that make cross-domain traffic possible.

For example, there is an assumption that classified networks, such as those used by US military and intelligence communities, are air gapped from the wider internet. Yet if we think of the definition of air gap – complete, physical segmentation – it rapidly becomes clear that this is simply not possible. Short of building a completely separate system of communication links (fiber, undersea cables, etc.) and avoiding all “over the air” transmission (satellite communications, wireless, etc.), the systems cannot be air gapped. Given that the US government does not operate a completely parallel internet, traffic from these networks “co-mingles” on public lines. While the traffic may be secured, and ingress and egress to on-site networks controlled via security solutions and diodes, traffic does exist and “cross-domain” traffic is possible, even if exceedingly difficult to achieve.

The problem with the above scenario, and what prevents the creation of an actual air gap for sensitive networks, is the distributed nature of the organizations involved. In order to enable not just networking within CIA headquarters or the Pentagon, but to extend such communication (including sensitive communication) to ships at sea and Forward Operating Bases (FOBs) in Afghanistan, connectivity outside a plausible, possible air gap environment must exist. Even something as seemingly isolated as a submarine still has connectivity to a network – via VLF signals while deeply submerged and via various over-the-air signals (connecting to various networks, tactical or otherwise) when surfaced or floating a communications buoy.



However, some environments exist where true air gaps are not only plausible, but actually required. The standard example is [reactor control for nuclear power generation](#). While other elements of nuclear power stations, such as those items directly interfacing with the overall electric system, will be networked (for reasons we'll explore shortly), reactor control systems are required to be isolated in nearly all countries that operate such facilities. Combined with significant [passive engineering safety controls](#), the actual reactor environment is designed to be operated as a completely isolated system, with strict measures in place for the introduction of outside material (e.g., software updates). Yet this environment represents something that the overall US military network is not: a single, centralized, specialized environment that can meaningfully be isolated from other systems. So long as heat can be exchanged and steam generated (indirectly, as in the [pressurized water reactor](#)), the reactor can “communicate” all that it needs to drive turbines and generate electricity with no further intervention needed.

What about pipelines, such as Colonial? To understand why an air gap is nonsensical in this case, we need to understand that a pipeline consists of multiple, communicating components which must work together as a system to enable monitored, safe, reliable operations. The combination of pipeline collection, [gate station](#), and [compressor station](#) components must communicate to a [control center](#) to govern the entire system, with sensor data and other feeds enabling visibility into operations. For a commercial pipeline that is thousands of miles long with hundreds of such facilities, an “air gapped network” is simply implausible, and if implemented would likely introduce sufficient operational friction and delay as to make operations less efficient and potentially even less safe. Short of building an entirely self-contained, geographically distributed network for the pipeline, a true air gap is at minimum prohibitively expensive. When combined with the need to communicate with suppliers and customers at machine speeds during product transfer and delivery, the air gap would need to be extended even further to up- and down-stream environments. In short, aside from building a completely separate, dedicated internet covering the entire oil and gas sector, along with substantial portions of electric generation, it is simply infeasible to implement a true air gap in pipeline (and related) environments.

One key aspect of this discussion is that critical infrastructure systems are rarely single point, isolated items in physical and network space. Like the pipeline and defense network examples above, the electric system (or “grid”) is similar in that it consists of multiple, interlocking and interdependent components from generation to [transmission to distribution](#) to function properly. Monitoring and communication at machine speeds enable for rapid responses and corrections to worrying or damaging conditions such as frequency or phase [anomalies that may imperil synchronization](#), or the disruptions created through interruptions in any part of the system. While it is possible to run such systems through pure physical controls, with fully manned substations and legacy [electromechanical relays](#) among other items, such operations are significantly less efficient, more expensive, and in most situations less safe than a modern [Energy Management System](#) (EMS) remotely coordinating multiple elements of physical grid operations with safety and protection provided by [digital protective relays](#). As with the other examples, the only way to “air gap” this network would be to build an entire dedicated network encompassing the entire electric system, from the largest publicly listed utility company to the smallest local distribution cooperative.

While air gaps are largely irrelevant, impossible, and even undesirable in many instances, we must recognize that connectivity and utilizing common infrastructure for communications still poses risks. Vocal proponents of air gap implementation will use a false dichotomy to argue their point, attempting to convince those with less knowledge of these systems that circumstances resemble an “all or nothing” proposition between complete connectivity and

complete isolation. Yet this is ridiculous and hardly representative of reality. As discussed with the military and intelligence network situation, you can have (reasonably) secure communications with a FOB in Central Asia from suburban DC provided the link and its components are properly established.

While a pipeline operator or water utility may not require that extreme level of security seen in protecting [JWICS](#) or other networks, relatively simple and widely available solutions can at least work to ensure attackers cannot hit Programmable Logic Controllers (PLCs) directly from the internet, and limit exposure in cases like the Colonial event. Short of deploying diodes and similar technology, robust network segmentation and authentication schema, such as use of a Virtual Private Network (VPN) with Multi-Factor Authentication (MFA), combined with fundamental Network Security Monitoring (NSM) will yield profound improvements in communication design.

The problem is not that PLCs and related equipment are accessible, it is that they are *trivially* accessible in many instances, such as the various intrusions into water treatment facilities in the [US](#) and [Israel](#) over the past year. That critical infrastructure devices can be enumerated via Shodan and directly accessed over the internet or from corporate IT networks is not a sign that such devices need to be air gapped. Rather it is evidence that such networks must be better designed and introduce improved controls. Once adversaries demonstrate the ability to consistently and easily subvert properly deployed VPNs and robust MFA implementation, then we as defenders and asset operators can talk about more robust solutions. Even then, something approaching the military model, of secure and encrypted communications tunneled over public networks, will likely be preferable than a true air gap solution necessitating the construction of a dedicated critical infrastructure network spanning continents.

Overall, critical infrastructure security has problems, and the Colonial Pipeline ransomware incident shows what several of them look like. But instead of immediately clamoring for an absolutist remedy that would be prohibitively difficult (and likely impossible) to implement, we as defenders and decision-makers should instead consider solutions closer-to-hand first. By applying fairly standard controls and network design principles, many of the supposed benefits of air gapped installations can be realized while also retaining the significant benefits of connectivity and communication over existing network infrastructure.

Source: <https://pylos.co/2021/05/13/mind-the-air-gap/>