# Solving a Brew Mystery: Digital Forensics With The Dragos Platform and OSIsoft PI System

**This paper presents a modern challenge of defending an industrial system, using situational awareness to detect and understand if an attack exists against the environment.**

Manufacturers have had to deal with sabotage since the dawn of the industrial revolution and the realities of the modern-day work environment. A changing cast of threats, technology, and human factors is compounding the forensics problem. Time to discover and unwind potential incidents can take weeks, if not months, of deep inspection by threat hunting experts and plant engineers.

One solution is to wait until the vendors' product lines provide data sources that support forensic reconstruction. This is beginning to occur now; for example, audit logs are increasingly included in embedded devices such as PLCs where they did not exist before. Unfortunately, industrial environments have a life cycle of 15 or 30 years and benefits won't be apparent for years or decades. Traditional security technologies such as firewalls, remote access solutions and managed environments, such as Microsoft Active Directory, are also becoming more accepted as good practices integrated by asset owners and vendors. This technology is essential and necessary, but also not sufficient in telling the entire story of an attack: impact to operations.

Another answer addressed here is to find novel data sources to evaluate and develop forensic timelines. There is a common phrase used in cybersecurity to explain how an attacker operates: "living off the land." This phrase describes the typical practice where an attacker will rely on the native functionality of the operating systems and software in a victim environment, rather than bring tools with them. Living off the land gives them the advantage of not raising alarms or suspicions by using the hosts, rather than abusing the hosts. This evasion is vital if the attacker will be accessing the environments for weeks or months. Secondly, using native functionality of a system, such as Linux or Microsoft Windows, is a transferable skill set. It is familiar to many, if not all, potential victim environments and is more approachable than a custom malicious software implant (requiring investment resources) that can be suddenly mitigated across the world once identified by researchers or the security community.

Defenders can also live off the land with similar benefits. It lowers the investment required to have the ability to detect and repel an attack, while lowering the threshold of training to technology. For industrial environments, living off the land can tell a more accurate and complete story at host, network, device and process perspective all through using the capabilities of the industrial environment rather than extending them.

## Introduction

Forensics is the act of analyzing a set of data to, in our case, determine the sequence of events surrounding a presumed cybersecurity incident. Once these events are discovered and evaluated, a story emerges to objectively describe the attack. This paper will describe a fictional facility and attack. We will then deconstruct the attack through native datasets of the environment, demonstrating how defenders can live off the land and the value of it to the investigation.

## The Case of the Inside Job

A fictitious brewing company "DBL HOP Brew Co." suspects foul play. The yield has fallen steadily over the last month, ever since a team member, Derrick, quit showing up for work without notice. Derrick was responsible for producing the brewery's core lineup of beers and worked the late-night shift. Derrick loved his beer and talking about the "dark web." The brewmaster, Avery, suspects tampering with the control system. Richard, their control engineer, has run numerous checks on the system–there is no evidence of tampering. Richard even upgraded the system to the latest version touting more security features. Still, Avery's years of experience tell her it must be some system issue.

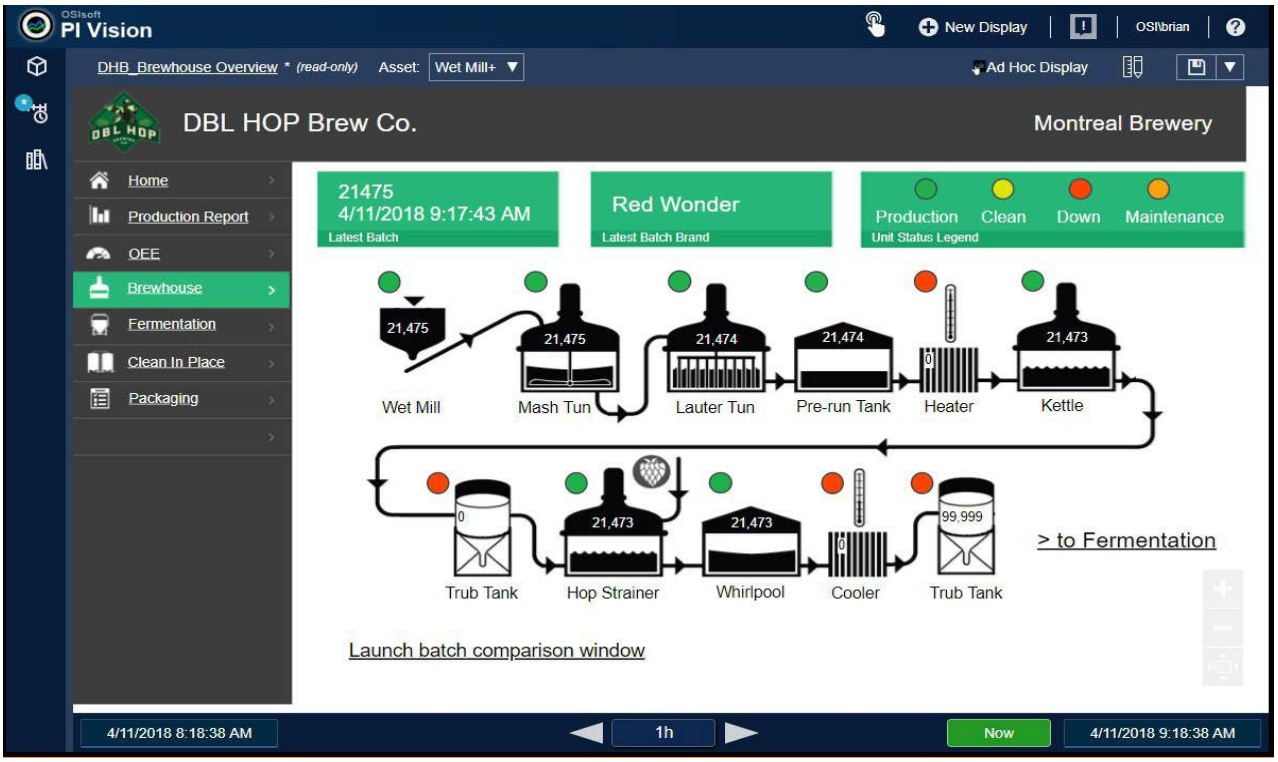


*Figure 1: DHBC's PI Vision HMI*

Greg, the owner of DHBC, spared no expense on his passion project; it is modern and highly automated. Process data from the brew system is collected and stored for analysis in the PI System by OSIsoft. This allows Avery and Richard to keep an eye on all process variables: volumes, masses, temperatures, times, etc. Event Frames have been created, as well, to help identify events like equipment startups and shutdowns, operator shifts, and product tracking batches. The system also has the Dragos Platform deployed for OT network security monitoring. This enables Avery and Richard to: **identify** all assets and communications in the brew system; **characterize** what those assets and protocols are; **detect threats in** the brew system using threat behavior analytics; **respond** to any potential cyber events using investigation playbooks; and **expedite recovery** so that the beer can keep flowing with an understanding of the root cause analysis of issues.

## Investigation

Avery tasks Richard with investigating the decrease in yield. He starts off with the following hypothesis: "an insider utilized control system features to negatively impact the beer production."

The first step in Richard's investigation is to verify if any weird internet traffic is seen to and from the brew system. The notification manager shows that there are not any new threat behavior analytic detections relating to external threats. Threat behavior analytics only occur when known attack tradecraft and methods have been observed by the Dragos Platform, which helps him understand the impact was not likely caused by an external source. Richard reviews the interactive map baselining features and quickly identifies what the pattern of normal operations looks like.

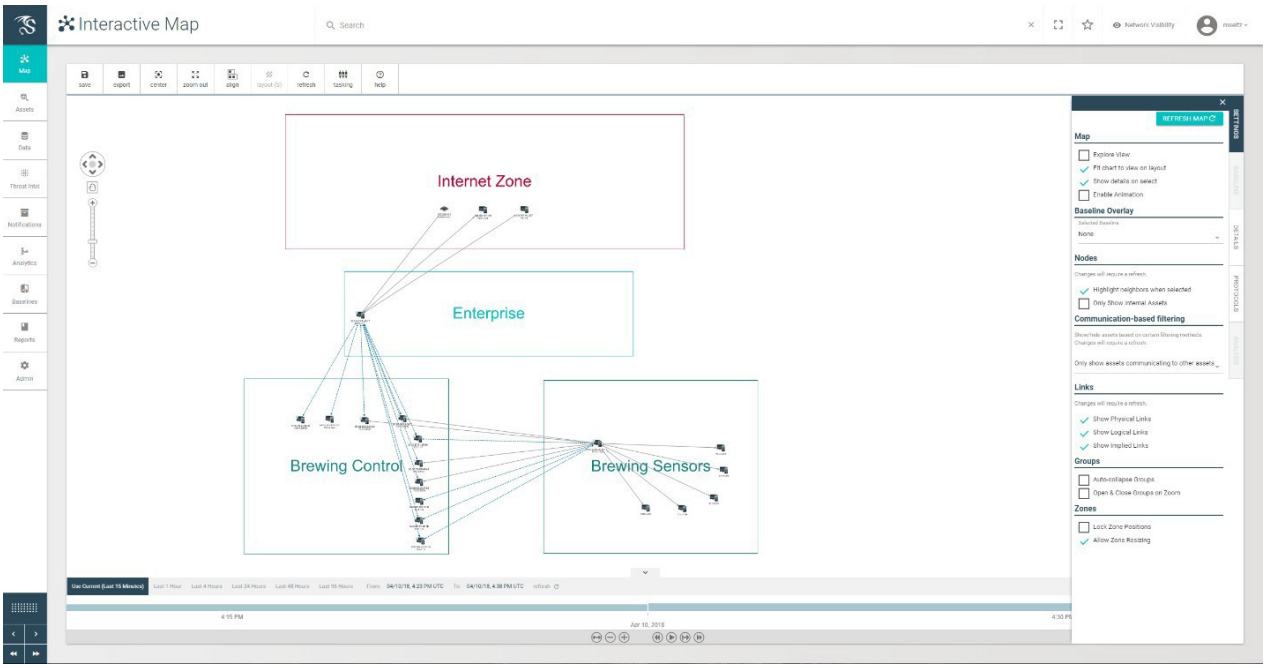


*Figure 2: Dragos Platform Interactive Map*

The remote communications before the decrease in yield are almost identical to the current, except for a few late-night web requests to brewing forums from the main brew system HMI computer. Derrick had been looking up some advice on fixing a problem with the grain ratios, unbeknownst to Avery. While this is slightly out of the ordinary, it doesn't seem like this web browsing activity was malicious or impacted operations directly. Richard concludes there aren't any glaring indicators of compromise while looking at remote communications and moves on with his investigation. The problem appears to be internal to the brew system.

Next, Richard looks at the host logs generated by the HMI. The Dragos Platform allows him to parse through the Windows logs, looking at the night in question. Feeling a bit rusty on Windows logs investigation, Richard loads up the relevant playbook. This playbook guides him through the high-level steps required for diving deep into the logs. He can clearly see that Derrick authenticated to the host and that he loaded up the HMI application. Through the playbook, he can also correlate these host logs with the previously observed remote

communications. No malicious processes appear to be running; however, he does find that Derrick installed a chat program that same night. Richard does not think this finding is relevant to his investigation and moves on.

Richard decides to look at the OT network activity on the Dragos Platform. He first loads a QFD centered around OT protocol communications to and from the HMI. He also tightens up his scope around the night shift. Traffic looks relatively normal as commands are being sent from the HMI to the various pieces of brewing equipment: load the grain mill, start the grain mill, stop the grain mill, transfer grist to mash tun, fill with boiling water, etc. However, traffic starts to look unusually light on the day Derrick allegedly walked off the job. It looks like one of the last commands issued was to load the grain mill—not a normal stopping point for the process.

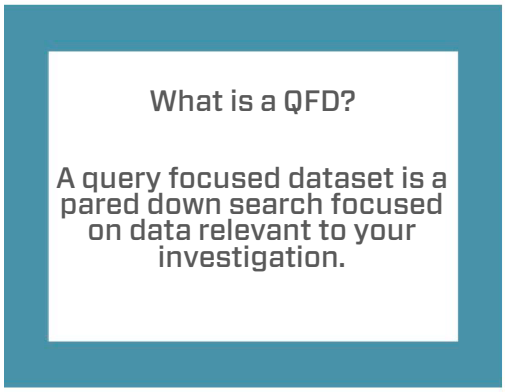


DHBC has a few motion-activated security cameras around the brewhouse. While in the Dragos Platform, Richard loads the investigation playbook that guides him through analyzing traffic generated by those cameras. The first step points him to the QFD where he can see several activity spikes which look normal and consistent with people walking around the brewery; however, he notes a sharp drop off around 10:00 pm. The next steps of the playbook lead him to confirm this by cross-referencing the footage in their security closet. Derrick can be seen checking in and starting his shift at 6:00 PM. Avery is also there, later than usual, filling in paperwork. Around 8:00 pm, Avery meets Derrick in the brewhouse, they have a beer together, and she leaves for the night. Richard fast-forwards through the footage until around 10:00 PM and confirms there are no more recordings.

As there is a clear impact to operations, Richard decides to look for ground truth: what does the process data look like before, during, and after Derrick walking off? Thankfully, the PI system is integrated with the Dragos Platform. This makes it easier for Richard to investigate, and he loads up a QFD used to search PI for the very last setpoint change on the day in question. Richard also notices through the notification manager that a threat behavior analytic fired on the night of the event, due to a sharp change in the process data preceded by abnormal network communications between the HMI and the grain mill. Having previously ruled out external communications, this is indicative of Richard's initial theory: an insider impacted operations. Richard now has a better estimate of time 0: 10:05 PM.

Richard compares batches before and after time 0. Using the PI data, he loads the historical data for the first few batches before and after and generates an average for volume, temperature, mass, original gravity, and final gravity data points. He then starts plotting out this data, and the trend shows a subtle but persistent change in batches. Not only that, the expected alcohol by volume for the batches is always low. Upon performing some mass balance analytics, Richard uncovers a problem in the amount of grain added to each batch. The grain weigh hopper isn't measuring the expected amount of grain.

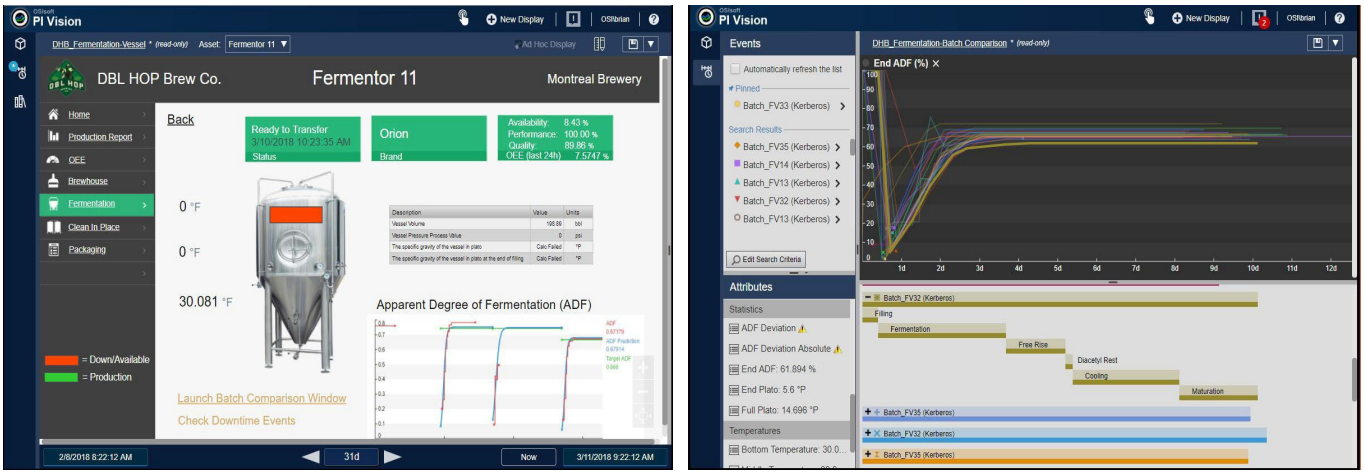


*Figure 3: Batch Comparison*

This leads Richard to perform calibration of the grain weigh hopper. He feeds a full bag of grain in the hopper, which should read 55 lbs. Unfortunately, it shows roughly one average adult's weight over the expected reading. Fearing the worst, Richard grabs his flashlight and inspects the hopper, uncovering the tragic events that happened the night Derrick disappeared.

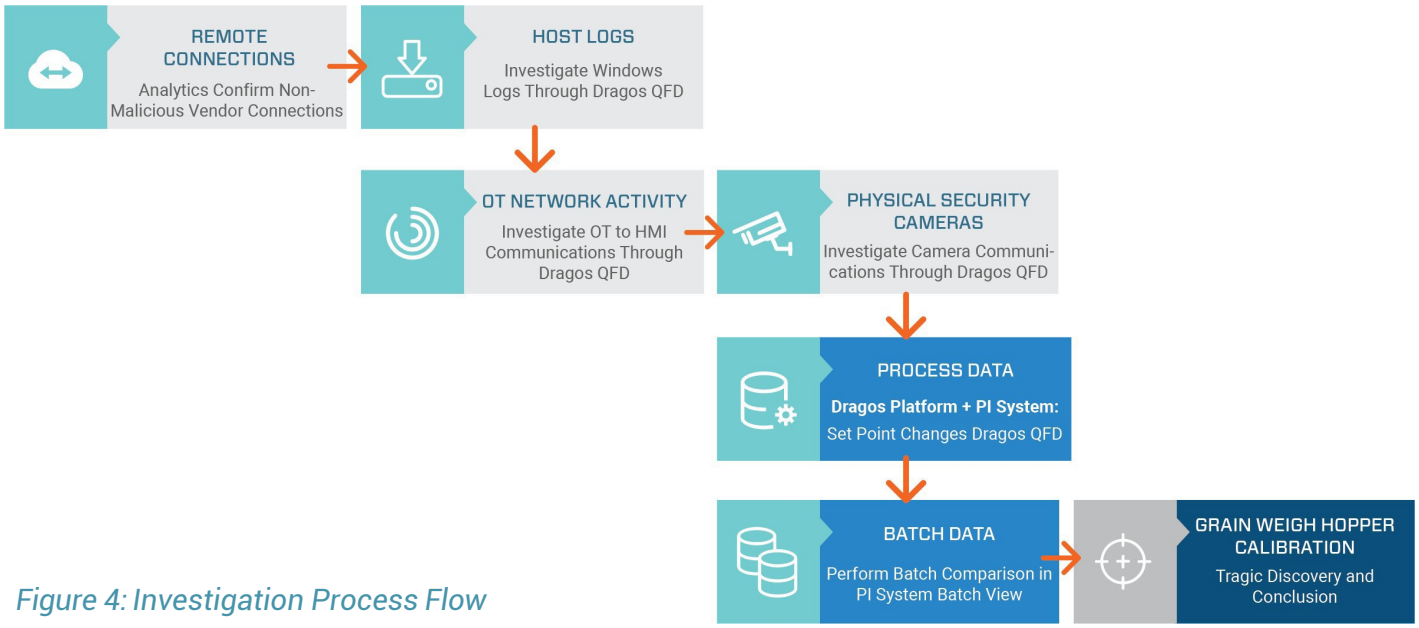


*Figure 4: Investigation Process Flow*

## Conclusion

Through this fictitious scenario, we show how defenders can use industrial data sources, namely process and batch data, to enrich and correlate with cybersecurity data during an investigation. The integration between the Dragos Platform and the PI System allowed our investigator to quickly and effectively parse network, host, and process data. The analyst can explore multiple data sources, which are typically siloed, and can pivot through them based on time, specific hosts, or other interesting data points. In the end, a conclusion to our investigation was achieved by utilizing existing data and systems from the environment, or living off the land.

## Acknowledgement

This brew story line credits the late Dr. Octave Levenspiel, who literally wrote the book on chemical reaction engineering. His elaborate quizzes would introduce seemingly superfluous details, like height and weight of the operator. Although his story lines stumped nearly everyone, learning to expect the unexpected was perhaps best lesson ever! Rest in peace, Octave.