

Supercharge Incident Response with DDI Visibility

Written by **Matt Bromiley**

November 2020

Sponsored by:

Infoblox

Is Today's Setup Working?

Has your organization suffered a data breach within the past 24 months? If so—and we believe many have—what process did your team use to respond to that breach? What data types were involved in the response process, and how long did the attacker have the advantage in your environment?

In the 2020 SANS Enterprise Cloud Incident Response (IR) Survey, approximately 58% of survey respondents indicated that an attacker had at least two days of undetected time within the environment, with a quarter of those respondents admitting their dwell time was *at least one month*!¹ While these numbers have come down from years past, the fact that we are still measuring dwell time in days and months is an indication that something needs to improve. Furthermore, this is only one metric—and only half the battle. Response teams must be able to effectively remove the attacker from the environment to make a difference.

Maybe it's time to explore a more efficient set of data and automated ways to access that data to assist the various IR stages and overall enterprise security. Many organizations suffer from a lack of data visibility and try to gather information about a breach manually after it occurs, which can be time consuming. This approach is not sustainable and is likely contributing to the long response times we see today. If the year 2020 has taught us anything, it's that in a month our environment may look nothing like it does today.

Flaws in IR do not necessarily mean the team is doing a poor job. The team may lack timely access to the data it needs to do its job effectively. This whitepaper looks at data that can empower your security team(s).

¹ "2020 SANS Enterprise Cloud Incident Response Survey," (September 2020), www.sans.org/reading-room/whitepapers/analyst/2020-enterprise-cloud-incident-response-survey-39805

Ironically enough, one of the simpler and more efficient ways to gain an advantage over attackers—and control of your environment’s security—is to utilize *the data you already generate and own*. In this paper, we explore how organizations should rely on and incorporate the following data points into nearly every aspect of their security approach:

- **DNS (Domain Name System)**—Requests and responses being made from everything running on your environment
- **DHCP (Dynamic Host Configuration Protocol)**—Network management protocol that dynamically assigns IP addresses to all the devices on your network
- **IPAM (IP Address Management)**—Planning, automating and managing the use of IP addresses and resources within your network

These three data points, which we will refer to as *DDI*, are the most valuable assets when it comes to understanding the scope of a breach within complex or modern enterprises. Many organizations suffer from a lack of visibility and/or environmental awareness; this data can help solve that problem, offering game-changing impact on any security program.

As you work your way through this paper, we ask you to consider the following:

- Are you comfortable with your enterprise visibility? Can you account for each asset on your network?
- How quickly can your security team determine the scope of a cyber incident?
- What environmental data points are available to your security team, and how are they utilizing them?

These questions, and more, inform the purpose of this whitepaper: Can we use what is already available to us to better secure our organization? It is time to find out!

Modern Problems Require Modern Solutions

Before we start diving into techniques and approaches for DDI, it is worth reflecting on how much the modern enterprise has changed. Prior to the year 2020, many organizations were already migrating services to the cloud, using SaaS applications and forging hybrid, global infrastructures. Many were also taking advantage of newer technologies and concepts, such as containerization and/or microservices, both of which come with their own inherent security concerns (which DDI data can also help solve).

Additionally, many organizations made infrastructure decisions based on current or impending regulatory requirements. For reference, consider how much of your security program was designed according to specifications laid out via NIST, NISD, CIS or other information security recommendations. During the design of your security program, how many assumptions did the architects make about on-premises versus cloud assets? Luckily, as we discuss in the next section of this paper, DDI data contributes to *all* these security frameworks, because most of them are rooted heavily in visibility and asset identification.

Fast forward to 2020, and the best-laid plans may have gone unnoticed. Many workplace transformations—adoption of SaaS, cloud migrations and the like—accelerated as working from home became the new norm and significant changes across the technological landscape became apparent. Entire enterprises have moved to remote operations, with much of the world figuring out how to deliver services online. Some industries have seen a devastating shakeup, while others have thrived during these troubling times.

One thing has not changed: the objective(s) of the information security team. Regardless of physical user presence, the impact of COVID-19 or an increase in online/virtual services, the information security team is still charged with protecting the organization, its users, data and customers while supporting workforce productivity or transformation initiatives.

During these ongoing changes, how was the security team tracking changes to the infrastructure and your assets? How did these changes affect their capability to detect and respond to incidents?

Even if we remove the pressures of COVID-19 and return to the business infrastructure you had in late 2019, the same question applies. Business operations can have long-term plans and implementations, but unpredictable, global pressures can upset the best-laid plans. While your users and customers may adjust to the changes, your security team's mission remains the same.

Unfortunately, some teams utilize different defense techniques based on where their users are located, creating inconsistencies that can lead to gaps in detection. We think it's prudent to rely on data points that are ubiquitous, regardless of where you conduct business. If your workforce is in-office today and remote tomorrow, your security approach can and should remain the same. Next, we analyze how DDI meets these and other requirements.

As we have mentioned, data points that provide ubiquitous visibility irrespective of any business changes are critical to effective security. DDI encompasses not only some of the most impactful data types (because it clearly increases your enterprise visibility), but it also carries one distinct advantage over more complex, even costlier solutions: You are *already generating this data* within your environment. Your network team(s) is (are) likely already using it! Security teams simply need a mechanism through which to harness and act on DDI data. Figure 1 offers a breakdown of each DDI component.

There is a strong chance that something about your organization changed significantly due to COVID-19. Did the pandemic force unwanted changes—or did it simply accelerate an inevitable timeline? Think about the opportunity to seize on new data points and strengthen your security team's capabilities.

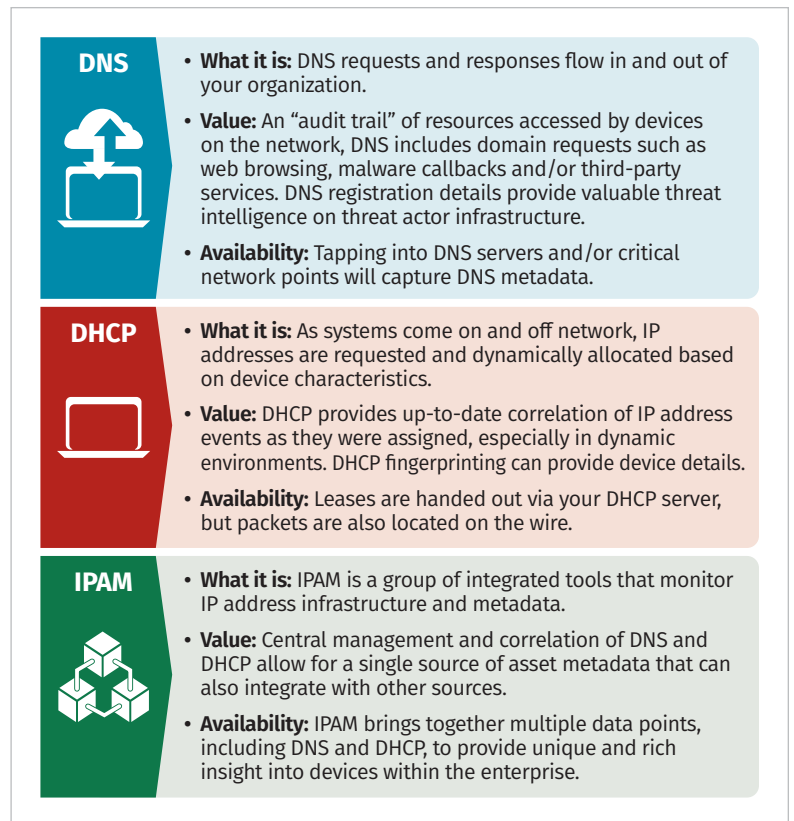


Figure 1. Breakdown of Each DDI Component

When we say you are already generating this data, that means you have systems providing IP addresses to devices as they connect to the network. Furthermore, to access resources, both internally and externally, these devices are already making DNS requests and receiving responses. This data needs to be harnessed, however. IPAM can go a long way toward this next crucial step. With IPAM, you are tracking every device with its IP addresses, network requests and internet destinations. Once you have this data harnessed, the next step is to utilize it to advance your security posture.

Enhancing Visibility and Asset Identification

When your security team starts harnessing DDI data, it will invariably discover new devices and services running within the organization. While we expect this realization—and quite frankly, that is the point—the DDI data is also a better representation of the modern enterprise, as we discussed previously. Gone are the days where system lifetime could be measured in years, months, days or even hours. Some operations may quickly spin up/down systems, services and capabilities, all for just minutes at a time!

How can you expect the security team to manually keep track of all these operations? They can't. Enterprise visibility suffers and creates more problems than it solves. This approach has led to an industrywide problem of redefining what visibility is. Some organizations are quick to label EDR or antivirus agent deployment as a representation of visibility instead of the solid network data within their own environments.

This labeling could not be further from the truth. There are numerous devices within any corporate infrastructure that cannot support an endpoint agent and, thus, will never show up in that population. Some organizations go as far as firewalling or segmenting areas of the enterprise where agents cannot be installed. Although this tactic is a good security practice, it does little to assist your security team with visibility into that particular asset. Figure 2 provides an example of various corporate devices and how DDI can assist in classification.

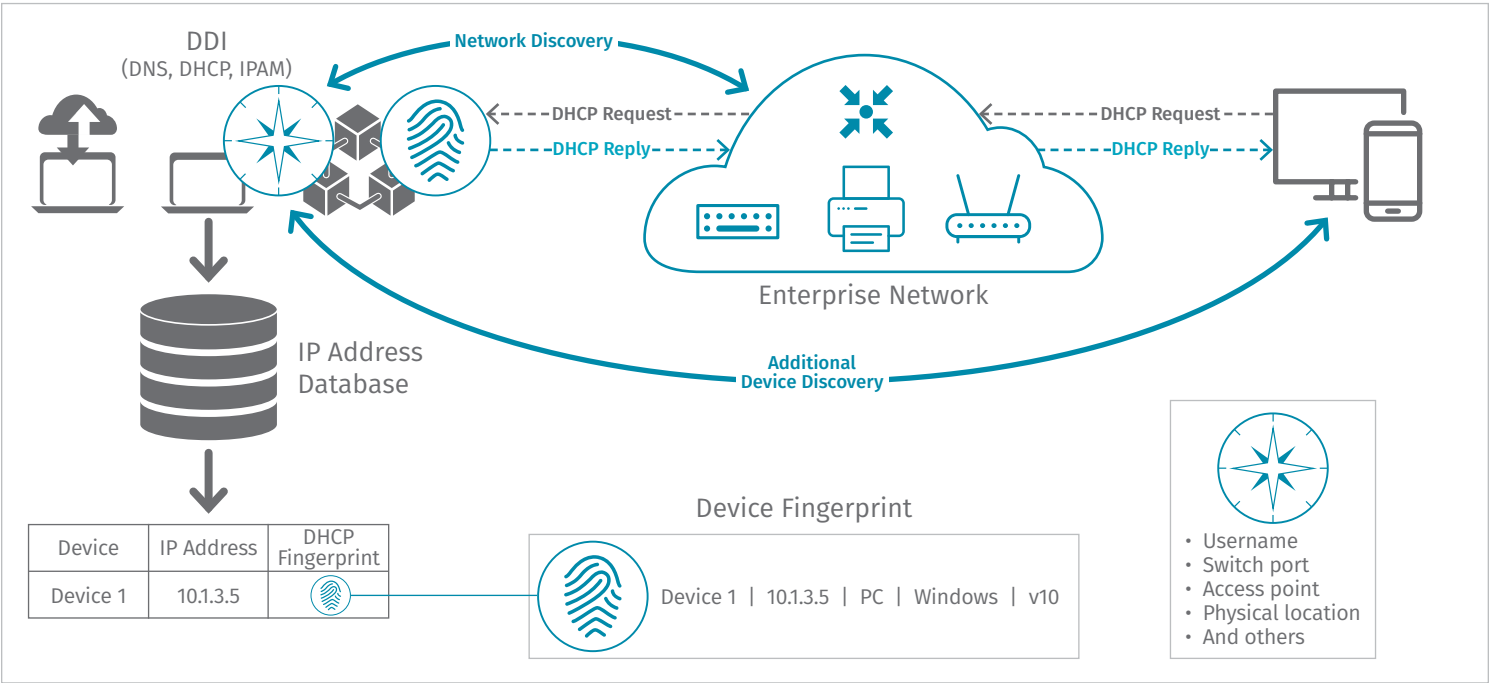


Figure 2. How DDI Can Assist in Classification of Various Corporate Devices

DDI data serves as the common ground for all devices requiring network access, which is 99% of today's enterprises. Even in air-gapped networks, some devices require connectivity, if only to each other and not the internet. After DDI data is harnessed, you can ask (and get a useful answer) at any point in time: Who is connected to my network, and what requests are they generating?

Enhancing the Six-Step Incident Response Process

When we break away from agent deployment and move to thinking of visibility from a DDI perspective, we also introduce a seismic shift in incident detection and response. When an organization relies on agent deployment as its source of visibility, it is placing trust in the detections and capabilities of that agent. Many devices within an organization cannot install agents: think IoT, OT and/or specialized, purpose-driven devices. While these devices cannot install an agent, they *can* serve as an entry vector for attackers. Essentially, your capabilities are only as strong as your agent allows them to be. You are not harnessing your own data—and attackers know this.

Attackers are aware of this weakness and use it to their advantage. Many attacker playbooks include searching for, avoiding and/or disabling endpoint security (again, assuming it is available). Furthermore, some attackers know that endpoint agents detect the latter stages of an attack, not the earlier stages. Thus, they adjust their techniques on the fly to accommodate the implementations found within your environment. But they cannot adjust their need for the network—whether they move east-west or north-south—and that is where our DDI data resides.

If you were to suddenly increase your security team's visibility into the environment via DDI, how would that change the team's approach to detection? What could your team do with harnessed, correlated DNS data? What if they could trace, in real time (via DHCP and IPAM), devices making suspicious requests? IR teams would be able to make faster and more efficient decisions in responding to cyber incidents.

To help identify where DDI data is useful to IR teams, let us begin with the typical SANS six-step incident response process, shown in Figure 3.

Next, we examine how DDI can enhance/improve each step along the way. We also pose “Key Questions” for each step. These provide your IR team with a means to assess its capabilities. Use these key questions to think about what data points you currently have and whether DDI would help you be more prepared.

Are you utilizing your available data to quickly correlate and respond to incidents, or are you simply waiting for a detection to fire? Waiting for an alert to fire is not visibility—it is reliance on a third party to tell *you* about *your* environment. Attackers know this and use it to their advantage.

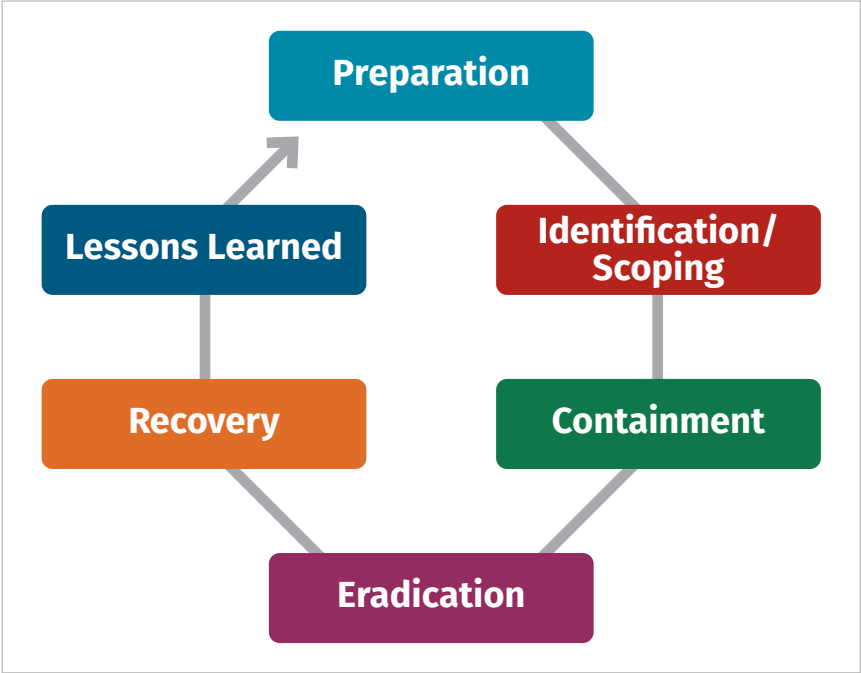


Figure 3. The SANS Six-Step Incident Response Process

Preparation

You will notice that the six-step IR process is repetitive, meaning the organization should improve from each incident.

Without a previous incident, preparation begins with understanding the environment. We have covered environment visibility and scope in a previous section, but it is worth restating: DDI visibility allows for more comprehensive incident detections and preparedness because the organization breaks away from depending on endpoint agents.

If an organization has experienced an incident, then the lessons learned from that incident are likely drivers for an increase in visibility, awareness and/or detections. As previously examined, DDI data provides unique options for these.



Key Questions to Ask

- ? How much awareness and visibility do we have into the environment?
- ? Without an incident, can we identify an IP address associated with an asset?
- ? What other metadata is available, and how can we enrich it to correlate against other available data points?

Identification/Scoping

After an incident, the primary focus of the IR team is to understand the scope of the attack and identify the affected systems. This step is one of the longest and most intensive during an incident, but it is one that must be completed as quickly as possible. This phase is often made long due to data collection, competing analyst priorities and other concurrent business activities. The longer identification takes, the more time an attacker has between detection and containment—and the more time he has to exfiltrate sensitive data. This is the very purpose of DDI data: to provide efficient identification of indicators associated with the attack and use them to identify the true systems affected quickly.

DDI data does not replace threat intelligence. In fact, DDI data is easily enhanced with threat intelligence and vice versa. With DDI visibility, you can begin making sense of *who* is talking to *where* and allow threat intelligence to help guide investigations and response.



Key Questions to Ask

- ? Where is (are) the asset(s) located?
- ? What type of asset(s) is (are) involved, and based on asset geography, how do local laws/regulations affect our ability to perform IR?
- ? Using DDI history and the known indicators of compromise, what other assets are involved?

Containment/Intelligence Gathering

Like identification and scoping, containment and intelligence gathering are rooted in observing the impact of a cyber incident. These cyclical subprocesses often feed off each other, which means you may bounce between scoping and intelligence gathering multiple times before identifying the full extent of an incident.

Of course, DDI data is crucial to this step as well. If an incident begins with detection of a malicious website, for example, identification may highlight one affected asset. Intelligence gathering analyzes the context of DDI data related to the original incident, which may lead to more indicators. Using this data to scope the incident may yield more indicators and return the analyst to the Identification phase. Figure 4 provides more context into this process.

After an incident has been properly scoped and indicators of compromise have been observed, this step also includes moving toward eradication of the threat. Depending on the type of incident and the response team, containment may be a slow-moving phase—to avoid tipping off the attacker.

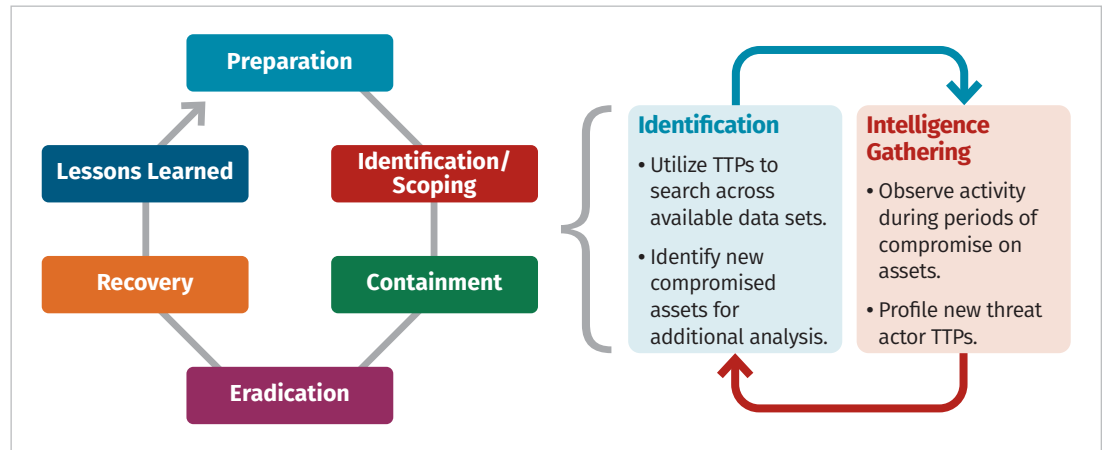


Figure 4. Cyclical Subprocesses Between Identification and Intelligence²



Key Questions to Ask

- ? What happened around the time of infection? Does DDI data (particularly DNS) provide any more indicators of compromise?
- ? Does external threat intelligence provide any insight we could use to identify additional indicators of compromise?
- ? With new indicators in hand, do we need to return to the Identification step to search for new systems?

Eradication and Recovery

When incident scoping has been effectively achieved and effective containment implemented, it is time to move the organization toward removing the threat from the environment. Eradication and recovery are often combined, because the organization remains in a heightened state of alert during and after removal of the attacker. The security team does not let its guard down after eradication, expecting a determined attacker to want to return to the organization.

² TTPs are tactics, techniques and procedures.

DDI data is a natural complement to these steps. Notice that earlier, we used DDI to detect and identify malicious activity and contain malicious communications. We continue to use the same data points to confirm that the attacker has left the environment and that malicious indicators of compromise are no longer observed. During the recovery stage, DDI will serve as the best way to monitor for persistent malicious activity.



Key Questions to Ask

- ? Do we see any additional signs of compromise in the environment?
- ? Are there other anomalies within the environment that may be indicative of an attacker returning?
- ? Did we effectively kick out the attacker?

Lessons Learned

As a final step, when an organization has successfully removed an attacker, the security team can begin to reflect on actions taken before and during the incident, as well as how the team can improve its processes going forward.



Key Questions to Ask

- ? Do we have mechanisms in place to detect the activity we just defended against?
- ? Did our analysis correctly identify and scope the point of entry *and* lateral movement, if any was present?
- ? Did we find gaps during our IR process, and how do we fix those gaps?
- ? What other data points would have made our response faster? How do we get access to those?

The impact of DDI data on the IR process cannot be understated. If you take these six steps and remove DDI data, your team is simply left pouring through giga- or terabytes of logs, hoping that correlation will allow you to piece the incident together. This takes time and, unfortunately, the longer your security team takes to investigate, the more time the attacker has in the environment.

Supporting IR Is Crucial to Success

Given all we have presented in this paper, the inclusion of DDI data in your organization's IR processes is not merely a nice thing to have. It should be regarded as a requirement for success. Currently, with environments constructed the way they are, we can no longer ask teams to secure the organization and respond to incidents without providing the data they need to be successful.

DDI data is successful when correlated and analyzed in concert. An asset is assigned an IP address via DHCP. That asset begins to make requests to and receive responses from domains via DNS. Other crucial metadata is collected and used to profile that asset. IPAM software integrates and tracks these data points, providing methods for responders to act on this data. This correlation should be automatic to enable analysts to respond, when necessary, in real time.

Case Study: A Tale of Two Analysts

Let us look at putting DDI data to work within a sample organization. It's one thing to know that DDI data is all around you—it's another to put it to work. In fact, one of the hardest tasks in information security is making the available data both accessible and actionable. In the following example, consider how your organization would react in each situation.

Alice and Bob are two security analysts who work at PKI Health Services, a medical provider. They manage a mixed environment that includes a wide range of servers, users and operating systems. As a healthcare provider, PKI is charged not only with protecting its users and systems, but also with protecting patient data. Because the healthcare industry has recently become the victim of prolific ransomware attacks, Alice and Bob are on high alert for an impending attack. In Figure 5, we examine their response to a potential ransomware incident.

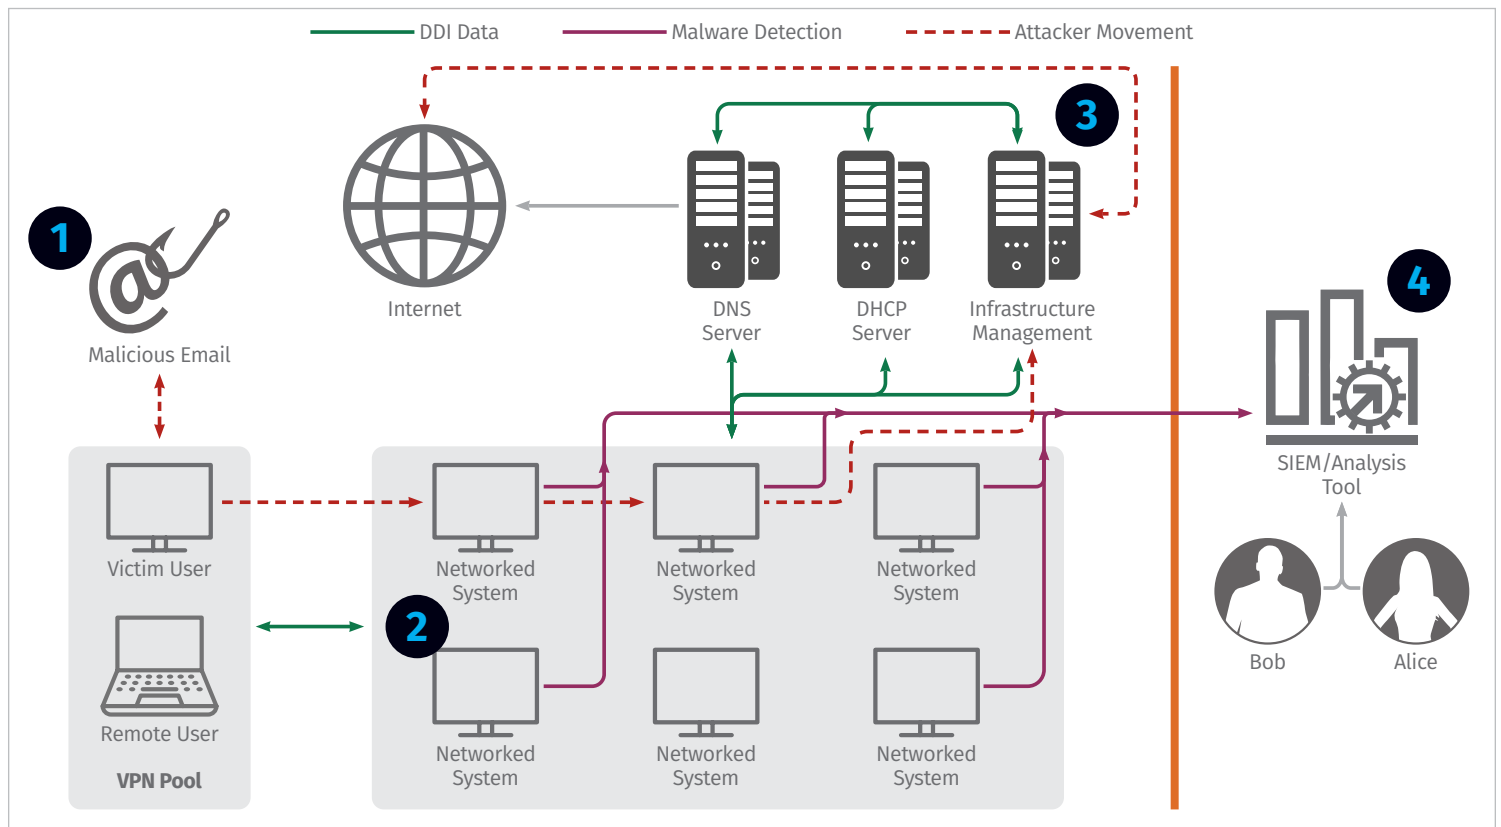


Figure 5. Case Study of Incident Detection Without DDI Data

It is easy to see that the security setup depicted in Figure 5 was based on third-party telemetry providing malware detection. When a user clicks through a phish (1) and allows ransomware to propagate throughout the environment (2 and 3), Bob and Alice are left awaiting malware detection telemetry (4), assuming there is any. The team operates in a reactive manner and is dependent on alerts to tell them where to look. Of course, they will need to perform a deep-dive investigation through the environment to identify the four compromised systems. **Did you notice that the original victim system, oftentimes referred to as *Patient 0*, is not sending endpoint telemetry to the security team?**

Let us look at PKI Health Services a few months later, after they have modeled their security team to rely on DDI data instead (see Figure 6).

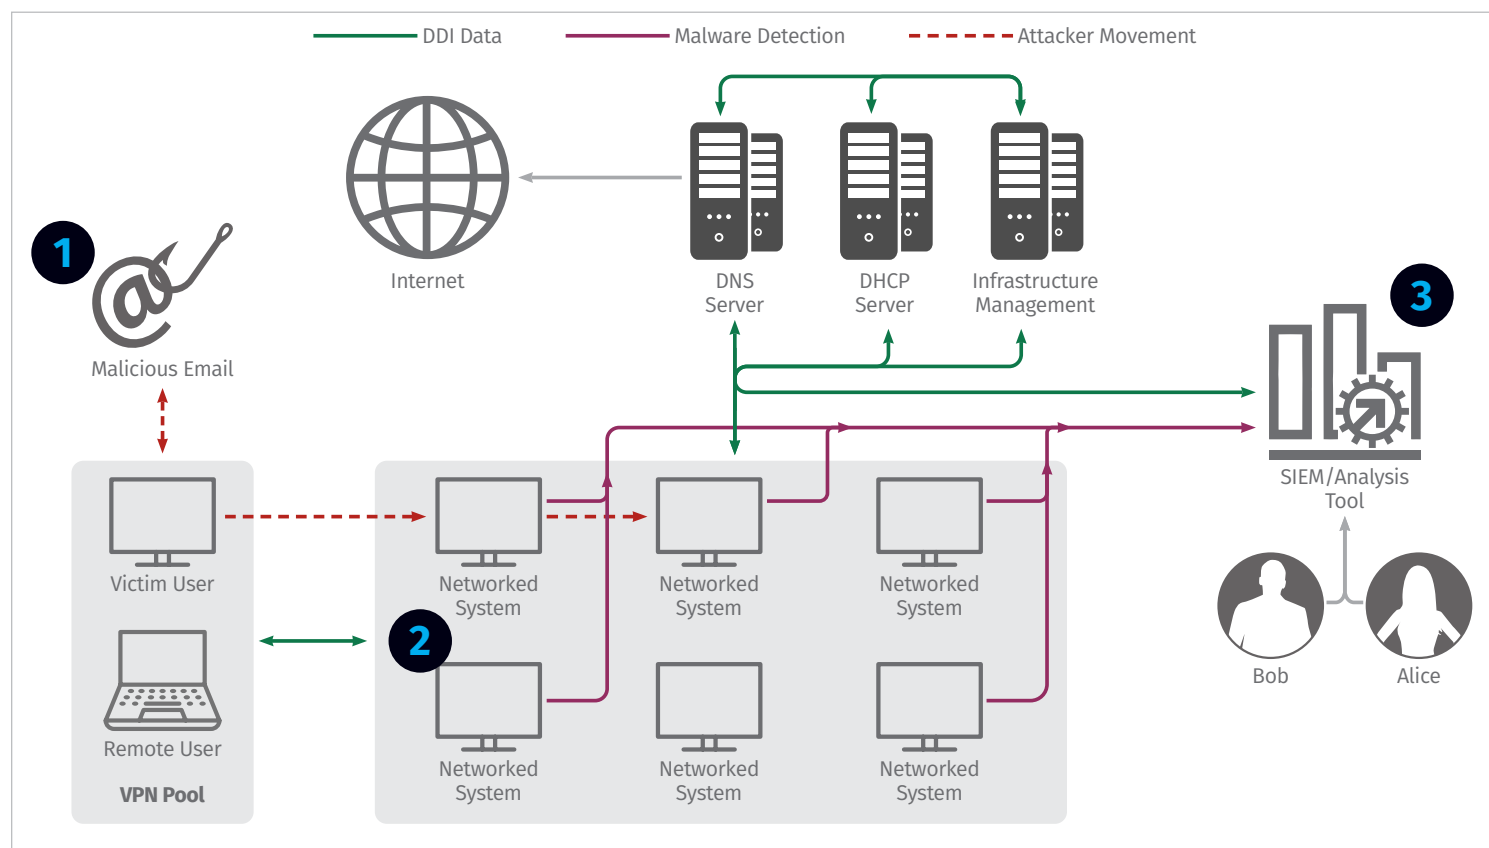


Figure 6. Case Study of Incident Detection with DDI Data

The differences in Figure 6 are subtle, but they are enough to have an impact on the security team's capabilities. Note that we did not introduce a new system; we simply provided the team with access to the data that allowed them to effectively assess and respond to an incident. When the victim user clicked through a phishing email (1), that activity was recorded via DDI data, which the security team is now receiving. Whereas the lack of endpoint visibility yielded a gap in detection, the team now has network data they can act on and pivot from. If the threat actor is still able to move laterally (2), the team is already aware of the unwanted presence. **They can quickly neutralize the attack, utilize the known network-based indicators to quickly scope the incident, and implement blocks where necessary.**

It is also important to note that we did not replace malware detection or endpoint telemetry with DDI data. In fact, we encourage the opposite—combine your other telemetry sources with DDI data, ensure that visibility gaps are covered, and determine which data points are enriched via correlative sources whenever possible. A win-win!

Closing Thoughts

In this whitepaper, we examined and discussed some of the most crucial data points for visibility within your environment: DNS, DHCP and IPAM, collectively known as *DDI*. This data provides the best insight into current and active assets within your environment, often going a step further to indicate physical geography, connection status and user activity. Concurrently, DDI data provides some unique visibility into threat activity, making it some of the best data to have on hand during response to a cybersecurity incident, regardless of incident severity.

Unfortunately, many security teams are handicapped when it comes to responding to incidents and doing so effectively. This is not to say that any team is doing anything *wrong*—it is likely that in the wake of an ever-changing business landscape, information security teams are simply trying to keep up. Furthermore, they may not be utilizing the best available data points to perform their job. Either way, it is time to make a change.

We should be providing our security teams with data points that are also adaptable and insightful. If teams are charged with defending every aspect of the organization, they should have visibility into every aspect. This visibility is hard to obtain with agent-based deployments but is readily available with a focus on network-based DDI, which is a foundational source of visibility data that can keep up as your organization continues to evolve the network to support modern workforce transformation efforts. With broad, reliable and consistent visibility, your security team will quickly realize that DDI data not only makes them more efficient at investigation and IR, but also allows them to think about and implement detections in a brand-new way.

About the Author

Matt Bromiley is a SANS digital forensics and incident response (IR) instructor, teaching [FOR508: Advanced Incident Response, Threat Hunting, and Digital Forensics](#) and [SANS FOR572: Advanced Network Forensics: Threat Hunting, Analysis, and Incident Response](#). He is also an IR consultant at a global IR and forensic analysis company, combining experience in digital forensics, log analytics, and incident response and management. His skills include disk, database, memory and network forensics; incident management; threat intelligence and network security monitoring. Matt has worked with organizations of all shapes and sizes, from multinational conglomerates to small, regional shops. He is passionate about learning, teaching and working on open source tools.

Sponsor

SANS would like to thank this paper's sponsor:

