INFOLOCK

When Shadows Cast Doubt:

The Hidden Problem of Shadow DLP

TABLE OF CONTENTS

CHAPTER 1
SHADOW DLP – NOT ANOTHER SUBSET OF SHADOW IT

Data Loss Prevention (DLP) has historically been more of a niche product, often overshadowed by other cybersecurity tools like antimalware and firewalls. Those tools were always a "must have", while DLP was commonly more of a "nice to have". The only notable exception, of course, has been organizations operating in highly regulated industries, such as healthcare or finance. However, something has changed in recent years.

The news regularly features headlines about data breaches, often affecting millions of people; and 2023 was a record-breaking year for such events, targeting organizations of all sizes and from various industries (Brunell, 2023). For example, mobile communications provider T-Mobile was attacked in early 2023, affecting 37 million user accounts (Alspach, 2023). And, in December 2023 alone, there were over 2.2 billion records affected, caused by 1,351 individual breaches (Ford, 2024). These are only the breaches officially published, the dark number can be expected to be substantially larger.

In response to data breaches continuing to be on the rise, lawmakers all over the world are creating new regulations or tightening existing ones to better protect people's personal data, even in less heavily regulated sectors (think GDPR). One key outcome of this development is that several new players offering data security as part of their bigger product portfolio have joined the market, making DLP the latest cybersecurity craze.

That's good, right? Well, yes and no. On the one hand, it is great to see DLP finally getting the attention it deserves. And making it more easily available to organizations historically not running a DLP program can definitely help to prevent data loss events. On the other hand, new issues have arisen as the market for DLP programs has increased. How should organizations deal with the fact that they might have two, three, or even more different DLP solutions readily available in their environment? What happens if several DLP tools are enabled without central oversight? Who will be responsible for the management and operation of all these different tools? Who will perform incident triage?

The recent trend of DLP tool proliferation is a phenomenon Infolock coined **Shadow DLP**. In this eBook, we will put a spotlight on Shadow DLP, examining it from different angles. In this first chapter, we will define what Shadow DLP is and how it may get introduced into any IT environment. In chapter 2, we will do a deep dive into all the different issues that Shadow DLP can cause. And, in the third and final chapter, we will introduce a phased approach to address Shadow DLP once and for all. Throughout this eBook, we will mention real-world examples from actual Infolock customers to further exemplify the root causes and results of Shadow DLP.

SHADOW DLP IS NOT SHADOW IT

Before we get started defining what Shadow DLP is, let's quickly clarify what it is not: it's not a subset of Shadow IT. Shadow IT can be defined as "any software, hardware or IT resource used on an enterprise network without the IT department's approval and often without IT's knowledge or oversight" (IBM, 2023). However, we can rewrite this statement to define what Shadow DLP actually is: **any software, hardware or IT resource offering DLP as a point solution or an add-on capability used on an enterprise network without the oversight of the corporate DLP program owners.**

There are two key differences between Shadow IT and Shadow DLP. First, Shadow DLP is limited to solutions that feature some kind of DLP capability offered either as a point solution that targets only one or a few specific use cases, or as an add-on to the service that this solution was purchased for. These are commonly cloud-based SaaS or IaaS offerings, such as a cloud proxy, email security, or cloud storage solutions. Next, and even more importantly, the IT department is fully aware of this capability, but not the DLP program owners. This means that the IT process owners start working directly with the solution vendor to get some level of visibility into their data streams, without further coordination with the people running the corporate DLP program. As a result, an isolated DLP integration is created.

You may wonder "How can this happen? Shouldn't those teams talk to each other to figure these things out before they could become a problem?" The answer is: yes, they should, and sometimes they do. In many such cases, organizational silos prevent information from making it to the right stakeholders. This is especially true in large organizations with various sections inside the IT department, where communication cannot flow directly as various directors, boards, and panels are getting in the way. This is particularly a problem when sections don't perceive there to be a logical overlap (why should the email team inform the endpoint team about an upcoming change?).

Another common way Shadow DLP gets introduced is when a salesperson forges a connection with someone like a CISO, or the director in a CISO's office, and persuades them that their DLP or add-on solution is the answer to one or more pain points. A purchase order gets signed without much further consultation with the responsible business units, ultimately resulting in yet another bright shiny tool in the toolkit.

Example Use Case

One of our customers had to replace their on-prem proxy environment with a cloud proxy solution. They had a very tight timeline to replace this essential piece of their network infrastructure, so they focused on proxy-related questions to make sure the new solution would satisfy their needs and use cases. Other integrations, such as with their primary DLP solution, were an afterthought. High-level questions like "Do we still have DLP visibility" were answered with that the new tool will have DLP, so

they would be covered. For the sake of time this answer was considered sufficient. Only, when they were just about to switch web traffic over, they realized that this would route it around their existing DLP infrastructure, while at the same time nobody bothered thinking about configuring the proxy's DLP feature. Thanks to our expertise as their Managed DLP (MDLP) Services Provider, we were able to create a solution to retain visibility without any outages. But this was a close call where they almost lost a substantial piece of their DLP visibility because the wrong questions were asked.

NOT ALL DLP ARE CREATED EQUAL

Before we dive any deeper, it's important to differentiate between the three (3) different types of DLP offered:

01

General DLP Solution

Its sole purpose is to identify potential data loss. It has a strong detection engine, a broad library of policy templates, comprehensive reporting capabilities, and provides detection methods for all relevant use cases, including data-in-motion (network), data-in-use (endpoint), and data-at-rest (discovery), for both on-premise and cloud-based use cases.

02

Point DLP Solution

Like a General DLP Solution, its sole purpose is to identify potential data loss. It has many of the same features as a General DLP Solution, but its coverage is limited to a subset of DLP use cases, such as being endpoint specific.

03

Add-on DLP Solution

Not dedicated to DLP but offers it as an add-on feature. Usually, a cloud-based networking or storage solution. Some only work with the solution they got shipped with, while others offer additional DLP coverage, usually in the form of a local agent.

The following table outlines a comparison of the three DLP types, to better exemplify their differences. Keep in mind that this is not an exhaustive list, rather a high-level overview to compare and contrast these different DLP offerings.

	General DLP	Point DLP	Add-on DLP
Vendor Development Focus	100% DLP	100% DLP	Secondary, superseded by primary SaaS/IaaS solution offering.
Implementation Effort	On-prem server and database management is very complex.	Varies depending on scope and coverage vector.	Cloud-based offering mostly managed by vendor.
Coverage	Versatile architecture to cover many vectors, but may have channel-specific gaps.	Only covers the specific vector it was purchased for.	Very well integrated with solution, but often has gaps for other vectors.
Detection	Often ships with advanced proprietary detection methods not found in other tools.	Very dependent on vendor. Most common detection methods are usually covered, but some may offer additional technologies.	Ships with the most common detection methods.
Incident Response	Dedicated console with elaborate filtering and triage option.	Dedicated console with elaborate filtering and triage options.	Part of the solution's management console, often lacking advanced filtering and triage options.
Reporting	Ability to create custom dashboards, charts, and metrics in management GUI.	Ability to create custom dashboards, charts, and metrics in management GUI.	Dashboards often predefined by vendor, but certain customizations possible.
Costs	Dedicated tool with a substantial price tag.	Dedicated tool with a smaller price tag, compared to General DLP.	Usually included at no or small extra fee, but some vendors charge extra for certain advanced features.
Summary	Best coverage, but often pricier and more complex to manage.	Decent compromise to address niche use cases, but not scalable as use cases expand.	Great solution if no other DLP coverage exists, but often lacks capabilities to expand to a full DLP program.

THE ROAD TO SHADOW DLP IS PAVED WITH GOOD INTENTIONS

Now that we have established a solid understanding of what Shadow DLP is, let's give an example on how it could sneak into any organization's environment. Let's check out the following chart.

Shadow DLP - Hypothetical Introduction Flow Chart

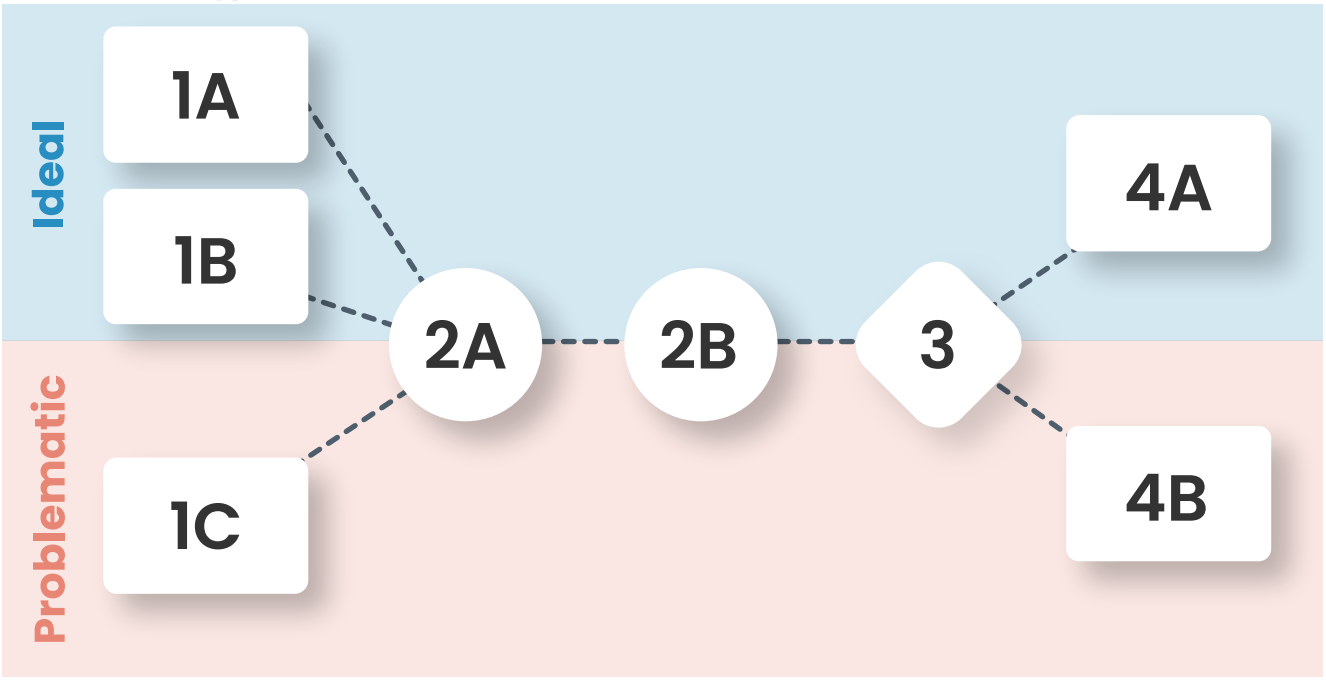

Each item in the flow chart above is meant to represent a hypothetical action in an organization's larger DLP environment where Shadow DLP is unintentionally introduced. Outcome 4A being the most ideal for addressing Shadow DLP, and 4B being the most problematic in terms of allowing Shadow DLP to run rampant in the environment. Definitions for each action can be found in the legend below.

Shadow DLP - Hypothetical Introduction Flow Chart Legend

- **1A)** A general DLP solution is acquired, and a DLP program plan is established.
- **1B)** A first point solution is purchased and enabled in production.
- **1C)** No existing DLP tools are in place at all.
- **2A)** A new networking tool is acquired and implemented that include Add-on DLP capabilities.
- **2B)** After some time, Networking team starts enabling DLP features that cover their IT Process use case(s).
- **3)** A SaaS storage solution gets licensed and turned into production. As part of the implementation, the owner of the SaaS process enables vendor-provided default DLP policies to "ensure protection."
- **4A)** DLP Program Owners become aware of Shadow DLP and start coordination/consolidation.
- **4B)** Various DLP tools continue to run side-by-side with no coordination efforts.

As we can see, how the story begins and ends may differ, but the middle part always follows a similar trajectory. Let's discuss these steps in a little more detail.

1) The journey either starts with no preexisting DLP Program at all, or an organization purchasing DLP to satisfy their initial data governance needs or address identified risks. This could be either a General DLP tool initiating a fully-fledged DLP program, or a point solution satisfying a small set of initial DLP use cases. This tool has a designated DLP Program Owner solely responsible for defining the detection policies and reviewing the resulting DLP violations. Additionally MDLP Services Providers are commonly used for these responsibilities, supporting the DLP Program Owner to further mature the DLP Program.

2) As time moves on, new tools are introduced into the environment, or existing on-prem tools get replaced with cloud-based alternatives. The primary focus of these tools is not DLP-related. Common examples are IAAS platforms, email security solutions, or cloud-based proxies. They all substantially change the way data flows through the environment or where it is stored. This poses the risk of reducing the established DLP tool's visibility into the organization's data streams leaving the corporate IT premises.

3) The IT Process Owners realize that there is a need to protect the data they process with their new tools. This may be due to the solution vendor highlighting this capability or a data breach reported on the news, just to name a few common reasons. Either way, the new tool's Add-on DLP capabilities are enabled, usually without coordination with DLP Program Owners. This yields in a DLP implementation that is substantially different from the baselines defined in a preexisting DLP program, often with a detection scope that is much more limited.

4) Eventually the DLP Program Owners may become aware of the new DLP tool(s) in their environment. This may happen during the final approval process when a ticket with a DLP flag ends up in their queue for approval, or when someone, such as the CISO, brings it to their attention. But often, due to project plans, deadlines, and/or budgetary constraints, they don't have a way of correcting this mishap. In a worst-case scenario, they might not even become aware of this development until much later. In the meantime, this Shadow DLP implementation provides a wrong sense of comprehensive security.

As we can see, there are several variables at play here. The result is the same though: multiple DLP tools are running simultaneously and without central oversight. Good old-fashioned vendor sprawl. In the next chapter, we'll go into more detail on why this isn't the best idea.

CHAPTER 2
THE MORE DLP WE HAVE, THE BETTER... RIGHT?

Before we tap into the bad and the ugly on why too many DLP solutions are not a great idea, it's important to highlight how multiple solutions can work together well, a perfect world scenario. For this to become a reality, all different DLP solutions must be:

- well-orchestrated, so they don't get in each other's way,
- well-configured, so they play to their individual strengths,
- well-understood, so that workflows can be harmonized across all tools while acknowledging unique differences for each tool, and
- centrally managed by the DLP Program Owner, to ensure consistent policy enforcement across the entire DLP landscape.

If all these conditions are met, multiple DLP solutions can work in harmony, playing to their unique strengths to provide comprehensive coverage. However, as we can see, these are tall orders that must be executed in highly complex environments. So, sadly, this scenario is rarely the case. Let's see what problems could be introduced with Shadow DLP.

TOO MANY COOKS IN THE DLP KITCHEN

As previously mentioned, the DLP Program Owner is often not directly involved when IT Process Owners start working with the Add-on DLP solutions that come with their tool. Instead, the IT Process Owners work directly with the vendors, either because the vendors advertise their DLP tool, or because the IT Process Owners feel the urge to use it because DLP came shipped with their product. They set up DLP and configure detection policies, not fully understanding the string of responsibilities that come with it. These responsibilities include creating the correct policies to start with, policy tuning, and incident response once the policies are not in test mode anymore.

As more and more on-prem solutions get replaced with cloud-based solutions featuring Add-on DLP features, here are some examples of who might introduce Shadow DLP into your environment, and how:

- Email team, as part of an email security solution
- Network team, as part of a cloud-based proxy solution

- Cloud infrastructure team, as part of an IaaS solution
- Cloud SaaS team, as part of a solution storing data in the cloud
- Endpoint team, as part of a machine compliance management solution

Ever heard the phrase "too many chefs in the kitchen?" It's a similar situation here, but much worse. It's too many chefs (application owners) in too many kitchens (DLP solutions), and each chef is specialized in a different cuisine (technical expertise). It's obvious that such a situation can easily turn into an unappetizing mess.

CONSISTENT IN BEING INCONSISTENT

What doesn't help is the fact that there is no common language standardization for DLP, meaning that the same thing can go by several (in some cases many) different names. For example, a regex-based detection rule can be called a Data Identifier, Sensitive Info Type, DLP Dictionary, Smart ID, etc. This means that even when teams exchange info about their specific DLP implementations, there's a good chance for miscommunication based on an inconsistent utilization of relevant terms.

This inconsistency extends to how the different DLP tools operate, and what exactly they offer. Here are some examples:

Policy Templates

- How many policy templates are provided for the DLP solution?
- What use cases do they cover? (Industry-specific, Region-specific, Legislation-specific, etc.)
- How similar are these across different vendors? *(For example, some HIPAA templates use U.S. Social Security Numbers as a key requirement, others use ICD Codes (International Classification of Diseases). So even if you have a HIPAA policy enabled in 2 different solutions, the results could be drastically different).*

Detection Patterns

- What Regular Expression (regex) language(s) are used in the different DLP solutions?
- How easily can the patterns for one solution be migrated to a different DLP solution?
- Does DLP vendor provide a tool to translate DLP rules from other vendors?

Severities / Priorities

- How is a severity/priority applied to an incident?
- How many severity levels are available?
- Can severities get harmonized across the different DLP tools? (If you have a scale of 1 – 5 for one tool, and (Low – Medium – High) for another, how do you harmonize those?)

Match Counting

- How are direct matches counted?
- How is supporting evidence like keywords counted?
- Is supporting evidence counting by default included, excluded, or optional?

Keyword Proximity

- How is keyword proximity determined to increase relevancy of results (could be by number of characters, words, tokens, or something else)

Policy Tuning

- How easily can policies get tuned?
- Are only high-level changes possible like excluding a recipient, or can you make granular lower-level changes, such as excluding a specific prefix?

Incident Response

- Are different workflow styles supported? (e.g., fan-in vs. fan-out)
- What incident response workflow tools are provided by the vendor:
  - central management console?
  - only via email notifications?
  - via a 3rd party integration like a SOAR?

These are just some examples to illustrate the complexity of any DLP solution, and the nuanced differences that can make an orchestrated effort a heavy lift. Failing to identify all relevant differences between the various tools in your environment can lead to a multitude of issues, such as:

- Discrepant incident response procedures that cannot guarantee that critical violations are identified and reported correctly.
- Poor and fragmented reporting that cannot provide the big picture across all DLP solutions.
- Inconsistent detection capabilities, resulting in detection blind spots.
- Poor policy tuning practices, as too many teams may make countless changes to their policy sets, which over time will increase the disparity among the different DLP policy sets.

IT'S NOT ONLY A DLP PROBLEM

DLP does not work in a vacuum, and often needs to work with other tools to accomplish proper data governance. In some instances, it must get integrated with other tools so that both can work to their full potential. In others, the information detected and generated by DLP is needed for other processes.

Data Tagging/Labeling is a good example of DLP needing to work with other tools in harmony. As a matter of fact, IANS identified data classification and tagging is a key challenge for any DLP implementation, even if you only use a single tool. They accurately called it "an extremely non-trivial task" (IANS Faculty, 2021). The big question is, if it is a challenge to enable proper data tagging with just one DLP solution, how would or could this even work with multiples? How can you ensure that your tagging tool of choice will work with all the different DLP solutions to ensure consistent label application? And what would the end user experience be? The answer will come down to testing, testing, and more testing, as it will be hard to get an answer from just assessing this complex kind of implementation in theory. This will take a substantial amount of time and bind personnel resources that then won't be available for other projects.

And this is for formal projects with proper training and preparation. But what about ad hoc requests, such as for eDiscovery? To be legally sound, the collected data must be as relevant and complete as possible. For example, a file upload might have been caught by several different DLP tools. This means, to get the full picture, you need information from Endpoint DLP, Network DLP, and CASB DLP, to fully understand origination, path, and destination of the upload. Each DLP solution will provide other contextual data points not available from any of the other tools (e.g., machine name, or who the data was shared with), so if you are missing this contextual data, the investigation may be considered incomplete. How can you ensure that all relevant information is compiled and arranged correctly to draw the right and complete picture to satisfy strict legal requirements? There are certainly strong eDiscovery tools out there to help consolidate information from various sources. Nonetheless, the increased complexity of multiple DLP tools won't make this job any easier. And keep in mind this only applies to DLP tools the investigation team is aware of. If Shadow DLP is in full swing, critical information could get missed simply because they didn't know it was there in the first place.

By looking at these two examples alone, it's easy to see that coping with Shadow DLP means that other projects and tool implementations are in jeopardy as well.

BEWARE OF VENDORS OVER-PROMISING

Many vendors offering a Point DLP or Add-on DLP solution promise that "they can do it all"; that their DLP solution is the only one you really need for your use cases. But this overconfidence in one's own capabilities can be misleading at best and devastating at worst. Before you blindly go down the rabbit hole, ask yourself (and the vendor!) some common-sense questions:

- What would be the benefit of making this solution our primary DLP tool?
- Is this only a financial benefit or are there additional features we can't leverage with our currently existing tool stack?
- Is it realistic to deprecate and fully replace our current primary DLP tool in a reasonable amount of time, and without making unnecessary compromises?
- What is the level of support to stand up and operate the new solution, and will hidden additional costs be incurred, such as for Professional Services?
- Who will be responsible for administering and managing the new solution, and who will perform incident triage?
- How will responsibilities change, and are these in line with the corporate reporting hierarchy (e.g., do you want the email team to take over DLP from GRC?)?

As you are trying to answer these questions, keep in mind that examples provided by any vendor trying to convince you of their DLP solution are always "perfect world". They promise little/short policy tuning cycles, high accuracy rates straight out of the box, and a small number of violations that need review/triage. Some even go as far as calling their DLP a "set it and forget it" tool, meaning that once it's in place and some initial policy tuning has been performed that everything else can be automated and nobody will ever have to look at it again. But we know that DLP is anything but that. It will require your ongoing attention, no matter if you used it for six months or ten years. And as we know, once you enable DLP policies in production, you open the floodgates and your first-level triage team will be drowning in DLP violations, even if you start with just a subset of identified data protection use cases.

No later than now you'll learn how good your selected DLP solutions is, based on how efficiently you can handle this influx of incidents. Because no matter how great the detection capabilities may be, if you don't have a good way to assess and triage the resulting events, your chosen DLP solution may proof to be ineffective after all.

But luckily, all is not lost. Now that you can properly identify what Shadow DLP is, it's time to start employing solutions to address it in your organization. We will discuss a phased approach to address Shadow DLP in chapter 3.

CHAPTER 3
OUT OF THE SHADOWS
HOW TO ADDRESS SHADOW DLP

As we can see, having multiple, uncoordinated DLP tools in an IT environment is a very bad idea. But not all is lost. You can clean up the mess by using a phased approach that applies method to the madness.

First, you'll need to gain an understanding of the magnitude of Shadow DLP in your IT environment. Next is to work on your DLP Program, to address the reasons that led to Shadow DLP in the first place. Afterwards, you can start leveraging your tools of choice by integrating them in a smart way to reduce duplication of efforts.

During this whole process, you can augment your path into a cleaner, greener DLP Program future by using a Managed DLP Services Provider to help you make more educated decisions.

PHASE 1: TAKE AN INVENTORY OF THE DLP SOLUTIONS IN YOUR ENVIRONMENT

No matter if you only fear, suspect, or already have some confirmation of Shadow DLP, the first phase of dealing with Shadow DLP is getting the full picture before you can start addressing the problem. You have to identify all the different solutions out there, and if they are in use. This means reaching out to the different business units managing your storage, network ,and endpoint infrastructure and ask them some initial questions:

- Do any tools you use feature DLP capabilities?
- If so, is DLP the tool's primary function, or is it a feature as part of a bigger tool set?
- Is the DLP function currently in use?
- If so, in what capacity (Proof of Concept, feature test, UAT, production, ...)?

This is just an initial set of questions to get the ball rolling. Once you get feedback, there will certainly be more targeted questions required to complete your inventory.

As the DLP Program Owner, you want to figure out all the different pieces to put the DLP puzzle together. To get the full picture, and before you can move on to the next phase, you should have answers to the following questions:

- How many different DLP solutions exist in our environment?
- How many are already in use, and in what capacity?
- What detection capabilities do they have / what use cases do they cover?
- What detection policies were defined?
- Do we have detection overlap between two or more tools?
- Who is responsible for each solution?
- Is incident triage performed, and if so, by whom, and what's the escalation process?

PHASE 2: CREATE A DLP PROGRAM IF YOU DON'T ALREADY HAVE ONE

If the answer to multiple of the initial inventory questions is "yes", there is the implication that a deeper communication and risk management problem exists in your organization. This second phase could be considered optional, but the fact that an IT business unit started using DLP without ever thinking about getting the DLP Program Owner involved is a bad sign.

Long story short, once you know the scope of Shadow DLP and the use cases it tries to cover, it is a great time to build a DLP Program around it. This will make all upcoming phases much easier.

Gartner provides a pretty decent article on how to create a successful DLP program in five steps. The key items we are trying to address in this phase are to scope the program (step 1) and design the initial architecture (step 3) (Goasduff, 2022). Step 1 is already done to a certain degree, based on the DLP solutions you identified during phase 1. Step 3 is realized in both phase 3, employing a primary DLP solution, and phase 4, leveraging APIs and other integrations. Without an overarching DLP Program in place, all coordination and tool harmonization efforts will be much harder, and there's a significantly higher chance that they will fail or be insufficient.

And even if you already have a DLP program in place for your primary DLP solution, you need to address the fact that some business units decided to implement a DLP tool on their own. The second step of Gartner's article is all about raising awareness, which is what was missing from your DLP program from the very beginning. So, before any substantial changes are made to any of the DLP tools covered by the (hopefully revised) plan, you will need to ensure that all relevant stakeholders are fully aware of the DLP program, it's purpose, and how adjustments are about to be made to fix the current situation. A well-defined DLP Program will help to identify relevant data detection and protection use cases, assign priorities, define incident triage workflows, and identify relevant stakeholders.

PHASE 3: EMPLOY A PRIMARY DLP SOLUTION AND ADD ADDITIONAL DLP TOOLS WHERE NEEDED

Once you have a full understanding of the different DLP solutions existing in your environment, and a relevant DLP Program in place to harmonize them, you can move on to the third phase. Here you will determine which tool or tools provide sufficient coverage for all your DLP use cases. Please keep in mind that these might have gotten expanded from the original DLP program plan based on conversations you had with the IT Process Owners during phases 1 and 2.

There are two different approaches you can take:

01.

Each DLP solution covers detection for the tool/use case it got shipped with (the perfect world example we mentioned at the beginning of chapter 2).

02.

You determine a primary DLP solution that will cover the majority of DLP use cases and use other DLP solutions for the use cases the primary DLP solution won't be able to cover sufficiently.

Based on Infolock's experience to date, we have found that the second approach is definitely preferable over the first one. Instead of doing a little bit everywhere, focus the majority of your DLP detection in one central solution. This is not necessarily the solution with the most detection capabilities, but the one that can satisfy your most important DLP use cases the best.

Once you have identified the primary DLP solution, the next step is to determine what potential detection gaps exist that this platform cannot close, but other existing tools could. Use these tools to complement your primary solution for these use cases. This way you establish the best detection possible without several DLP tools getting in each other's way or generating duplicate incidents for the same file transfer. If additional coverage is necessary, determine who will be responsible for managing these solutions, and who will coordinate all detection and incident triage efforts to enable a unified DLP experience.

Once the determination has been made on which tool or tools to use, the legwork starts. This includes:

- Informing all DLP stakeholders on the final decision on which tools to keep and which to disable, as well as a transition plan to move the relevant detection methods into the right tools.
- Determining the final set of detection policies to use, as well as a means to ensure consistent detection across all approved DLP solutions.
- Assigning responsibilities regarding DLP program management, tool administration, and incident response workflows.
- Establishing a reporting methodology that incorporates data from all approved platforms in a meaningful way.

Example Use Case

When done right, adding a second DLP solution can be very successful. For example, we manage a customer with a longstanding DLP environment, using a General DLP solution for all use cases. However, a few years ago they migrated to Microsoft 365 and switched their office communication tool from a different solution to Microsoft Teams. They wanted to gain visibility for Teams chats, which their existing DLP tool could not sufficiently provide at the time. Reviewing their licensing structure, we determined that Microsoft's own DLP solution would be able to provide coverage for this new use case. As a result, we expanded their DLP program to include this new tool, replicated all relevant policies from the primary DLP solution, and found ways to enable efficient incident triage and reporting. While not a Shadow DLP example, it's a great way to show how two DLP tools can run side-by-side to provide expanded coverage.

PHASE 4: UTILIZE APIs AND OTHER INTEGRATIONS TO ORCHESTRATE DLP

One problem you'll face running multiple DLP solutions side-by-side is that relevant information is spread all over the place. This makes consistent incident triage, reporting, and overall DLP oversight a challenge. Following are two potential options to tackle these problems.

Option 1 Direct Integrations

One way is to leverage APIs and other protocol-based integrations to connect the DLP solutions directly. The goal is to use one DLP solution (preferably your primary) as a single pane of glass as much as possible.

Example Use Case

We have a customer that used a General DLP Solution for all their identified use cases. About a year ago they replaced their on-prem proxy infrastructure with a cloud-based proxy solution, which meant that their on-prem DLP detectors were no longer able to inspect web traffic. One solution could have been to use the endpoint agent to inspect web traffic, but this would have limited visibility more than what we were comfortable with.

So, we went down a different route. As the cloud proxy had its own DLP engine, we decided to replicate the detection policies from their General solution in the new solution. Afterwards, we created a Secure ICAP tunnel to send the violations detected by the proxy's Add-on DLP to their General DLP solution. This way, we omitted having to do incident response in two different solutions. At the same time, we were able to continue using our established reports in the General DLP Solution to show full DLP coverage, including the web traffic violations from the cloud proxy.

Unfortunately, not all vendors support this kind of connection to or from their competitors' DLP solutions, or not to a satisfying degree. Nonetheless, it's worth trying, as it's best to use one DLP solution for as much incident response as possible and to get the most comprehensive reports.

Option 2 External solutions

Alternatively, you can leverage a Security Information and Event Management (SIEM) or Security Orchestration, Automation and Response (SOAR) solution:

- A SIEM is a good idea if you need a single source of truth for incident reporting. You can correlate events across different tools (e.g., an upload detected on the endpoint can be correlated to a file transfer detected by your proxy or CASB solution). It won't necessarily help much with incident response but can help to enable consistent reporting across the different DLP tools in place.

- A SOAR tool is an interesting option if you want to centralize and streamline your incident response procedures. Connecting and configuring the different DLP solutions to leverage SOAR scripts, etc., might be tricky and will likely require some Professional Services work. But once established, your incident triage team can work with a single pane of glass to review incidents from various different DLP tools.

Depending on the number of DLP solutions in scope, there may be some substantial legwork that needs to be done in the SIEM/SOAR tool to normalize the data. This will be important so that all DLP data follows a consistent attribute mapping scheme. Without this, you are merely moving key Shadow DLP issues to a different portal.

By following the four phases we just discussed, you'll be able to determine the scope of Shadow DLP, establish a proper DLP Program Plan, identify the right tools for the job, and implement and orchestrate them in a way that makes sense. But there's one more thing you can do, if you want to increase chances to pull this off.

LEVERAGE A MANAGED DLP SERVICES PROVIDER TO HELP YOU ADMINISTER THE DIFFERENT TOOLS

Managing a single DLP solution can already be a daunting task. It is shipped with a lot of modules which need continuous fine-tuning to operate efficiently. Additionally, resulting incidents must be reviewed within a reasonable time frame to ensure that broken business processes are addressed without unnecessary delay. The level of effort increases exponentially with each new DLP solution you add to the stack.

Your primary incident response team is often already overwhelmed with the alerts reported from all your different security tools (phishing, malware, etc.) so adding more DLP violations to the queue will increase the risk of incident fatigue and missing critical violations in the noise (think of needle in the haystack). While this fan-out approach is easier to manage, it puts a lot of stress on a single team.

Alternatively, you can try to divide-and-conquer by having the IT Process Owners run the Add-On DLP for the tools they are responsible for, called the fan-in approach. Unfortunately, this is usually not a sufficient solution either. First, you are merely outsourcing some of the issues your incident response team faces to different internal teams. Next, these teams usually don't have the relevant expertise to perform incident triage.

This means their workflows might be insufficient or incomplete, and there's a substantial risk that relevant DLP violations are misclassified or not investigated in an appropriate amount of time. Lastly, there is the potential for limited oversight from the DLP Program Owner, which increases the chances that DLP won't be operated in a way compliant with the DLP Program.

The solution to all these problems could be a MDLP Services Provider. A good MDLP Services Provider can help you during each phase of the Shadow DLP identification and mitigation process. They have:

- the required expertise to manage different DLP solutions in a consistent way;
- the technical acumen to translate a core set of detection rules to work with different DLP solutions, which will ensure consistent detection methodology;
- insight into various different customer environments so they can identify and address new DLP trends faster;
- enough human resources to deal with the larger number of violations that could result from running multiple DLP solutions and provide the necessary coverage your organization may require; and
- Service Level Agreements, that ensure relevant DLP-related tasks are completed within specified timeframes, as in accordance with your organization's DLP Program Plan and expectations.

While an MDLP Services Provider comes with a certain price tag, it can save you money in the long run as it frees up internal resources from work-intensive, repetitive tasks so they can better apply their subject-matter expertise to other, more important topics.

A FINAL THOUGHT

As you can see, there are several things that can be done to identify and address Shadow DLP. The most important thing is to determine if you experience this phenomenon, and to what degree. At the end of the day, you can't fix what you don't know about. What you do from there is certainly up to you. This will depend on the size of the problem, the number of stakeholders involved, and ultimately the willingness to change things for the better. But we know that there is a way to fix the problem. It may not be easy, but what in life ever is?

As a final piece of advice, don't fall for the urge to use all your DLP tools just because you have them. In many cases, DLP was not the reason you got the tools. Less is certainly more here, as long as you are able to cover all your identified data loss use cases.

If you encounter Shadow DLP, employ the phases outlined in this chapter to bring order to the chaos and restore your DLP program to a better and more effective state. And, if you fear you can't pull it off yourself, Infolock is here to help!

Schedule a zero-commitment agnostic review of your data protection needs and we'll send you a customized overview of your DLP options and technical recommendations.

Schedule your call HERE

or visit infolock.com/review

GLOSSARY

DLP Incident - Also often called a DLP Alert, this is the result of a policy **violation** detected by a configured DLP rule.

DLP Program - The organization's overarching strategy to use data loss prevention to mitigate risk by identifying business processes that handle sensitive information. Additionally, it outlines measures to safeguard this information effectively, and assigns responsibility to manage all facets of the DLP solution.

DLP Program Owner - An individual or business unit, usually part of the Governance or Information Security umbrella, responsible for the overall DLP Program within that organization. They are also often responsible for the operation of the DLP tools.

Fan-in Incident Response - Initial triage is performed by individual business units all across the organization, and feedback is sent back to central department for analysis and reporting.

Fan-Out Incident Response - Initial triage is performed by one business unit. Depending on verdict, additional business units are involved as need be.

General Data Protection Regulation (GDPR) - The European Union's data privacy and security law, protecting EU citizens' personal data all over the world.

IT Process Owner - An individual or business unit, usually under the IT department umbrella, responsible for a specific IT-related business process. This could be things like networking, email, storage, or server administration, just to name a few.

Primary DLP - The solution selected by the DLP Program Owner to cover all relevant DLP-related use cases. This commonly is a General DLP solution, but it may be just a Point or Add-on DLP solution based on number and scope of use cases.

Security Information and Event Management (SIEM) - A central tool that ingests logs and events from a multitude of security tools to provide real-time analysis and reporting capabilities.

Security Orchestration, Automation and Response (SOAR) - A central tool designed to accept events from a multitude of security solutions, to improve efficiency and effectiveness of incident response workflows.

REFERENCES

Alspach, K. (2023, June 29). The 10 Biggest Data Breaches of 2023 (So Far).

Retrieved from CRN.com:
https://www.crn.com/news/security/the-10-biggest-data-breaches-of-2023-so-far

Brunell, T. (2023, December 20). 2023: A Year of Record-Breaking Data Breaches.

Retrieved from IdentityIQ:
https://www.identityiq.com/data-breaches/2023-a-year-of-record-breaking-data-breaches/

Ford, N. (2024, January 05). List of Data Breaches and Cyber Attacks in 2023 – 8,214,886,660 records breached.

Retrieved from IT Governance Blog:
https://www.itgovernance.co.uk/blog/list-of-data-breaches-and-cyber-attacks-in-2023

Goasduff, L. (2022, Febuary 02). Build a Successful Data Loss Prevention Program in 5 Steps.

Retrieved from Gartner IT Insights:
https://www.gartner.com/en/articles/build-a-successful-data-loss-prevention-program-in-5-steps

IANS Faculty. (2021, April 22). Data Loss Prevention: Challenges and Strategies.

Retrieved from IANS Research:
https://beta.iansresearch.com/resources/all-blogs/post/security-blog/2021/05/03/data-loss-prevention-challenges-and-strategies

IBM. (2023). What is shadow IT?

Retrieved from
https://www.ibm.com/topics/shadow-it