Data labeling for all Businesses

Simple Purview Implementations

Abstract

Data labeling and classification is a critical security control for any business, especially in the Copilot age. In this guide, we will walk through a few simple deployments that can be used for about any SMB customer.

Disclaimer: This guide is provided AS IS for educational purposes only, without any warranty, whether express or implied. Every person and organization is responsible for their own cybersecurity hygiene. You are using this guide at your own risk.

Dom Kirby
domkirby.com

TLP:CLEAR - Disclosure is not limited.

About Purview Information Protection

Microsoft Purview Information Protection **(PIP)**, formerly known as **Microsoft Information Protection (MIP)**, is a suite of tools and capabilities designed to help organizations **discover, classify, protect, and govern sensitive information** across their digital landscape. Whether data lives within your organization or travels beyond its boundaries, Purview Information Protection provides the necessary tools to safeguard it.

Core to the functionality of Purview Information Protection is the concept of **sensitivity labels**. This functionality will help you meet the requirements of **CIS Safeguard 3.7: Establish and Maintain a Data Classification Scheme**. The challenge with 3.7 is that many small businesses may not seem mature enough to implement a data classification scheme. However, I think one or a combination of these implementations will help you get more customers over the hump.

Note:** For brevity's sake, this guide will not walk through the configuration of each individual label and policy. I will provide the settings for each label, but you should reference the product documentation for guidance on creating them. **Always practice in a lab first!

What is Data Classification?

Most of the time I introduce this concept, peoples' mind goes straight to "classified" information in the government context. Frankly, the concept is similar. The Government broadly follows the Bell LaPadula model (no write down), which is hard to apply to the corporate world. Nonetheless, the general idea stays the same.

Data Classification means we understand our data types (and their sensitivity) and label it appropriately. This helps us:

- ✓ Understand our data.
- ✓ Apply proper protection measures to sensitive data.
- ✓ Clearly mark our data so that others know what type of data they are working with.
- ✓ Restrict **who** can access our data to those who **need** to access it.

The beauty of PIP is that it works out of band. I can send you a file with a data classification, and you cannot open it unless you are on the ACL (which can be applied in several ways we will cover).

General Label Architecture

Before we dive into building sensitivity labels, let us talk about what they are capable of. Sensitivity labels can be used in several ways from simple marking to robust encryption and access control. Here are the key features available with each label:

- **The Basics: All your labels will have a user-faced name and description.** Planning these is essential, as you will need to train your users and make a naming convention that is clear and easy to understand.
- **Scope: When making a label, we must choose the scope to which it can be applied.**

  - **Items:** This will be the core of our focus for this guide. **Items** allow you to apply labels to *Files and/or Emails*. These are the most common types of data we need to protect. For simplicity's sake, I often recommend (for SMB) that all labels apply to Files and Emails.
  - **Groups & sites:** This needs a little setup work (we will get into that later) but it is a powerful tool. This functionality allows you to label an entire Microsoft 365 Group or SharePoint site, automatically labeling the content there.
  - **Schematized data assets (preview)**: This is a new functionality I will not be addressing in this guide. It will allow for the protection/labeling of SQL, Azure SQL, Azure Synapse, Azure Cosmos, AWS RDS, and more.
- **Application Options:** What would we like the labels to do once applied?
  - **Items:**
    - Apply or remove encryption: Will apply encryption and access control to a document or email.
    - Apply content marking: Will add headers, footers, and/or watermarks to documents and emails (although watermarks do not apply to emails)
  - **Groups and Sites:**
    - Privacy and external user access: Controls the levels and methods of access available to internal and external users accessing content within a labeled group or site.
    - External sharing and conditional access: Allows for limiting external sharing and applying Conditional Access Policies to labeled groups or sites.

A Note on Auto-Labeling

Auto labeling is a great destination for this journey. Depending on your licensing, you can apply granular rules that will automatically apply sensitivity labels. However, the point of this guide is to help you get your customers to take baby steps. With that in mind, we are not going to cover automatic labeling in this guide.

Tenant Preparation

Before we dive into creating labels and label policies, we need to prepare our tenant. After ensuring you have the correct licensing for implementing Purview Information Protection, consider the following preparations.

Enable Groups and Sites Protection

Before you can apply labels to Modern Groups and/or SharePoint sites, you'll need to complete some preparatory steps. This will require a little bit of PowerShell, so ensure you're using a Windows machine for this. Follow the steps at https://learn.microsoft.com/en-us/purview/sensitivity-labels-teams-groups-sites to configure this capability.

Get your Groups Right

The simplest implementation we'll look at doesn't have much reliance on Entra ID groups. However, making sure you have groups based on information boundaries is an important pre-requisite for more complex label configurations. You'll want to make sure you have Entra ID groups based on the customer's working groups (departments, tiger teams, etc.).

Deploy Microsoft 365 Applications

Volume license or retail copies of the Microsoft 365 (Office) suite are prone to challenges when working with PIP labeled content, especially when it comes to using the latest features (such as co-authoring of protected content). As such, I **strongly** recommend you deploy the Microsoft 365 subscription (click-to-run) applications within the customer's environment instead of retail or volume copies of Office.

Optional: Enable Desktop Co-Authoring for Protected Files

Microsoft is now rolling out the capability to co-author protected files using Office desktop apps. This was a longtime shortcoming of working with protected content (although you've been able to do so on the web for some time). **However, this comes with a list of caveats.** Please review these docs: https://learn.microsoft.com/en-us/purview/sensitivity-labels-coauthoring and make a careful decision before enabling this feature. It makes key changes to your metadata structure and will interrupt many third-party DLP products.

Implementation Scenarios for Small Business

In this guide, we will cover four implementation scenarios of increasing complexity. We will start with the simplest of implementations (for the least data mature customers) and move into a more organizational chart aligned model with multiple levels of confidentiality for mature customers with different classes of sensitive data.

Assumption: Protection Scope

Remember we're going for simple here. I did not specify a label scope in the suggested settings below. I recommend scoping all the labels to *Files and Email*, and *Groups and Sites*.

Scenario 1: Let's Just Mark What's Confidential

The simplest deployment of PIP involves just a few labels. It is best suited for customers who aren't yet mature enough to have defined data access control lists or don't yet have defined "departments." It's important to note that this implementation is *better than nothing* and will help with outside risks. However, it does basically nothing to protect against *insider risk*, so it should be considered a "steppingstone" in the right direction.

Labels

- **Personal**: This label applies no protections; it is a way to classify something that is personal. People are naturally going to work with "personal" content on their work PC, giving them an outlet to do so and clearly mark that data is helpful. *Note: This doesn't mean you're getting away from the fact that there's no privacy on company storage. It simply is a means to mark content that isn't necessarily work-related.*

- **Public:** This label also doesn’t apply any protections. This label is meant to classify data that *is* work related but *is not* confidential. Common examples of this may include general work communications, marketing materials, etc. This will also include any type of communication sent outside the organization.
- **Company Confidential:** This is our first protective label. This label will apply encryption but will allow access to all internal users of the company. It should be used to classify information sensitive and thus only used by internal team members.

Note: *Any time I reference “Company” such as “Company Confidential,” you can interpret it as is or replace Company with the name of the company (which can help provide clarity). An example would be “Contoso Confidential.”*

Label Settings

Label	User Description	Encryption	Content Marking
Personal *Priority 0*	Non-business data, for personal use only.	*No Encryption*	*No Content Marking*
Public *Priority 1*	Business data that is not sensitive in nature.	*No Encryption*	*No Content Marking*
Company Confidential *Priority 2 - Highest*	Confidential data that requires protection, which allows all users full permissions. Data owners can track and revoke content.	**Yes** • Assign Permissions Now • User access to content expires *Never* • Allow offline access *always* • Permissions: See Below	*Optional: Consider adding a footer that states “Confidential” or something to that effect.*

For Company Confidential permissions: The permissions for Company Confidential are intentionally broad in scope. That is because this is for a less mature SMB that doesn’t have a lot of data boundaries in place. With that in mind, when assigning permissions:

- Choose the “Add all users and groups in your organization” option

  + Add all users and groups in your organization

  - As a result, **any** authenticated user to your tenant will be able to access this data.
- Typically, you’ll want to choose “Co-Author” permissions in this label so that all users can work with the data.

At the end, your permissions table should look like this:

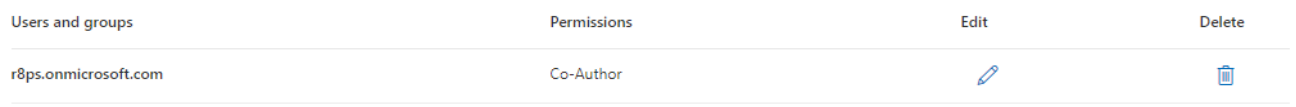

Users and groups	Permissions	Edit	Delete
r8ps.onmicrosoft.com	Co-Author	✎	🗑

Of course, the onmicrosoft.com domain should be your tenant's onmicrosoft.com domain (not my lab's).

> ***A note on priority:*** *The priority of your labels starts at zero and go up from there. The higher priority the label, the more important Purview considers that label. Be careful with this. If "Public" is a higher priority label than "Confidential," you end up in a scenario where sensitive content is downgraded because of priorities. Labels should increase in priority based on the sensitivity of the data they are meant to be labeling.*

Label Policy

Labels are published through *Label Policies*. One or more label policies will affect what labels a user has available to them, and which ones may be applied by default. In this very simple method, we only need one label policy.

Policy Settings:

- **Labels to publish:** All the labels we just made.
- **Admin units:** Skip (this is an E5 feature).
- **Users and groups:** All users and groups (unless you want to roll this out to a pilot group)
- ***Recommended* settings:**
    - Require users to apply a label to their emails and documents: selected. This will force users to label their content *somehow* (even if most of it is public for now).
- ***Recommended* default labels:**
    - **Documents:** Public
    - **Emails:** Same as document; Require users to apply a label to their emails; Email inherits highest priority label from attachments (this will automatically label an email Company Confidential if a Company Confidential document is attached).
        - **Optional:** Recommend users apply the attachment's label instead of automatically applying it. This causes Outlook to give a "heads up" that their chosen email classification is lower than its attachment(s) instead of just automatically relabeling the email.
    - **Meetings:** None (unless you have Teams Premium)
    - **Fabric and Power BI:** None; Out of scope of this guide.
- **Name:** Something descriptive like "Contoso Global Labeling Policy."

Summary

As you can see, this method is extremely simple. The entire point is to just give the customer *some way* to mark things that are for internal consumption only. It's an effective way to start the labeling conversation and expose users to the concept. Importantly, it can build a foundation for growth.

Scenario 2: Contractors

Many businesses use independent contractors to help them scale their talent. In most scenarios, there are legal requirements around how contract staff are managed versus full-time employees. Additionally, many businesses wish to control what information an independent contractor is working with (to limit it to their contract's scope).

This scenario takes a small step up in complexity by introducing **label groups** and **sublabels**. It tackles the "contactor challenge" by grouping employees in a separate group from contractors, which will be encompassed by "All Users."

Pre-Requisite

Before we set this up, we need to understand or build the mechanism for separating contractors from FTEs. There are a couple of ways to do this:

- **Employees Group** (Preferred): My usual approach is to have a specific group that is for employees, and one for contractors. It is important to remember that PIP labels are *inclusionary* and not *exclusionary*. As such, we need to *include* our FTEs in access rights.
    - **Dynamic Group:** I usually like to mark contractors as such in the user directory. This can be done by using a username convention (like Microsoft's "v-" method) or by using the "Employee type" field available in Entra ID. In the example below, Grady is a *contractor* at Contoso:

      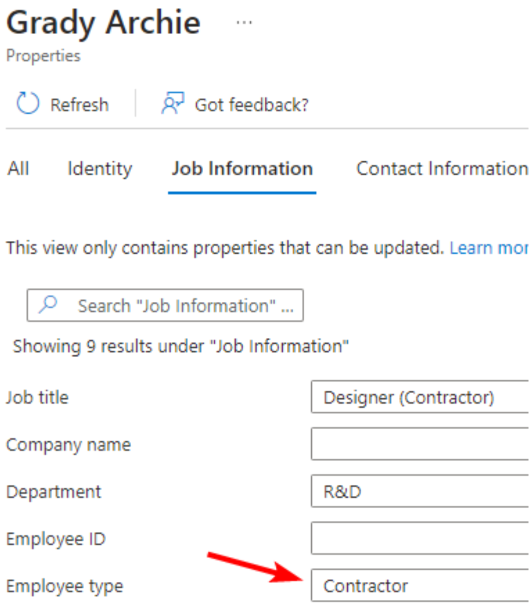

      Grady Archie

      Properties

      This view only contains properties that can be updated. Learn mor

      Showing 9 results under "Job Information"

Job title	Designer (Contractor)
Company name	
Department	R&D
Employee ID	
Employee type	Contractor

      Contoso has opted to identify contractors by specifying the employee type as Contractor. This means that we can build a dynamic group of employee type equals employee, full time, or whatever our designation is for FTEs (we could also do a "not equals contractor" lookup)

Labels

This scenario expands just a little bit on top of scenario 1 by introducing a label group and sublabels for all users and employees.

- **Personal**: This label applies no protections; it is a way to classify something that is personal. People are naturally going to work with “personal” content on their work PC, giving them an outlet to do so and clearly mark that data is helpful. *Note: This doesn’t mean you’re getting away from the fact that there’s no privacy on company storage. It simply is a means to mark content that isn’t necessarily work-related.*
- **Public:** This label also doesn’t apply any protections. This label is meant to classify data that *is* work related but *is not* confidential. Common examples of this may include general work communications, marketing materials, etc. This will also include any type of communication sent outside the organization.
- **Company Confidential:** This label will not apply any protections, but its sublabels will.
    - **All Users**: This will be the same as “Company Confidential” in scenario 1 but will show as “Company Confidential\All Users.”
    - **Employees Only**: This will be the label that limits access to only those who are members of our Employees group and will show as “Company Confidential\Employees Only.”

Label Settings

Label	User Description	Encryption	Content Marking
Personal *Priority 0*	Non-business data, for personal use only.	*No Encryption*	*No Content Marking*
Public *Priority 1*	Business data that is not sensitive in nature.	*No Encryption*	*No Content Marking*
Company Confidential *Priority 2*	Confidential data that requires protection, select a sublabel for the appropriate audience.	*No Encryption (applied by the sublabels)*	*No Content Marking*
Company Confidential\All Users *Priority 3* *Note that this should be created a sublabel under Company Confidential called “All Users.” This will achieve the above display name.*	Confidential data that requires protection, which allows all users full permissions. Data owners can track and revoke content.	*Same as “Company Confidential” in Scenario 1*	*Consider markings in the footer of content.*
Company Confidential\Employees Only *Priority 4 - Highest* *Note that this should be created a sublabel under Company Confidential called “Employees Only.” This will achieve the above display name.*	Confidential data that requires protection, which restricts access to Company employees only (contractors will not be able to view this content)	**Yes** • Assign Permissions Now • User access to content expires *Never* • Allow offline access *always* • Permissions: See Below	*Consider markings in the footer of content that specify the content is for employees only (such as “Company Confidential: Employees Only”)*

Company Confidential\Employees Only Permissions: Configuring this label will be very similar to Company Confidential in scenario 1. However, instead of adding all users, we will add our Employees group we (hopefully) have made.

Label Policy

This label policy can be identical to our scenario 1 policy. You'll just want to make sure you include all **five** of our labels we created above (the parent Company Confidential label counts as one).

Policy Settings:

- **Labels to publish:** All the labels we just made.
- **Admin units:** Skip (this is an E5 feature).
- **Users and groups:** All users and groups (unless you want to roll this out to a pilot group)
- ***Recommended* settings:**
    - Require users to apply a label to their emails and documents: selected. This will force users to label their content *somehow* (even if most of it is public for now).
- ***Recommended* default labels:**
    - **Documents:** Public
    - **Emails:** Same as document; Require users to apply a label to their emails; Email inherits highest priority label from attachments (this will automatically label an email Company Confidential if a Company Confidential document is attached).
        - **Optional:** Recommend users apply the attachment's label instead of automatically applying it. This causes Outlook to give a "heads up" that their chosen email classification is lower than its attachment(s) instead of just automatically relabeling the email.
    - **Meetings:** None (unless you have Teams Premium)
    - **Fabric and Power BI:** None; Out of scope of this guide.
- **Name:** Something descriptive like "Contoso Global Labeling Policy."

Summary

The aim of this scenario is to maintain simplicity from the first scenario, but also account for a common scenario of user personas. You could easily apply "contractors" to something different if needed, that was just the most common persona I saw in practice.

Scenario 3 (Medium Complexity): Departments

All right let's turn up the heat just a little bit. Let's say that our customer has defined departments within their organization chart, and those departments each have data that is sensitive to them and the piece of the business they run. This scenario loosely applies the concept of **mandatory access control**[1] where administrators set access boundaries.

[1] Mandatory access control stipulates that one cannot change the classification of an object. The policies as written here do allow that, which is why this is a "loose" application of MAC.

Pre-Requisites

This is a more complicated implementation of labeling and has more pre-requisites as a result. The first thing we must be able to do is identify people's departments **using security groups**. This example uses departments, but this could be applied to many components of an org chart.

Recommendation: Fill out the Directory

If you have a mature customer with departments, consider filling out the user directory to reflect titles, departments, managers, etc. This adds a **lot** of value to Entra ID (such as by filling out the org chart in Teams) but also gives us a lot of data from which we can make *dynamic groups*.

Recommendation: Use Dynamic Groups

For departments that you intend to use for PIP, I recommend that you create **dynamically assigned groups** based on the **department field in Entra ID user objects**. This will automatically add and remove users as they are brought on or change departments. The group type is up to you; If you are using all the collaboration features (like Team Sites for departments), consider using **Microsoft 365 Groups**. Both M365 and Security Groups will be available in PIP labels for permissions.

Labels and Settings

For simplicity's sake, we'll simply build this on top of scenario 1's labels. However, if you've already implemented scenario 2, you can simply build upon that.

Label	User Description	Encryption	Content Marking
Personal *Priority 0*	Non-business data, for personal use only.	*No Encryption*	*No Content Marking*
Public *Priority 1*	Business data that is not sensitive in nature.	*No Encryption*	*No Content Marking*
Company Confidential *Priority 2*	Confidential data that requires protection, which allows all users full permissions. Data owners can track and revoke content.	**Yes,** same as scenario 1 Company Confidential.	*Optional: Consider adding a footer that states "Confidential" or something to that effect.*
Department Confidential *Priority 3*	Data that is confidential to a specific department. Select a sublabel in this group for the appropriate department.	**No**, this is a label group.	*N/A*
Department Confidential\Finance *Priority 4 - Highest* *Note that this should be created a sublabel under Department Confidential called "Finance" or the*	Data that is confidential to the Contoso Finance department.	**Yes** See below	*Recommendation: Apply footer markings that indicate confidentiality and the department. Example:*

appropriate department name. This will achieve the above display name.	*Contoso Finance Confidential*
Repeat the above sublabel for each department you would like in scope.	

The first three labels mimic scenario 1, so we won't cover those here. For the **department sublabels**, we will need to use our departmental groups. Mimic the settings from "Company Confidential\Employees Only" above but select your departmental groups *instead of* your employees' group.

Label Policies

You'll notice that the word "policies" in this section is suddenly plural. I've upped the complexity level on you for scenario 3! If we were to create a single label policy that everyone gets all labels, people are going to make mistakes. Imagine being in Sales and accidentally marking something for Finance. Now you can't access your own document! As such, we will create a **base policy** and then create **department** policies that stack on top.

> ***Note:*** *This is where label priorities become particularly important. Labels will maintain their priority even when stacked for users via multiple policies.*

Policy 1: Base Policy

Our base policy will look a lot like our other policies.

Policy Settings:

- **Labels to publish (DIFFERENT FROM OTHER SCENARIOS):** Personal, Public, Company Confidential (including the sublabels of Company Confidential if you are building on scenario 2).
- **Admin units:** Skip (this is an E5 feature).
- **Users and groups:** All users and groups (unless you want to roll this out to a pilot group)
- ***Recommended* settings:**
  - Require users to apply a label to their emails and documents: selected. This will force users to label their content *somehow* (even if most of it is public for now).
- ***Recommended* default labels:**
  - **Documents:** Public
  - **Emails:** Same as document; Require users to apply a label to their emails; Email inherits highest priority label from attachments (this will automatically label an email Company Confidential if a Company Confidential document is attached).
    - **Optional:** Recommend users apply the attachment's label instead of automatically applying it. This causes Outlook to give a "heads up" that their chosen email classification is lower than its attachment(s) instead of just automatically relabeling the email.
  - **Meetings:** None (unless you have Teams Premium)
  - **Fabric and Power BI:** None; Out of scope of this guide.

- **Name:** Something descriptive like "Contoso Global Labeling Policy."

Note: You may wish to make some different decisions on default labeling when applying scenario 3.

Policies 2 and Beyond: Departmental Policies

Once you have your base policy, you will need to create *departmental* policies that publish the correct departmental labels to the correct department members. **Repeat** this policy for each department you have labels for.

Policy Settings:

- **Labels to publish (DIFFERENT FROM OTHER SCENARIOS):** Department Confidential and the sublabels that apply to that department. You **must** include **both** the top label (Department Confidential) and all appropriate sublabels.
- **Admin units:** Skip (this is an E5 feature).
- **Users and groups (DIFFERENT FROM OTHER SCENARIOS):** The proper departmental group for the department you are publishing to.
- ***Recommended* settings (DIFFERENT FROM OTHER SCENARIOS):** Inherited from base policy.
- ***Recommended* default labels (DIFFERENT FROM OTHER SCENARIOS):** Inherited from base policy.
- **Name:** Something descriptive like "Contoso [Department] Labeling Policy."

Other Considerations

Tiger Teams

Dealing with tiger teams or cross-departmental teams can be a challenge. In many orgs, we would just default to "Company Confidential" to enable this cross-collaboration. You can also consider scenario 4 which offers customized permission ideas.

Executive Access

Depending on your customers' structure, the department idea gets difficult at the executive level. For example, the Chiefs might all be in a "G&A" department. To take this into account, you'll need to consider what data executives need access to and plan accordingly. The approach may be to create an "Executives" group and add it to the permissions of your departmental labels.

Summary

In scenario 3, we've really climbed the maturity ladder. This scenario simply will not work if the customer does not have an operationally mature org chart. Nonetheless, I wanted to include it as a potential growth target for your customers. For a complex company, this is still a rather simple implementation of segmentation of information by the department.

Scenario 4 (High Complexity): Pull out all the stops!

Okay, we're getting hypothetical here. *I'm not going to get into detailed configuration in scenario 4*. However, up until this point, we've only covered data sensitivity in terms of *who should access data* and the fact that it *is confidential*. Scenario 4 expands the concept of confidentiality to **levels** of confidentiality.

Essentially, different data might have a different impact on the company if we lose control of it. Scenario 4 addresses that.

In addition to **confidentiality levels** (similar to the Bell LaPadula concept of Secret, TS, etc.; minus the "no write down" concept of the state machine), scenario 4 brings in the idea of **customized access lists** (discretionary access control) which empowers users to apply granular access when it is appropriate (such as for a tiger team document, or even sensitive information shared externally).

> ***Note:*** *This is a very enterprise style of implementation with added management overhead. In fact, it is essentially based on Microsoft's enterprise labeling examples.*

Enterprise Labeling Examples

As a reminder, I am not going to walk through this implementation, but I do want to share it as a thought exercise you can carry forward if you want. If you'd like to play with this enterprise labeling scenario, stand up a Microsoft 365 Developer tenant. This scenario relies on the concept of *everything is labeled*.

Label	Description	Protections	Markings
Personal		Same as other scenarios	
Public		Same as other scenarios	
General *This is a label group; a sublabel must be chosen.*	Business data that is not intended for public consumption. However, this can be shared with external partners, as required. Examples include a company internal telephone directory, organizational charts, internal standards, and most internal communication.	None	None
General\Anyone	Organization data that isn't intended for public consumption but can be shared with external partners if appropriate. Examples include customer conversations that don't include sensitive info or released marketing materials.	None	None

General\All Employees (unrestricted)	Organization data that isn't intended for public consumption. If you need to share this content with external partners, make sure it's appropriate with data owners and relabel content as General/Anyone (unrestricted). Examples include a company internal telephone directory, organizational charts, internal standards, and most internal communication.	None	None

Label	Description	Protections	Markings
Confidential *This is a label group; a sublabel must be chosen.*	Sensitive business data that could cause damage to the business if shared with unauthorized people. Examples include contracts, security reports, forecast summaries, and sales account data.	None	None
Confidential\Anyone (unrestricted)	Data that does not require protection. Use this option with care and with appropriate business justification.	None	Footer: Classified as Confidential
Confidential\All Employees	Confidential data that requires protection, which allows all employees full permissions. Data owners can track and revoke content.	Encrypted; all tenant users can access	Footer: Classified as Confidential

Confidential\Trusted People	Confidential data for internal/external sharing that can be reshared by trusted recipients.	Encrypted; *data* owner *selects users **or** email recipients have permission*	Footer: Classified as Confidential
Highly Confidential *This is a label group; a sublabel must be chosen.*	Very sensitive business data that would cause damage to the business if it was shared with unauthorized people. Examples include employee and customer information, passwords, source code, and pre-announced financial reports.	None	None
Highly Confidential\All Employees	Highly confidential data that allows all employees view, edit, and reply permissions to this content. Data owners can track and revoke content.	Encrypted; all tenant users can access	Footer: Classified as Highly Confidential
Highly Confidential\Specified People	Highly Confidential data that requires protection and that can only be viewed by specified people, with limited rights.	Encrypted; *data owner selects users **or** email recipients have permission*	Encrypted; recipient selects viewers or email recipients have permission

As you can see, this model is built on the idea of **confidentiality level** and uses **discretionary access control** to empower data owners to control who has access. You could blend this scenario with scenario 3, allowing for both **mandatory access control and discretionary access control**. For example, "Highly Confidential\HR" could be a label used for employee data and would rely on **mandatory** access control to ensure only HR has access. Conversely, HR could use the **discretionary** label of "Highly Confidential\Specified People" if they need to share a file with a manager or employee outside of HR for a limited purpose.

***Discretionary** access control can be applied internally and externally. For example, your "Confidential\Trusted People" label could apply a "Do Not Forward" setting to an external Outlook email, thus bringing the external parties securely into the mix.*

Confidentiality Levels

Confidentiality levels will mean different things to different people. However, the general rule of thumb dictates that the **level** of confidentiality is weighted by the **impact** of that data leaking. This can be done with a simple qualitative calculation:

Impact Level	Impact	Confidentiality Level
0	No Impact (the information is public)	**Public**
1	Limited to no impact, but the information is nonpublic	**General**
2	Some impact: may release news early or something to that effect (think of a press release draft)	**Confidential**
3	Moderate impact could have implications on business operations (contracts, compliance reports, etc.)	
4	Significant impact; could create competitive challenges or privacy violations etc.	**Highly Confidential**
5	Damning; Legal Troubles; Our Competitors will beat us if we leak this; the board will be pissed... etc. Financial information, credentials, source code, etc.	

These are illustrative examples of impact levels. Each organization should decide their own risk tolerance and impact levels.

Corporate Policy Implications

I know we just got super nerdy with the available implementations. But before you apply all the nerdy bits, we must start with the business bits. **Corporate policy should dictate how Purview Information Protection – or any other data classification product – is implemented**. Applying it the other way around will create confusion when you look to implement the rest of an information security program.

As such, you should **first** work with your client to create a Data Classification Policy, and then implement Purview Information Protection to align with that policy. This will also make it easier to create training aids, as they can reference the client's Data Classification Policy. Data classification can also be covered in a broader Data Management Policy.

Training Users

The concept of labeling data is likely to be new to your customers' teams. When rolling this out, you'll need to take user training into consideration. In general, here are the phases I would use to roll out PIP / information labeling:

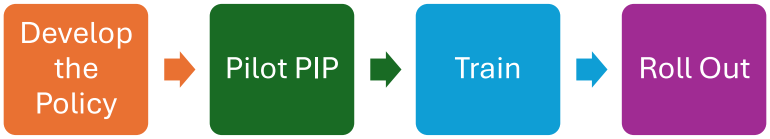

Piloting and Power Users

Like any good technology rollout, your PIP implementation project can benefit from a few good power users. Pick some savvy users from around the company and train them on the new policy and technology first. Work with company leadership to empower them to coach others and call them out when something is perhaps mislabeled.

Broad Training

Once your power users have been using the tech for a while, it's time to roll it out! Before you do that, work with leadership to host a companywide training session. Run through all the labels and what they mean, when to use them, and how they relate to the policy. Also cover the resources available to them (next section).

Job Aides/Quick Reference Card

The policy is an operational necessity to ensure that information labeling is tucked neatly into the information security program for the customer. However, nobody wants to read that thing every time they have a question. To solve for this, consider making job aid or "quick reference card" for labeling. This should be easily accessible to everyone who needs to use labels and be clear and to the point. Something like this:

Label	**Use When**
Public	The information you're working on is meant for public disclosure or is not sensitive in nature.
Company Confidential	Use when the information you're working with is confidential and should remain internal. Everyone at Company can access this information, but nobody outside of company.
Department Confidential\[Your Department]	Use these labels when you are working on sensitive information that needs to remain in the confines of your department.

That's really all users need to know about the labels at a glance (the policy should cover the details). In addition to the labels themselves, consider the following content:

- Quick reference information about using labels in various apps (Outlook, Word, etc.).
- If applying labels to groups and sites, consider documenting that for users to see and understand what that means.
- What to do when something seems under-classified?

Closing it Out

I hope that you've found this guide useful, and that you'll use it to get more of your customers on the data classification train.

A Note on Access Control Schemes

My references to mandatory and discretionary access control may include some liberties I have chosen to take. I apologize in advance to the CISSPs reading this. I'm simply hoping to identify the most closely aligned access control model to the topic at hand.

Got Questions?

The best way to get in touch with me is via LinkedIn at https://linkedin.com/in/dominickirby. You can also reach me and check out my other content on my website: https://domkirby.com.

Other PIP Implementation: Traffic Light Protocol

Doing a lot of *external community collaboration* with sensitive data? Check out my guide on implementing PIP to the Traffic Light Protocol standards: https://domk.pro/piptlp.