



COHESITY

ZSCALER AND COHESITY DEPLOYMENT GUIDE

Contents

Terms and Acronyms	4
About This Document	6
Zscaler Overview	6
Cohesity Overview	6
Audience	6
Software Versions	6
Request for Comments	6
Zscaler and Cohesity Introduction	7
ZIA Overview	7
ZPA Overview	7
Cohesity Data Cloud Overview	8
Cohesity Resources	9
Cohesity and Zscaler Data Loss Prevention Integration	10
Cohesity License Requirements	10
Zscaler DLP Indexed Document Match (IDM)	10
Configuring the Index Tool with VMware	11
Index Tool VM Specifications and Sizing Guidelines	11
Download the Index Tool Image	11
Adding an Index Tool Configuration	12
Configuring the Index Tool VM (VMware)	12
Updating the Index Tool VM	15
Running the Index Tool VM in Explicit Proxy Mode	16
Requirements for Explicit Proxy Mode	16
Index Tool VM Commands	17
Adding an Index Tool Configuration	17

Creating an Indexed Document Match Template	18
Creating an IDM Template	18
Creating a Scheduled IDM Template	19
Modifying an IDM Template	21
Deleting an IDM Template	21
Defining IDM Match Accuracy for Custom DLP Dictionaries	22
Understanding DLP Engine	23
Adding Custom DLP Engines	23
Configuring DLP Policy Rules with Content Inspection	25
Defining DLP Inline Policy Rules	25
Configure Cohesity Data Cloud	28
Analyzing Zscaler DLP Engine Logs	30
Zscaler Data Discovery Report	32
Best Practice for Securing Your Drop Server	33
Appendix A: Requesting Zscaler Support	35

Terms and Acronyms

The following table defines acronyms used in this deployment guide. When applicable, a Request for Change (RFC) is included in the Definition column for your reference.

Acronym	Definition
CA	Central Authority (Zscaler)
CDP	Continuous Data Protection
CSV	Comma-Separated Values
DHCP	Dynamic Host Configuration Protocol
DLP	Data Loss Prevention
DNS	Domain Name Service
DPD	Dead Peer Detection (RFC 3706)
DR	Disaster Recovery
GLBA	Gramm-Leach-Bliley Act
GRE	Generic Routing Encapsulation (RFC2890)
ICMP	Internet Control Message Protocol
IdP	Identity Provider
IKE	Internet Key Exchange (RFC2409)
IPS	Intrusion Prevention System
IPSec	Internet Protocol Security (RFC2411)
ISO	ISO Image
MFA	Multi-Factor Authentication
ML	Machine Learning
NTP	Network Time Protocol
PCI	Payment Card Industry
PFS	Perfect Forward Secrecy
PSK	Pre-Shared Key
RBAC	Role-Based Access Control
SaaS	Software as a Service
SSH	Secure Shell
SSL	Secure Socket Layer (RFC6101)
SSO	Single Sign-On
TLS	Transport Layer Security
VDI	Virtual Desktop Infrastructure
XFF	X-Forwarded-For (RFC7239)
YARA	Yet Another Recursive Acronym
ZPC	Zscaler Posture Control (Zscaler)
ZDX	Zscaler Digital Experience (Zscaler)
ZIA	Zscaler Internet Access (Zscaler)
ZPA	Zscaler Private Access (Zscaler)

Trademark Notice

© 2025 Zscaler, Inc. All rights reserved. Zscaler™ and other trademarks listed at zscaler.com/legal/trademarks are either (i) registered trademarks or service marks or (ii) trademarks or service marks of Zscaler, Inc. in the United States and/or other countries. Any other trademarks are the properties of their respective owners.

About This Document

The following sections describe the organizations and requirements of this deployment guide.

Zscaler Overview

Zscaler (NASDAQ: [ZS](#)) enables the world's leading organizations to securely transform their networks and applications for a mobile and cloud-first world. Its flagship Zscaler Internet Access (ZIA) and Zscaler Private Access (ZPA) services create fast, secure connections between users and applications, regardless of device, location, or network. Zscaler delivers its services 100% in the cloud and offers the simplicity, enhanced security, and improved user experience that traditional appliances or hybrid solutions can't match. Used in more than 185 countries, Zscaler operates a massive, global cloud security platform that protects thousands of enterprises and government agencies from cyberattacks and data loss. To learn more, see [Zscaler's website](#).

Cohesity Overview

Cohesity is an American privately held information technology company headquartered in San Jose, California with offices in India, Ireland, and Costa Rica. The company develops software that allows IT professionals to back up, manage, and gain insights from their data across multiple systems or cloud providers. To learn more, refer to [Cohesity's website](#).

Audience

This guide is for network administrators, endpoint and IT administrators, and security analysts responsible for deploying, monitoring, and managing enterprise security systems. For additional product and company resources, see:

- [Zscaler Resources](#)
- [Cohesity Resources](#)
- [Appendix A: Requesting Zscaler Support](#)

Software Versions


This document was authored using the latest version of Zscaler software.

Request for Comments

- **For prospects and customers:** Zscaler values reader opinions and experiences. Contact partner-doc-support@zscaler.com to offer feedback or corrections for this guide.
- **For Zscaler employees:** Contact z-bd-sa@zscaler.com to reach the team that validated and authored the integrations in this document.

Zscaler and Cohesity Introduction

Overviews of the Zscaler and Cohesity applications are described in this section.

 If you are using this guide to implement a solution at a government agency, some of the content might be different for your deployment. Efforts are made throughout the guide to note where government agencies might need different parameters or input. If you have questions, contact your Zscaler Account team.

ZIA Overview

ZIA is a secure internet and web gateway delivered as a service from the cloud. Think of ZIA as a secure internet on-ramp—just make Zscaler your next hop to the internet via one of the following methods:

- Setting up a tunnel (GRE or IPSec) to the closest Zscaler data center (for offices).
- Forwarding traffic via our lightweight Zscaler Client Connector or PAC file (for mobile employees).

No matter where users connect—a coffee shop in Milan, a hotel in Hong Kong, or a VDI instance in South Korea—they get identical protection. ZIA sits between your users and the internet and inspects every transaction inline across multiple security techniques (even within SSL).

You get full protection from web and internet threats. The Zscaler cloud platform supports Cloud Firewall, IPS, Sandboxing, DLP, and Isolation, allowing you to start with the services you need now and activate others as your needs grow.

ZPA Overview

ZPA is a cloud service that provides secure remote access to internal applications running on a cloud or data center using a Zero Trust framework. With ZPA, applications are never exposed to the internet, making them completely invisible to unauthorized users. The service enables the applications to connect to users via inside-out connectivity rather than extending the network to them.

ZPA provides a simple, secure, and effective way to access internal applications. Access is based on policies created by the IT administrator within the ZPA Admin Portal and hosted within the Zscaler cloud. On each user device, software called Zscaler Client Connector is installed. Zscaler Client Connector ensures the user's device posture and extends a secure microtunnel out to the Zscaler cloud when a user attempts to access an internal application.

Zscaler Resources

The following table contains links to Zscaler resources based on general topic areas.

Name	Definition
ZIA Help Portal	Help articles for ZIA.
ZPA Help Portal	Help articles for ZPA.
Zscaler Tools	Troubleshooting, security and analytics, and browser extensions that help Zscaler determine your security needs.
Zscaler Training and Certification	Training designed to help you maximize Zscaler products.
Submit a Zscaler Support Ticket	Zscaler Support portal for submitting requests and issues.

The following table contains links to Zscaler resources for government agencies

Name	Definition
ZIA Help Portal	Help articles for ZIA.
ZPA Help Portal	Help articles for ZPA.
Zscaler Tools	Troubleshooting, security and analytics, and browser extensions that help Zscaler determine your security needs.
Zscaler Training and Certification	Training designed to help you maximize Zscaler products.
Submit a Zscaler Support Ticket	Zscaler Support portal for submitting requests and issues.

Cohesity Data Cloud Overview

Cohesity Data Cloud is a cloud data management platform for securing and managing enterprise data that reduces your attack surface, lowers costs, and minimizes risk. Cohesity Data Cloud is available as self-managed software and SaaS with rich features, including:

- **Modern Backup and Recovery:** The most comprehensive, modern, web-scale data management and backup and recovery solution to protect cloud-native, SaaS, and on-premises data at scale. You get instant recovery at scale and with fully hydrated snapshots and CDP.
- **Traditional and Modern Workloads:** Support for VMs, databases, files, containers, cloud-native, SaaS, storage, and traditional workloads.
- **Defend Against Ransomware Attacks:** Multilayered security architecture with Zero Trust security, including granular RBAC, MFA, SSO, immutable snapshots, and ML-based ransomware attack detection. Protect and recover against ransomware with threat protection, cyber vaulting, and ML-powered data classification.
- **Threat Protection and Data Classification:** Highly curated and managed threat feeds, trained with ML, threat detection and response to your specific needs by augmenting the extensive library of over 117,000 behavioral patterns, create YARA rules defining Indicators of Compromise (IoC), or import custom rules. Highly accurate ML-based engine to classify sensitive data, automatically or on-demand, including personally identifiable information (PII), PCI, and HIPAA.
- **Global Search and Unified Management:** Reduce recovery point objectives to minutes by eliminating slow, chain-based backups. A single management platform offering multilayered security architecture, unifying operations with integrated solutions for backup, CDP, DR, search, ransomware attack detection, and vulnerability scanning into a single scalable environment.
- **Cohesity-managed Cloud Vault:** Cohesity FortKnox is a SaaS cyber vaulting and recovery solution that gives your data an additional layer of managed security and protection against cybersecurity threats.
- **Cohesity CloudArchive Direct:** Policy-based data archival to meet long-term data retention, compliance, and regulatory requirements.
- **Cohesity Cloud Services:** Cohesity-managed data security and management with SaaS that runs multiple cloud data services, including backup, cyber vaulting, threat defense, data classification, DR, and more on a single multi-cloud platform.
- **Business Continuity:** Simplify business continuity and disaster recovery with automated failover and failback orchestration for your mission-critical workloads. Get your critical applications online after a breach or outage through automated orchestration.
- **Security Integrations:** Cohesity integrates with leading perimeter and end-point security vendors, giving you greater visibility and actionable alerts in your Security Operations Center (SOC).

- Deployment: Software-defined solution for on-premises, public cloud, backup as a service, and edge sites.
- API-first Extensibility: Derive business insights with the Cohesity Marketplace partner ecosystem. Streamline operations and easily integrate on-premises and cloud environments with prebuilt automated workflows and API integrations.

Cohesity Resources

The following table contains links to Cohesity support resources.

Name	Definition
Cohesity Support	Online support for Cohesity customers.
Cohesity Documentation	Online documentation for Cohesity cloud services.
Cohesity User Group	Online community for Cohesity users.

Cohesity and Zscaler Data Loss Prevention Integration

The two major solution components of the Cohesity and Zscaler integration are Cohesity Data Cloud and Zscaler DLP. Cohesity Data Cloud does backup, recovery, and sensitive data discovery. The sensitive data is then sent to the Zscaler DLP solution and fingerprinted and indexed. The Zscaler Client Connector blocks or allows the transmission of the data as per the user-configured policy.

The following diagram shows a conceptualization of the integration.



Figure 1. Cohesity and Zscaler integration

Cohesity License Requirements

The following Cohesity licenses are required for the Zscaler and Cohesity integration:

- Cohesity DataProtect license
- Cohesity Data Classification License


Zscaler DLP Indexed Document Match (IDM)

Zscaler uses [Indexed Document Match \(IDM\) templates](#) (government agencies, see [Indexed Document Match \(IDM\) templates](#)) to fingerprint which critical documents contain sensitive data in your organization. Fingerprinting and indexing your documents creates a document repository that the Zscaler service uses to identify which documents partially or completely match the DLP policy when evaluating outbound traffic.


Creating an IDM template requires [Zscaler's Index Tool](#) (government agencies, see [Zscaler's Index Tool](#)). You can upload text-based files and non-text-based documents (e.g., binary files). After an IDM template is created, you can then apply the template to a custom DLP dictionary. When adding the template to the dictionary, you must choose the match accuracy level for the template. Match accuracy is the level of accuracy (i.e., the percentage of similarity) that a document must meet to count as a match for an indexed document.

Configuring the Index Tool with VMware

The following sections describe configuring the Index Tool with VMware.

 New or clean deployment of the Index Tool requires a VM image running on ZIA version 24.

Follow the directions to complete configuring the Index Tool.

 Since the Index Tool provides access to highly sensitive information, ensure that everyone who has access to it is authorized and authenticated.

Index Tool VM Specifications and Sizing Guidelines

If your index templates include less than 300 million records, Zscaler recommends the following configuration:

- Hypervisor: VMware ESX/ESXi version 6.0 or later.
- CPUs: 4 CPUs. Zscaler requires 4 CPUs because the CPUs ensure that hash generation performance is not impacted.
- RAM: 16 GB
- Disk: 600 GB
- VM Network: 1 Virtual NIC

If your index templates include more than 300 million records, Zscaler recommends the following configuration:

- Hypervisor: VMware ESX/ESXi version 6.0 or later.
- CPUs: 4 CPUs. Zscaler requires 4 CPUs because the CPUs ensure that hash generation performance is not impacted.
- RAM: 64 GB
- Disk: 1 TB
- VM Network: 1 Virtual NIC

Download the Index Tool Image

You must download the Index Tool VM before you configure it.

If your index templates include less than 300 million records, you can download the Index Tool VM image from the ZIA Admin Portal. To download the Index Tool VM:

1. Go to **Administration > Index Tool**.
2. Click **Download Index Tool**.




Figure 2. Zscaler Index Tool download page

Adding an Index Tool Configuration

To add an Index Tool configuration:

1. Go to **Administration > Index Tool**.
2. Click **Add Index Tool Configuration**. The **Add Index Tool Configuration** window is displayed.
3. In the **Add Index Tool Configuration** window:
 - a. **VM Name**: Enter a unique name for the virtual machine (VM).
 - b. **Status**: Make sure that the VM is **Enabled**.

Figure 3. Zscaler Index Tool configuration wizard

4. Click **Save** and [activate the change](#).

After you save the SSL Certificate for the configuration, you can download it from the Index Tool page or from the Edit Index Tool Configuration window.

Configuring the Index Tool VM (VMware)

To configure the Index Tool VM:

1. Make sure you have added an [Index Tool Configuration](#) (government agencies, see [Index Tool Configuration](#)). You need this configuration to complete the VM setup.
2. In ESX/ESXi, install the Index Tool VM image you downloaded previously.
3. Boot up the VM and log in as user `zsroot`. The initial root password for this user is randomly generated.

```
Tue Jan 23 23:04:27 UTC 2018
+-----+
|  zsroot password = 7BVdAB1g  |
|  Change password ASAP by    |
|  running following in CLI   |
|  sudo zadb change-password  |
+-----+

FreeBSD/amd64 (zadb) (ttyv0)

login: zsroot
Password: █
```

Figure 4. Zscaler Index Tool—CLI

4. To change the root password:

- a. Entering the following command:

```
sudo zadb change-password
```

- b. Enter the initial root password that was randomly generated for you.

```
Last login: Tue Jan 23 22:52:38 on ttyv0
Copyright (c) 1992-2013 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
    The Regents of the University of California. All rights reserved.

FreeBSD 8.4-RELEASE (SMKERNEL) #0: Tue Dec  3 17:46:10 PST 2013

Welcome to EDM Client UM!

Please change the default zsroot password as soon as possible.

To setup the UM with your account, follow the user guide.

Several useful commands:
- sudo zadb configure-network
- sudo zadb configure <client cert bundle zip filename>
- sudo zadb stop
- sudo zadb start
- sudo zadb status
- sudo zadb restart
- sudo zadb update-now
- sudo zadb force-update-now
- sudo zadb troubleshoot
[zsroot@zadb ~]# sudo zadb change-password
Password: █
```

Figure 5. Zscaler Index Tool—CLI wizard

- c. Enter the new root password.

```
    The Regents of the University of California. All rights reserved.

FreeBSD 8.4-RELEASE (SMKERNEL) #0: Tue Dec  3 17:46:10 PST 2013

Welcome to EDM Client UM!

Please change the default zsroot password as soon as possible.

To setup the UM with your account, follow the user guide.

Several useful commands:
- sudo zadb configure-network
- sudo zadb configure <client cert bundle zip filename>
- sudo zadb stop
- sudo zadb start
- sudo zadb status
- sudo zadb restart
- sudo zadb update-now
- sudo zadb force-update-now
- sudo zadb troubleshoot
[zsroot@zadb ~]# sudo zadb change-password
Password: █
-----
Changing local password for zsroot
New Password: █
```

Figure 6. Zscaler Index Tool—Password Reset

- d. Re-enter the new root password. After changing the password, you must log in to zsroot again using the new password.

5. (Optional) By default, the VM starts using DHCP to obtain the IP address and default router information. If there's no DHCP server available, you can configure it manually:
 - a. Enter the following command:


```
sudo zadp configure-network
```
 - b. For nameserver, enter `c` to change the IP address and press `Enter`.
 - c. Enter the IP address and press `Enter`.
 - d. If you want to add a new nameserver, enter `y`, otherwise enter `n`, and press `Enter`. The VM restarts the network and checks the connection.

```
[zsroot@zadp ~]$ sudo zadp configure-network
Password:
-----
nameserver:8.8.8.8 (Options <c:change, d:delete, n:no change>) [n]c
nameserver (Resolver IP address) [8.8.8.8]: 10.32.112.10
Do you wish to add a new nameserver? <n:no y:yes> [n]: n
ifconfig_em0 (IP/CIDR or DHCP (1.2.3.4/24, DHCP)) [DHCP]: 10.66.103.177/24
defaultrouter (IP or NO for DHCP (1.2.3.4, NO)) [NO]: 10.66.103.254
hostname (Host name of this VM (zadp)) [zadp]:
Network configuration has been changed, restart network, please wait...
Network changes has been successfully applied.
Checking network connection by pinging zscaler.com.
Successful pinging zscaler.com, network looks running fine.
Syncing system date and time, please wait...
Syncing system date and time has been completed.
[zsroot@zadp ~]$
```

Figure 7. Zscaler Index Tool—System Network Configuration

6. Return to the ZIA Admin Portal and go to **Administration** > **Index Tool**.
7. Locate the [Index Tool Configuration](#) (government agencies, see [Index Tool Configuration](#)) you added previously, and under the SSL Certificate column, click **Download**.

No.	VM name	Status	SSL Certificate
1	Index Tool 1	Enabled	Download

Figure 8. SSL Certificate

8. Copy the SSL client certificate.zip file to the VM and install it:
 - a. This example uses `scp` to copy the file (for example,


```
scp EdmClientCertificate.zip zsroot@10.66.108.100:~/
```

```
scp <SSL_certificate_zip_filename> zsroot@<vm_ip>:~/
```
 - b. Enter the following command to install the SSL certificate (for example,


```
sudo zadp configure EdmClientCertificate.zip
```

```
sudo zadp configure <SSL_certificate_zip_filename>
```
 - c. Enter the domain name used for the Index Tool's fully qualified domain name (FQDN). For example, if the Index Tool is reachable from `indextool.mycompany.com`, then the domain name entered here is `mycompany.com`. The self-signed certificate is generated for `*.mycompany.com`.

```

[zsroot@zadp ~]$ sudo zadp configure EdmClientCertificate.zip
Password:
-----
Stopping zadp service if running...
killing process 864 chown.
ZADP service has been stopped
-----
Client certificate file has been installed into /sc/conf/zscaler_edm_certificate
.crt and /sc/conf/zscaler_edm_key.key.
-----
Generating default configuration file...
Default configuration file has been generated at /sc/conf/sc.conf.
-----
Self signed certificate generation for webui
-----
Please enter domain name will be used for self signed cert: █

```

Figure 9. Zscaler Index Tool – Install SSL Certificate

- d. Enter a passphrase, then re-enter the passphrase to confirm it. You are prompted to type the full path name to the text file where the passphrase is stored. You can also press Enter twice to accept the default location and file `/home/zsroot/zscaler_zadp_webui_certificate_pass.txt`.

If the service was configured properly, the service:

- Checks if the network is configured correctly.
- Installs the SSL client certificate you specified.
- Generates a self-signed SSL server certificate.
- Downloads the latest install package.
- Starts the service.

9. (Optional) If you need to install a self-signed or custom SSL server certificate:

- a. Enter the following command to install the server certificate:

```
sudo zadp install-server-cert
```

- b. Enter the full path to the PEM-formatted certificate file.

- c. Enter the following command to restart the Index Tool service:

```
sudo zadp restart
```

Go to <https://<IP Address of the Index Tool VM>> to access the Index Tool. After the Index Tool service has started, you can log in with your ZIA Admin Portal login credentials and create Index Templates to use when creating DLP dictionaries.

To learn more, see [Creating an Exact Data Match Template](#) and [Creating an Indexed Document Match Template](#) (government agencies, see [Creating an Exact Data Match Template](#) and [Creating an Indexed Document Match Template](#)).

Updating the Index Tool VM

If you successfully configured the Index Tool, the service automatically downloads the latest install package before it starts. To manually update the service:

1. Enter the following command to stop the service:

```
sudo zadp stop
```

2. Enter the following command to update the service:

```
sudo zadp update-now
```

3. Enter the following command to start the service:

```
sudo zadp start
```

Running the Index Tool VM in Explicit Proxy Mode

You can run the tool in explicit proxy mode:

1. Log in to the VM as user `zsroot`.
2. Enter the following command:


```
sudo zadp configure-network
```
3. For **Do you require a proxy server configuration?**, enter `y` and press `Enter`.
4. For **proxyserver**, enter the IP address of your proxy server (e.g., `proxy.zscaler.net`) and press `Enter`.
5. For **proxyport**, enter your proxy port number (e.g., `9443`) and press `Enter`. The VM then tests the connection and after this is successful, the configuration is complete.

To remove the explicit proxy configuration:

6. Enter the following command:


```
sudo zadp configure-network
```
7. For **Do you require proxy server configuration?**, enter `n` and press `Enter`.
8. For **Do you want to delete current proxy configuration?**, enter `y` and press `Enter`.

Requirements for Explicit Proxy Mode

If you're using explicit proxy mode, DNS and NTP connections are not tunneled. You need an internal DNS server to run in this mode. The Index Tool must have DNS resolution for the current Zscaler Central Authority (CA) IP, update server, and the NTP server. The Index Tool host also must query a DNS server to resolve the following settings:

- `smcacluster.<Zscaler Cloud Name>`
- `update1.<Zscaler Cloud Name>`
- `update2.<Zscaler Cloud Name>`
- `zistribute.<Zscaler Cloud Name>`
- The NTP server. By default, the Index Tool VM has the following FQDNs for NTP servers configured:
 - `0.freebsd.pool.ntp.org`
 - `1.freebsd.pool.ntp.org`
 - `2.freebsd.pool.ntp.org`

You can override these FQDNs to your internal IP address in your DNS server configuration or using other methods.

In addition, since the proxy configuration doesn't allow authentication, you must configure the proxy server to allow specific IP/MAC addresses without user and password authentication.

The proxy server must also allow SSL bypass for communication from the VM to a specific set of IP addresses. These IPs are listed at `config.zscaler.com/<Zscaler Cloud Name>.net/edm`. You can find your cloud name in the URL that your admins use to log in to the Zscaler service. For example, if an organization logs in to `admin.zscalertwo.net`, then that organization's cloud name is `zscalertwo`. You would go to `config.zscaler.com/zscalertwo.net`. To learn more, see [What is my cloud name for ZIA](#) (government agencies, see [What is my cloud name for ZIA](#)).

Index Tool VM Commands

You can enter the following commands to configure, update, and troubleshoot your VM.

Command	Description
<code>sudo zadb stop</code>	Stops the Index Tool service.
<code>sudo zadb start</code>	Starts the Index Tool service.
<code>sudo zadb restart</code>	Restarts the Index Tool service.
<code>sudo zadb status</code>	Displays whether the Index Tool service is running or stopped.
<code>sudo zadb update-now</code>	Updates the Index Tool service. The service must be stopped before you can run this command.
<code>sudo zadb force-update-now</code>	Forces the Index Tool service to update to the latest version regardless of what version is on the VM. The service is automatically stopped before the update begins.
<code>sudo zadb troubleshoot</code>	Runs a series of checks to help troubleshoot issues, such as checking the installed certificates, the zcloud server configuration, all services, and whether or not an update is needed.

Adding an Index Tool Configuration

To add an Index Tool Configuration, follow the steps shown previously in [Adding an Index Tool Configuration](#).

Creating an Indexed Document Match Template

Using the Index Tool, you can create, modify, or delete an Indexed Document Match (IDM) index template.



You can create up to 64 IDM templates for your organization. The largest file you can upload to an IDM template is 400 MB. You can index up to 100 GB of files for your organization.

Creating an IDM Template

You can create scheduled or manual IDM templates.



The integration with the Cohesity platform is focused on Scheduled IDM Templates and does not support the Manual IDM Template method.

- **Scheduled IDM Templates:** A scheduled IDM template allows you to set up an SSH connection and schedule updates between the template and your organization's storage server for the critical documents.
- **Manual IDM Templates:** A manual IDM template allows you to manually upload your organization's critical documents.

To create a Scheduled IDM template:

1. Go to <https://<IP Address of the Index Tool VM>> to access the Index Tool. Log in to the Index Tool with your ZIA Admin Portal login credentials.

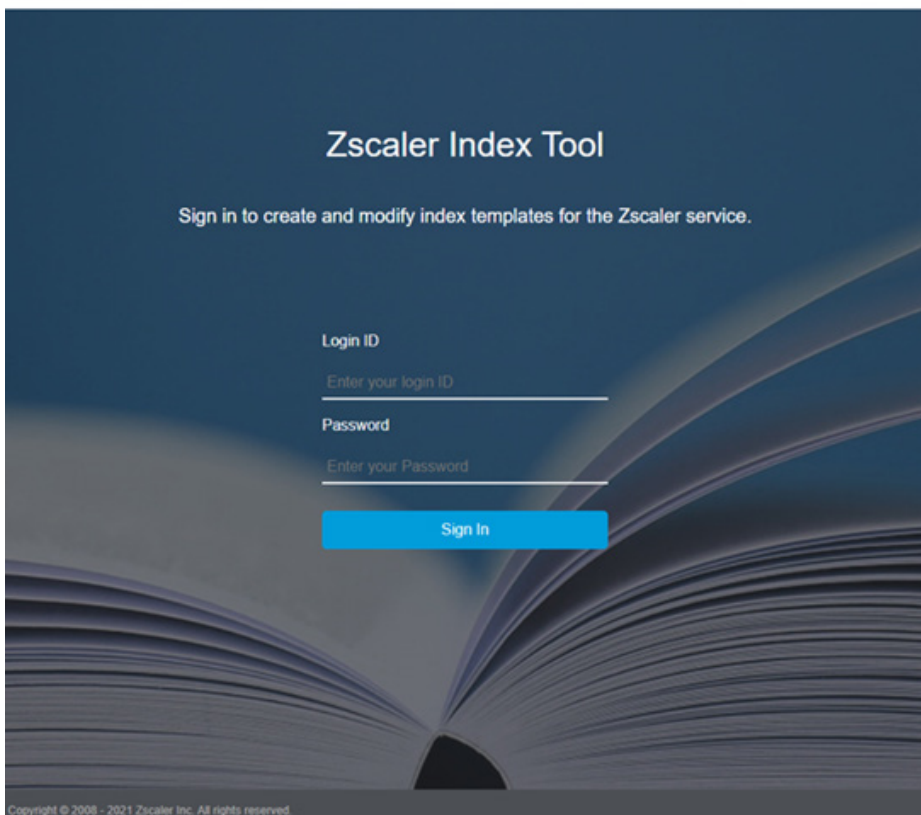


Figure 10. Login to Index Tool—Web UI

2. Click **Indexed Document Match Templates**.
3. On the **Indexed Document Match Templates** tab, click **Create New Template** and then select **Scheduled IDM Template**.

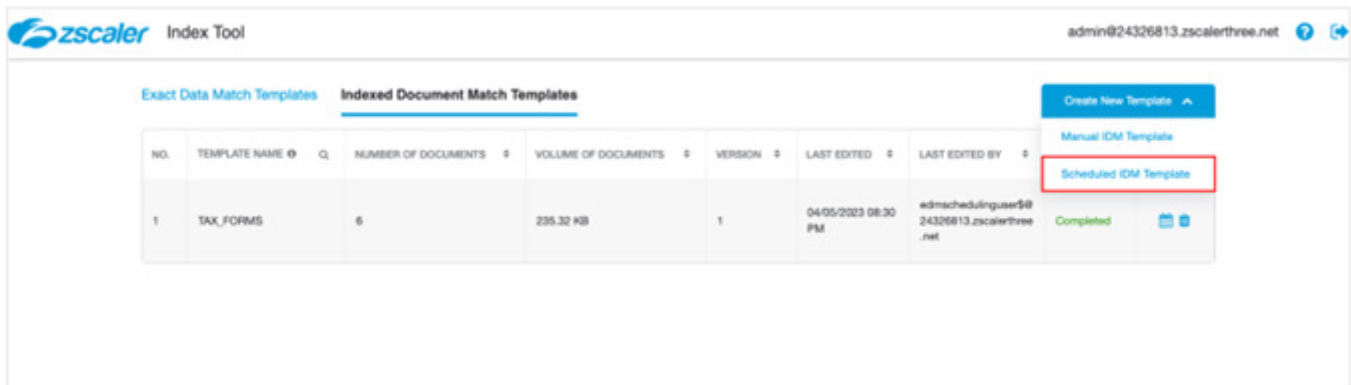


Figure 11. Scheduled IDM Template

Creating a Scheduled IDM Template

When you create a scheduled IDM template, you must set up an SSH connection between the template and your organization's document server. You must then schedule regular updates between the template and server.

In the Scheduled Indexed Document Match Template window:

1. Under **General**:
 - a. Enter a **Template Name**. After the template is saved, the name appears in [Indexed Document Match](#) (government agencies, see [Indexed Document Match](#)) in the ZIA Admin Portal.
 - b. For **Host**, enter the IP address or domain for the document server.
 - c. Enter the **Port** for the document server.
 - d. Enter the **File Path** for the directory where the documents are located in the document server.

Scheduled Indexed Document Match Template

General

Template Name	Host ?	Port
<u>Example IDM Template</u>	<u>12.34.56.78</u>	<u>22</u>
File Path		
<u>/user/documents/123456</u>		

Figure 12. Configure Scheduled Indexed Document Match Template

2. Under **SSH Configuration**:
 - a. Click **Download** to download the SSH key.
 - b. Copy the username. Use this username to create a user in your document server.
 - c. Go to your document server and complete the following steps:
 - Create a user with the username you copied from the Index Tool.
 - Add the downloaded SSH key to the user's trusted keys.
 - Ensure that the user has read access to the directory in the specified file path.

- d. Click **Verify** in the template to verify the SSH setup configuration. You cannot save the template until the setup is configured properly and verified.

SSH Configuration

Step 1: Download the SSH key
Download

Step 2: Copy the username
zdm24326813

Step 3: Give the Index Tool access to the documents

1. In the document server, create a user with the copied username
2. Add the downloaded SSH key to the user's trusted SSH keys
3. Ensure that the user has read access to the directory in the file path

Verify

Figure 13. SSH Configuration

3. Under **Schedule**:
- Repeat:** Choose if the update repeats **Monthly**, **Weekly**, or **Daily**.
 - Every:** If you selected **Monthly** or **Weekly**, choose when in the selected period the update repeats. For example, if you selected **Monthly**, you can choose the day of the month to update or if you selected **Weekly**, you can choose to have the update happen every Friday.
 - Time:** Select what time the schedule update happens.
 - Time Zone:** Select the time zone your update happens.
 - Update Now:** Select to immediately update the template.

Schedule

Repeat: Monthly

Every: 6

Time: 6:30 AM

Time Zone: GMT-8:00 Pacific Standard Time (PST)

Update Now ?

Your current time is 11:59 AM PST

Figure 14. Configure Schedule

- (Optional) Enter a **Description** for the template.
- Click **Save**.

After saving the template, you are redirected to the Indexed Document Match Template page, and the tool processes the template. If the template was created properly, Completed is shown in the Status column. If the template was created, but the documents weren't indexed yet, then Created is shown. If the template was not created properly, then Error is shown.

After an IDM template is created, it appears in [Indexed Document Match](#) (government agencies, see [Indexed Document Match](#)) of the ZIA Admin Portal, where you can view the template's details or delete it. You cannot change the template name after creation. To change the name, you must create a new template.

Modifying an IDM Template

To submit new documents or delete indexed documents for scheduled IDM templates, you must make the changes in the document server. The template updates at the scheduled time, or you can schedule an immediate update for the template in the Index Tool. To reschedule a scheduled template's update, click the **Calendar** icon in the **Actions** column.









NO.	TEMPLATE NAME 	NUMBER OF DOCUMENTS 	VOLUME OF DOCUMENTS 	VERSION 	LAST EDITED 	LAST EDITED BY 	STATUS	ACTIONS
1	Scheduled IDM template	5	464.82 KB	6	03/16/2021 10:14 PM	admin@safemarch.com	Completed	 

Figure 15. Configure Schedule

Deleting an IDM Template

To delete a scheduled IDM template:

1. Log in to the Index Tool.
2. On the **Indexed Document Match** tab, locate the template you want to remove. Click the **Search** icon to search for a specific template.
3. In the **Actions** column, click the **Delete** icon.
4. In the confirmation window that appears, click **Delete**.

You can also delete the IDM template from the ZIA Admin Portal:

1. Go to **Administration > Index Templates**.
2. In the **Indexed Document Match** tab, locate the template you want to remove.
3. Click **Delete**.




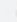
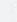


Index Templates							
Exact Data Match		Indexed Document Match					
No.	Template Name 	VM Name 	Number of Documents 	Last Modified 	Modified By 		
1	Scheduled IDM template	Index Tool	5	3/16/2021, 10:14 PM	admin@safemarch.com		

Figure 16. Index Templates

4. Click **OK**.

Defining IDM Match Accuracy for Custom DLP Dictionaries

You can add Indexed Document Match (IDM) templates to custom DLP dictionaries that represent critical documents that you want to protect in your organization. When adding an IDM template, you must also choose the match accuracy level for the template in the dictionary. To learn more, see [Creating an Indexed Document Match Template](#) and [Adding Custom DLP Dictionaries](#) (government agencies, see [Creating an Indexed Document Match Template](#) and [Adding Custom DLP Dictionaries](#)).

To add an IDM template:

1. Go to **Administration > DLP Dictionaries & Engines**.
2. On the **DLP Dictionaries** tab, click **Add DLP Dictionary** or click the **Edit** icon for an existing dictionary. The **Add/Edit DLP Dictionary** window appears.
3. In the **Add/Edit DLP Dictionary** window:
 - a. From **Dictionary Type**, choose **Indexed Document Match**.
 - b. From the **Index Template**, select the IDM template you want to use for the dictionary.
 - c. Choose the **Match Accuracy** for the IDM template you selected. This is the level of accuracy (i.e., the percentage of similarity) that a user-uploaded document must meet to count as a match for an indexed document.
 - **Low**: Zscaler detects a low-accuracy match when one of the following occurs:
 - A user-uploaded document matches at least 40% of an indexed document.
 - An indexed document matches at least 70% of a user-uploaded document.
 - **Medium**: Zscaler detects a medium-accuracy match when one of the following occurs:
 - A user-uploaded document matches at least 70% of an indexed document.
 - An indexed document matches at least 90% of a user-uploaded document.
 - **High**: Zscaler detects a high-accuracy match when a user-uploaded document matches at least 90% of an indexed document.

Figure 17. DLP Dictionary

4. Click **Save** and [activate the change](#).

Understanding DLP Engine

A [DLP engine](#) (government agencies, see [DLP engine](#)) is a collection of one or more DLP dictionaries. The Zscaler service provides predefined DLP engines and supports custom DLP engines:

- **Predefined DLP Engines:** The Zscaler service provides 5 predefined engines (HIPAA, GLBA, PCI, Offensive Language, and Self-Harm & Cyberbullying). These engines contain default DLP dictionaries. For example, the PCI engine contains the Credit Cards and Social Security Number dictionaries. You can also edit predefined engines. To learn more, see [Editing Predefined DLP Engines](#) (government agencies, see [Editing Predefined DLP Engines](#)).

The Cohesity data protection integration leverages both Zscaler's Custom DLP Dictionary in combination with Custom DLP Dictionaries.

- **Custom DLP Engines:** You can create custom DLP engines to detect content that is relevant to your organization. You can create a maximum of 47 custom DLP engines. To learn more, see [Adding Custom DLP Engines](#) (government agencies, see [Adding Custom DLP Engines](#)).

Adding Custom DLP Engines

Adding a custom DLP engine is one of the tasks you can complete when configuring [DLP policy rules](#) (government agencies, see [DLP policy rules](#)). To learn more about the ranges and limitations for custom DLP engines, see [Ranges & Limitations](#) (government agencies, see [Ranges & Limitations](#)).

To add a custom DLP engine:

1. Go to **Administration > DLP Dictionaries & Engines**.
2. In the **DLP Engines** tab, click **Add DLP Engine**. The **Add DLP Engine** window is displayed.
3. In the **Add DLP Engine** window, enter the **Name** for the custom DLP engine.
4. For **Engine Builder**, add operators and DLP dictionaries to [build an expression](#) (government agencies, see [build an expression](#)). You can see your expression in the **Expression Preview**.

The screenshot shows the 'ENGINE BUILDER' interface. At the top, there's a section titled 'EXPRESSION'. Below this, there's a dropdown menu currently set to 'ALL'. Underneath, there's a 'Select a dictionary' dropdown menu with a red 'i' icon and an 'x' to its right, indicating an error or warning. Below that is a '+ ADD' button. At the bottom of the 'EXPRESSION' section, there's an 'Expression Preview' section which currently shows an empty field with a '0' character.

Figure 18. DLP Engine

5. Under **Expression**:

- a. Click **Add** to add a **Dictionary** or a **Subexpression**. Click **Remove (X)** to delete dictionaries or subexpressions. If you click **Dictionary**, you must select the custom DLP dictionary that contains the Index template associated with the Cohesity tenant as shown in [Defining IDM Match Accuracy for Custom DLP Dictionaries](#). Certain dictionaries require you to set a value for the [Configuring the Match Count](#) (government agencies, see [Configuring the Match Count](#)). You can enter any value less than 1000.

The screenshot shows the 'Edit DLP Engine' window. The 'Name' field is 'Cohesity-Engine'. Under 'ENGINE BUILDER', the 'EXPRESSION' section shows a tree view starting with 'ALL'. Below 'ALL' is a sub-expression 'Cohesity IDM' with a remove button (X) and an 'ADD' button. The 'Expression Preview' shows '!(Cohesity IDM > 0)'. The 'DESCRIPTION' field contains 'Cohesity DLP Engine - KR IDM1'. At the bottom are 'Save', 'Cancel', and 'Delete' buttons.

Figure 19. Indexed DLP Engine



For the root subexpression, only the All (AND) and Any (OR) operators are allowed.

- b. Continue adding dictionaries and operators to the expression as needed. At each level, you can create up to 4 subexpressions, use up to 4 operators, and add up to 16 dictionaries per operator.
6. (Optional) For **Description**, enter any additional notes or information. The description cannot exceed 255 characters.
 7. Click **Save** and [activate the change](#).

Configuring DLP Policy Rules with Content Inspection

You can use Zscaler's DLP engines to detect data, allow or block transactions, and notify your organization's auditor when a user's transaction triggers a DLP rule. If your organization has a third-party DLP solution, Zscaler can forward information about transactions that trigger DLP policy to your third-party solution via secure Internet Content Adaptation Protocol (ICAP). However, Zscaler does not take ICAP responses from your DLP solution.

Zscaler only monitors or blocks content according to the policy you configure, then forwards information about transactions so that your organization can take necessary remediation steps.



The Zscaler DLP engines support files up to 400 MB and can scan the first 100 MB of the extracted text. The maximum size also applies to files extracted from archive files.

To configure a DLP policy rule with content inspection, make sure you have completed the following steps:

- [Defining IDM Match Accuracy for Custom DLP Dictionaries](#)
- [Adding Custom DLP Engines](#)

Defining DLP Inline Policy Rules

To create a DLP inline policy to inspect traffic matching the DLP engine associated with the DLP Dictionary containing the Indexed Documents delivered by Cohesity:

1. Go to **Policy > Data Loss Prevention**.
2. Click **Add** and select **Rule With Content Inspection**.
3. In the **Add DLP Rule** window, enter the following DLP Rule attributes:
 - a. **Rule Order:** Policy rules are evaluated in ascending numerical order (Rule 1 before Rule 2, etc.), and the Rule Order reflects this rule's place in the order. You can change the value, but if you've enabled Admin Ranking, then the assigned Admin Rank determines the Rule Order values you can select.
 - b. **Admin Rank:** Enter a value from 0–7 (0 is the highest rank). Your assigned admin rank determines the values you can select. You cannot select a rank that is higher than your own. The rule's Admin Rank determines the value you can select in Rule Order, so that a rule with a higher Admin Rank always precedes a rule with a lower Admin Rank.
 - c. **Rule Name:** Enter a unique name for the DLP rule or use the default name.
 - d. **Rule Status:** An enabled rule is actively enforced. A disabled rule is not actively enforced but does not lose its place in the Rule Order. The service skips it and moves to the next rule.
 - e. **Rule Label:** Select a rule label to associate it with the rule.
 - f. **DLP Engines:** Select the Cohesity DLP Engine created in the section [Adding Custom DLP Engines](#).
 - g. The **Match Only** option only takes effect if the rule's **Action** is **Allow**. You can select **Match Only** to configure how engines must trigger for the service to take action.

Figure 20. DLP Inline Policy



You can choose up to 4 DLP Engines per inline policy.

If your organization requires notification alerts sent to an auditor's mailbox or other automated system for analysis or incident creation, you can configure an email notification for the rule.



If you do not select an auditor and notification template, a notification is not sent for this rule.

1. For **Auditor Type**, select whether the auditor is from a **Hosted** database or is **External** to your organization.
2. Select the **Auditor**:
 - a. If the auditor is from a hosted database, select or search for the auditor.
 - b. If the auditor is external, enter the auditor's email address.
 - c. Select a **Notification Template** if you configured one. See [Configuring DLP Notification Templates](#) (government agencies, see [Configuring DLP Notification Templates](#)). You can also search for a notification template or click the **Add** icon to add a new notification template.
 - d. **OCR**: Enable this option to allow Zscaler's DLP engines to scan images in files. If this option is disabled, the DLP rule doesn't apply to image files.



To enable this option for your organization, contact your Zscaler Account team.

- e. **Inspect Downloads**: Enable this option to allow DLP inspection for content downloaded from specific cloud apps. If this option is enabled, you must also choose **Any** for **URL Categories** and at least one cloud app for **Cloud Application**. If disabled, the DLP rule only applies to content sent to cloud apps.

3. (Optional) For **DLP Incident Receiver**, complete one of the following:
 - a. If you don't have a third-party DLP solution or don't want to forward content, leave the **Zscaler Incident Receiver** or **ICAP Receiver** field as **None**.
 - b. If you want to forward the transactions captured by this policy rule to an on-premises DLP incident receiver:
 - For **Incident Receiver**, select whether the DLP incident receiver is an ICAP receiver or a Zscaler Incident Receiver.
 - Select the applicable ICAP Receiver or Zscaler Incident Receiver from the drop-down menu. You must configure your [ICAP receivers](#) or [Zscaler Incident Receivers](#) (government agencies, see [ICAP receivers](#) or [Zscaler Incident Receivers](#)) in order to complete this step.
4. Select the **Action** for the rule. You can **Allow** or **Block** transactions that match the rule. If you select **Allow**, the service allows and logs the transaction. If you select **Block**, the service blocks and logs the transaction.
5. (Optional) **Configure Client Connector Notification**: You can **Enable** or **Disable** Client Connector notifications for the rule when violations occur. The field is only available if you enable the **Web DLP Violations** option for your organization on the **End User Notifications** page in the ZIA Admin Portal and you select **Block** as the **Action** for the rule. To learn more, see [Configuring Browser-Based End User Notifications](#) (government agencies, see [Configuring Browser-Based End User Notifications](#)).
6. Click **Save** and [activate the change](#).

For example, if a policy rule using predefined Zscaler DLP engines is configured (as shown in the following image), the Zscaler service blocks all the files that:

- Contain medical information
- Are over 1000 KB in size
- Are sent by users in the Operations group through Gmail

The Zscaler service sends an email notification regarding the policy violation to your organization's auditor but doesn't forward information to an incident receiver.

Configure Cohesity Data Cloud

To configure the Cohesity Data Cloud, perform the following configuration steps.

Log in to Cohesity Data Cloud with your credentials.

1. Click **Security > Security Center**.
2. From the left-side navigation, click **Integrations > Browse Integrations**.
3. In the **Data Loss Prevention** section, click **Configure** under the Zscaler integration.
4. On the **Zscaler Configuration Guide** dialog, review the Zscaler prerequisite steps. Click **Continue**.
5. On the **Zscaler Configuration** dialog, do the following configuration:
 - a. In the **Select Document Servers** section, search to find and select one or more document servers where you want to store the sensitive files detected by Data Classification scans. Click **Next: Configuration Options**.

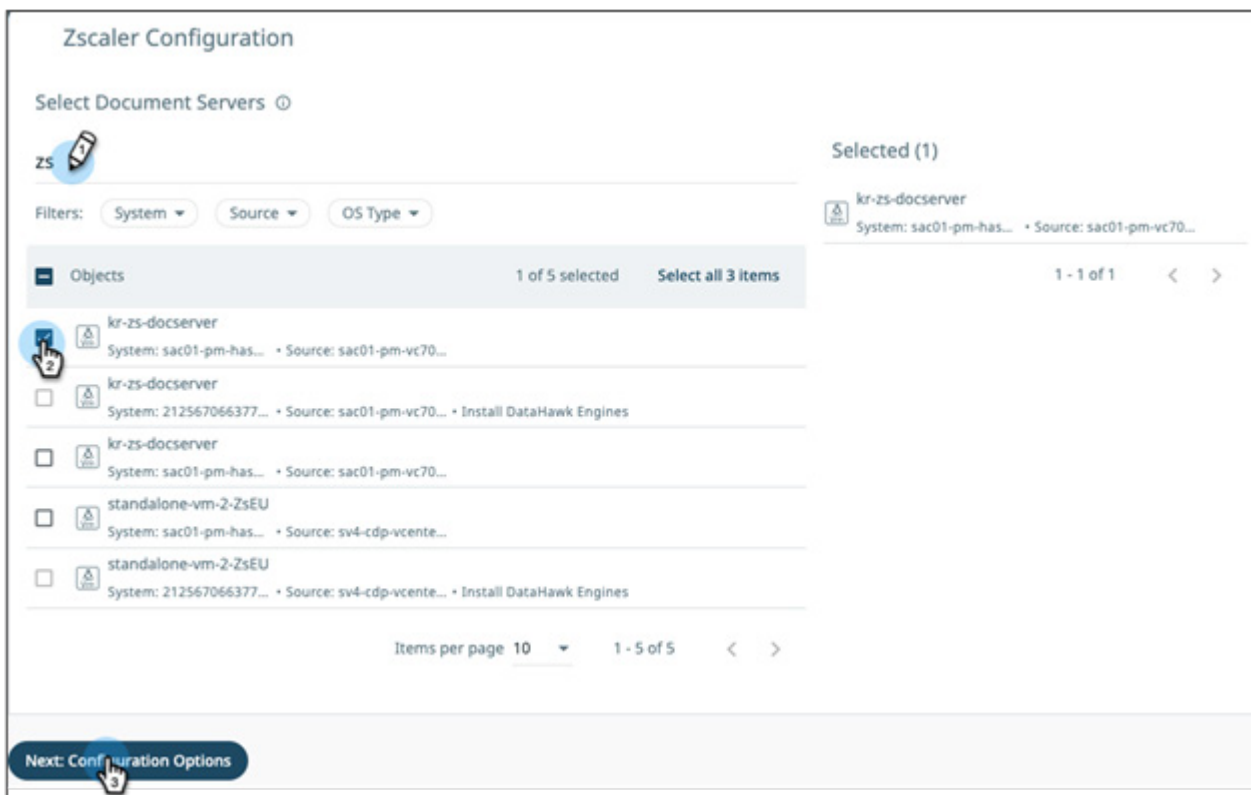


Figure 21. Cohesity Zscaler Configuration

- b. Select one of the following options
- **None:** Select this option if you do not want to filter the sensitive files you want to store on the document server.
 - **Include the following:** Select this option and choose the **File Type** and **Patterns** you want to include to filter the sensitive files. Only the sensitive files that match the File Type and the Patterns are stored on the document server.
 - **Exclude the following:** Select this option and choose the **File Type** and **Patterns** you want to exclude to filter the sensitive files. This excludes the sensitive files that match the File Type and the Patterns from the data classification.
- c. Under **Submission Status**, enable **File Submission: Active** to securely store sensitive files on the document server identified by Cohesity Data Classification.



Sensitive files discovered in the data classification scans are performed after the Zscaler configuration is stored on the document server.

- d. Click **Configure**.

Zscaler Configuration

1 Document Servers VMs 1 Clusters Configured

Files Submission

Submission Filters

None Include the following Exclude the following

File Type

.jpg .png .gif .bmp .tiff

Patterns

Email IPv4 Public IPv4 Address IPv6 Url US SSN IMEI MAC Address Credit Card

Tokens, Keys and Secrets Product Keys Explicit Password Religion (English) US Ethnicity (Wide) Full Name Phone

Date of Birth (DOB) US Address US Full Name US Bank Account Number Age Policy Number

Cleartext Password near Term Email and Cleartext Password US Passport Number US Visa Number

Submission Status

File Submission: Active
Files from only new data classification scans will be submitted. Files from previous scans are not submitted.

Configure Cancel

Figure 22. Cohesity Zscaler Exclusions Configuration

After the Cohesity Data Cloud integration with Zscaler is complete, you can view the status of sensitive files submitted on the document server on the Zscaler Activity page. To access this page, go to **My Integration > Zscaler** and click the **Activity** tab.

Analyzing Zscaler DLP Engine Logs

After the Zscaler DLP Inline policy is enabled, the Zscaler engines are triggered any time a violation is identified.

To visualize the DLP inline policy logs:

1. Go to **Analytics** > **Web Insights**.

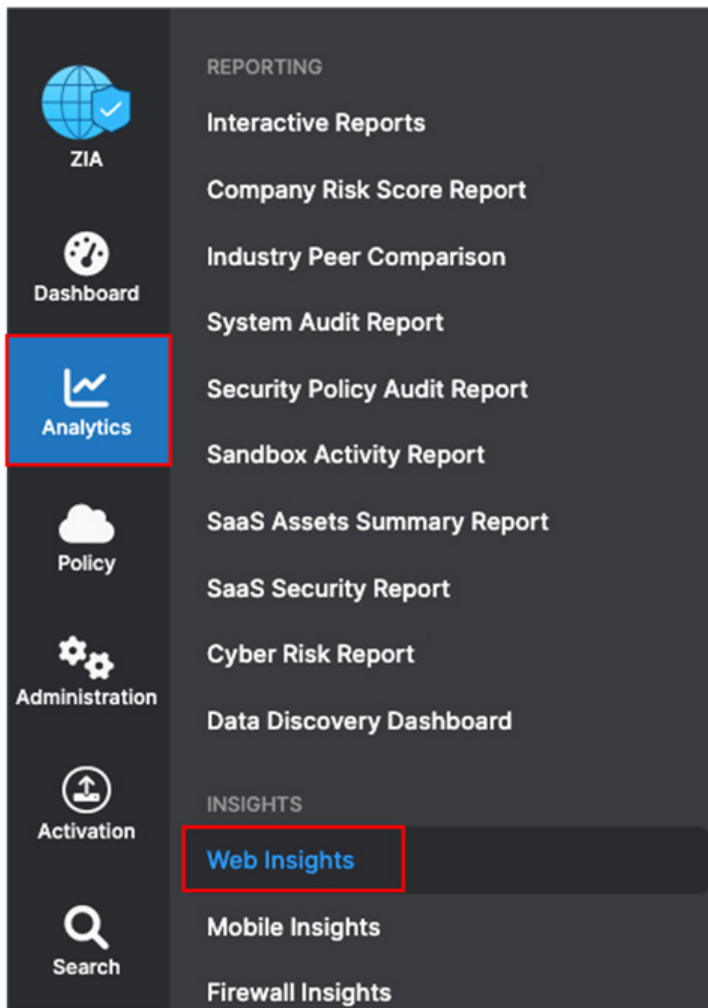


Figure 23. ZIA Web Insights

2. Filter the logs based on the specific DLP engines (i.e., Cohesity-Engine):
 - a. Select the **Timeframe** for which to filter logs.
 - b. Select **Add Filter**.
 - c. Search for **DLP Engine**.
 - d. Select the **DLP Engine** name created in the previous sections (e.g., Cohesity-Engine).
 - e. Select **Add Filter** again.
 - f. Select **Policy Action**.
 - g. Select **Block**.
 - h. Click **Apply Filters**.

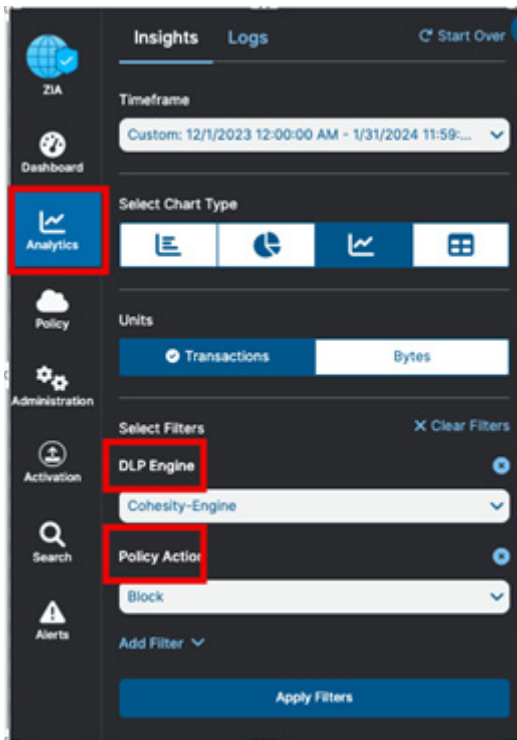


Figure 24. ZIA Log Filter

With the log filters applied, ZIA lists all the logs matching the selected criteria. If one or more violations are detected, it displays several types of information related to the violation such as: User, Policy Action, URL Category, and Cloud Application.

Event Time	User	Policy Action	Location	URL Category	DLP Dictionaries	DLP Engine
Thursday, December 21, 2023 9:40:44 PM	pocadmin1@dev.c...	Violates Compliance Category	Road Warrior	Professional Services	Cohesity IDM (100)	Cohesity-Engine
Thursday, December 21, 2023 9:42:12 PM	pocadmin1@dev.c...	Violates Compliance Category	Road Warrior	Professional Services	Cohesity IDM (100)	Cohesity-Engine
Thursday, December 21, 2023 10:03:03 PM	pocadmin1@dev.c...	Violates Compliance Category	Road Warrior	Professional Services	Cohesity IDM (100)	Cohesity-Engine
Thursday, December 21, 2023 10:03:23 PM	pocadmin1@dev.c...	Violates Compliance Category	Road Warrior	Professional Services	Cohesity IDM (100)	Cohesity-Engine
Tuesday, January 02, 2024 12:47:25 AM	pocadmin1@dev.c...	Violates Compliance Category	Road Warrior	Professional Services	Cohesity IDM (100)	Cohesity-Engine
Tuesday, January 02, 2024 12:47:38 AM	pocadmin1@dev.c...	Violates Compliance Category	Road Warrior	Professional Services	Cohesity IDM (100)	Cohesity-Engine

Figure 25. ZIA Data Protection Insight Logs

In addition to User, Policy Action, URL Category, and Cloud Application, Zscaler also lists the DLP Engine (name) and DLP Dictionary that triggered the policy action.

Zscaler Data Discovery Report

Zscaler offers a detailed view of potential DLP violations via its Data Discovery Report page.

To visualize the Data Discovery Report:

1. Go to **Analytics > Data Discovery Report**. On this page, you can apply several filters to obtain high-level visibility of the potential DLP violations that have occurred over time, up to 90 days.




Figure 26. Data Discovery Report

2. In the **Files in Top Eight DLP Engines** widget, find the engine to see its violation details (i.e., Cohesity-Engine).

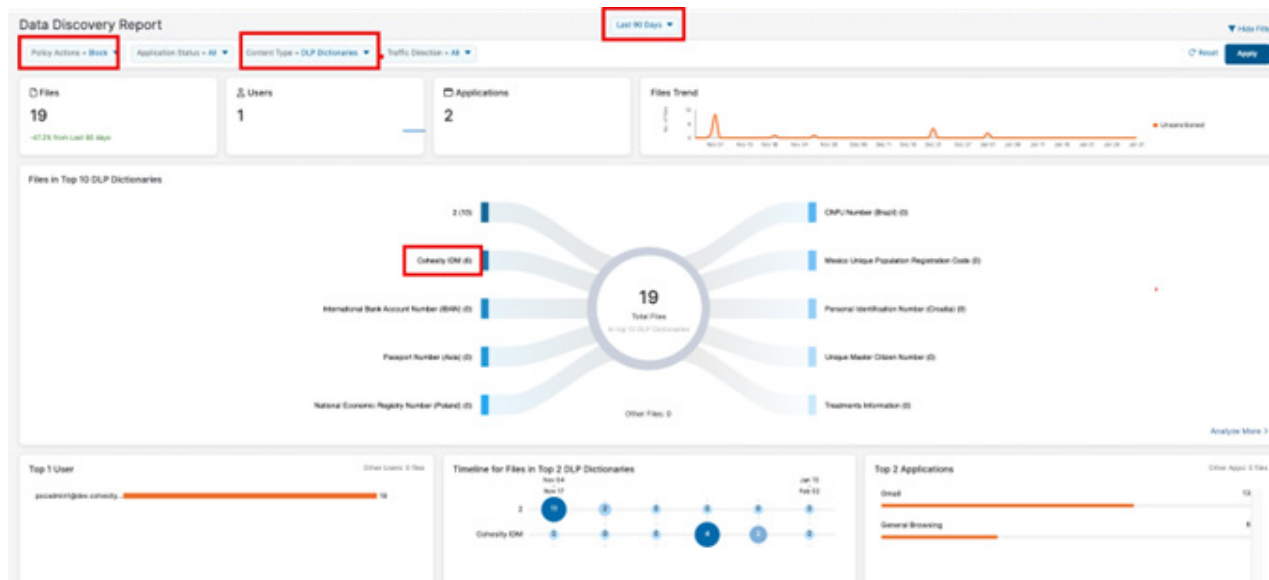


Figure 27. ZIA DLP Data Discovery Report

3. On the **Data Discovery Report** page, visualize the potential violations originated from files restored by Cohesity based on a graphical timeline.
4. In the **Files in Top Eight DLP Engines** widget, select the engine name that you want to see the details.
5. The **Data Discovery Details** displays several associations.
 - a. **Content Type:** DLP Engine triggered.
 - b. **Applications:** Cloud Applications that triggered the DLP Engine.
 - c. **Users:** The users who have potentially violated a DLP inline policy.

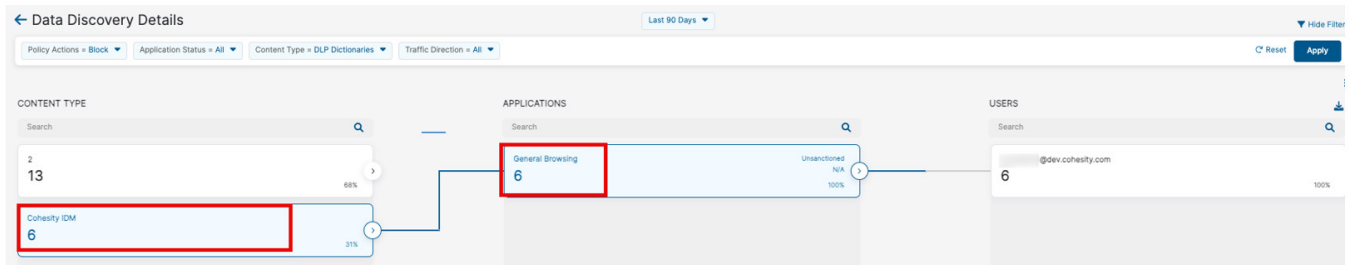


Figure 28. Data Discovery Details

Best Practice for Securing Your Drop Server

The following section helps you set up the security functionality within your Drop Server.

Zscaler recommends using a supported and licensed OS that can run current updates for security patching. Since this server handles sensitive information, control and monitor access to it with the least privilege necessary. Ensure that everyone who has access to it is authorized and authenticated.

The examples used are not exhaustive. Zscaler doesn't recommend consulting your internal information security for recommendations, best practices, and considerations for your environment.

Microsoft Windows 11 Example

Securing a Windows 11 virtual appliance involves a variety of measures to protect against security threats. The following are step-by-step instructions to help you harden your system and limit the local firewall.

1. Ensure that you have a clean installation of Windows 11. You can obtain a Windows 11 installation ISO from Microsoft or a preconfigured virtual appliance image. After you have the image or ISO, create a new virtual machine on your cluster and install Windows 11.
2. After Windows 11 is installed, ensure that automatic updates are enabled. This ensures your system receives security patches and updates as they become available. Go to **Settings > Update & Security > Windows Update** and enable the **Automatically download updates** option.
3. Change the Default Administrator Account. By default, Windows 11 creates an Administrator account with a well-known username. Attackers can easily target this. To change the default Administrator account, create a new user account with administrative privileges and delete the original Administrator account.
4. Windows 11 includes built-in antivirus software called Windows Defender. Ensure that Windows Defender is enabled to protect against malware and other security threats. Go to **Settings > Update & Security > Windows Security** and enable the **Real-time protection** option.
5. Windows 11 also includes a built-in firewall that can help protect your system from unauthorized access. To enable the Windows Firewall, go to **Settings > Update & Security > Windows Security > Firewall & network protection** and enable the **Windows Firewall** option.

6. After the Windows Firewall is enabled, configure it to allow only necessary traffic to and from the virtual appliance. To do this, go to **Settings > Update & Security > Windows Security > Firewall & Network Protection > Advanced** settings. From there, you can create inbound and outbound rules to allow or block specific types of traffic.
7. In addition to Windows Defender, consider installing additional anti-malware software to provide an extra layer of protection. There are many third-party options available, such as Malwarebytes, Norton, or McAfee. Ensure that the software is updated regularly and that scans are run regularly.
8. Add any tooling recommended by your internal information security team for monitoring activity on the host and the ability to respond in the event of an attempted attack.

Ubuntu Linux Example

Similarly, to the Windows guide, the following are step-by-step instructions to help you harden your Ubuntu virtual appliance. If you have a prebuilt image or virtual instance, the organization already has a secured configuration.

1. Before starting the hardening process, ensure that your Ubuntu virtual appliance is up-to-date. Open a terminal and run the following commands:

```
sudo apt update
```

```
sudo apt upgrade
```

This updates the system with the latest security patches and bug fixes.

2. Identify and disable any unnecessary services running on your Ubuntu virtual appliance. Use the following command to list the active services:

```
systemctl list-units --type=service --state=running
```

Disable any services that are not required by running the command:

```
sudo systemctl stop <service-name>
```

```
sudo systemctl disable <service-name>
```

3. Enable and configure a firewall to control incoming and outgoing network traffic to the host. Ubuntu uses ufw (Uncomplicated Firewall) by default. Open a terminal and run the following commands:

```
sudo ufw enable
```

```
sudo ufw allow ssh
```

```
sudo ufw default deny incoming
```

```
sudo ufw default allow outgoing
```

This enables the firewall and sets the default policies for incoming and outgoing traffic.

4. It's essential to secure SSH access to the virtual appliance. Edit the SSH configuration file by running the following command:

```
sudo nano /etc/ssh/sshd_config
```

Inside the file, make the following changes:

- Set PermitRootLogin to no to prevent remote root login.
- Change the default SSH port (optional but recommended).
- Set PasswordAuthentication to no and use SSH key-based authentication.
- Save the changes and restart the SSH service:

```
sudo systemctl restart ssh
```

Appendix A: Requesting Zscaler Support

If you need Zscaler Support to provision certain services or to help troubleshoot configuration and service issues, it is available 24/7/365.

To contact Zscaler Support:

1. Go to **Administration > Settings > Company Profile**.




Figure 29. Collecting details to open support case with Zscaler TAC

2. Copy your **Company ID**.

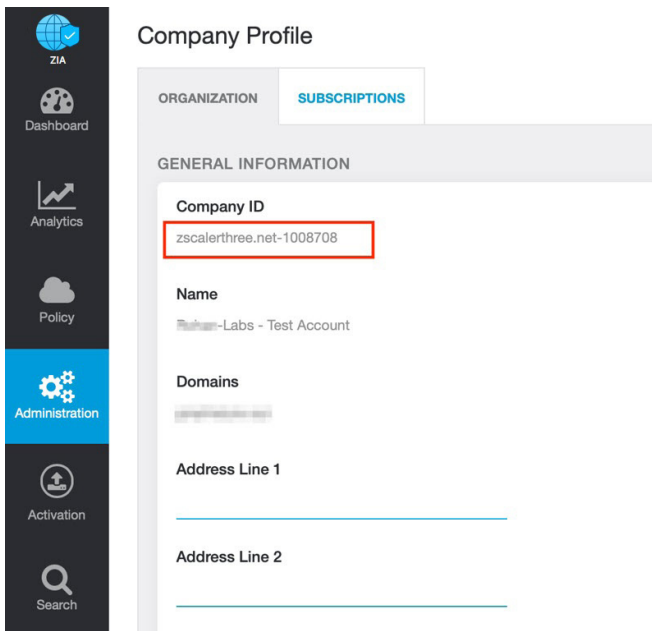


Figure 30. Company ID

3. With your company ID information, you can open a support ticket. Go to **Dashboard > Support > Submit a Ticket**.




Figure 31. Submit a ticket