SASE Deployment

Made Simple with Cato

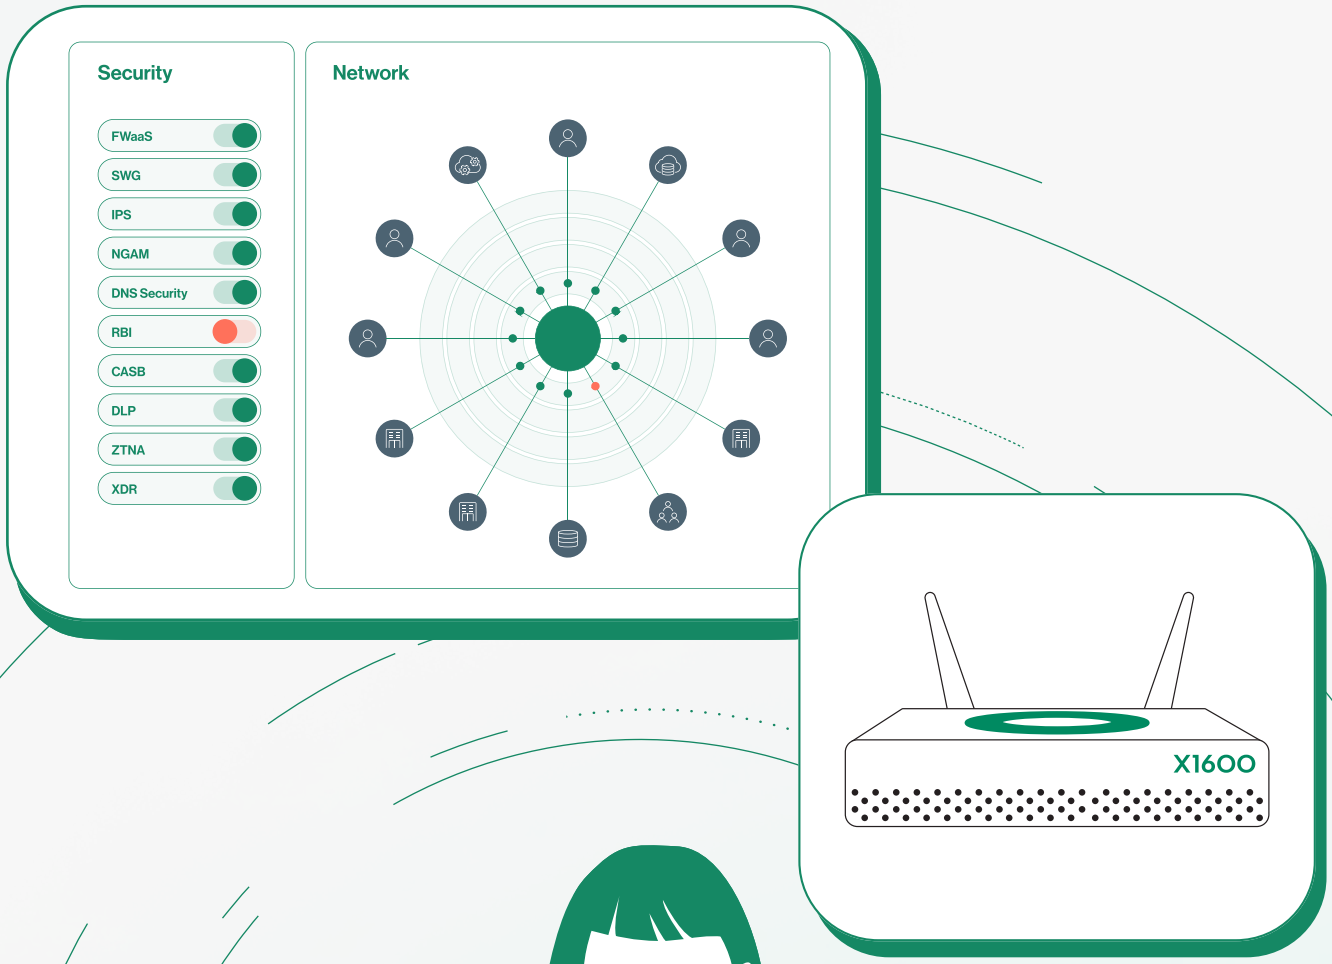

CATO NETWORKS

SASE Deployment

Made Simple with Cato

IT is under increasing pressure to own and run a complex, fragmented security infrastructure. Each incremental requirement pressures IT resources to select, deploy, integrate, and manage a new point solution while sustaining an ever-growing technology stack's posture, performance, global reach, and resiliency - all while providing the technical foundations for the business to grow and be competitive. The Cato SASE platform was developed to address that very pain. Cato's cloud-native full security stack and global private backbone will enable your organization to be more flexible and agile going forward.
This checklist will guide you along the general steps of deployment and help you avoid some common missteps that can cause delays during deployment.

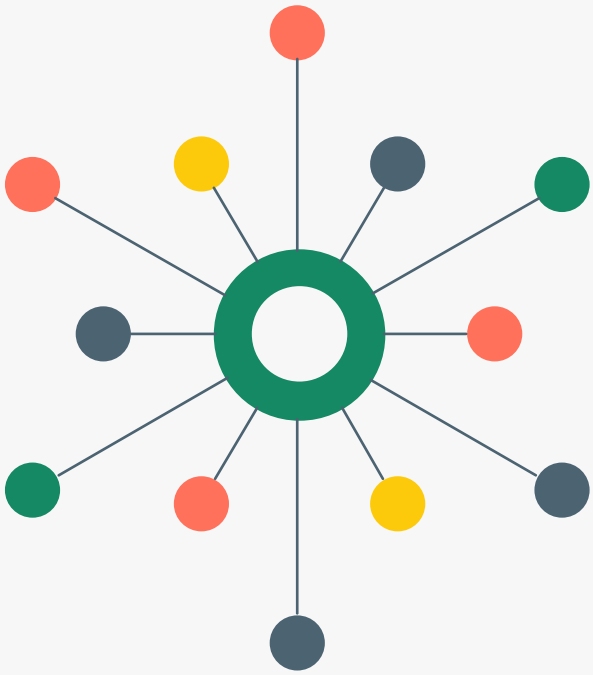

Before You Begin

Deploying any new security or networking technology can be challenging and stressful, especially if the technologies are from different vendors and don't work together toward your goal of SASE adoption. Cato's SASE platform is the opposite of that approach, where the network and security controls are converged onto a single platform that is cloud-native and easy to deploy and manage.

This paper is divided into five specific deployment stages for Cato SASE. This checklist covers the most commonly used SASE features at publication.

Stage 1: **Planning**

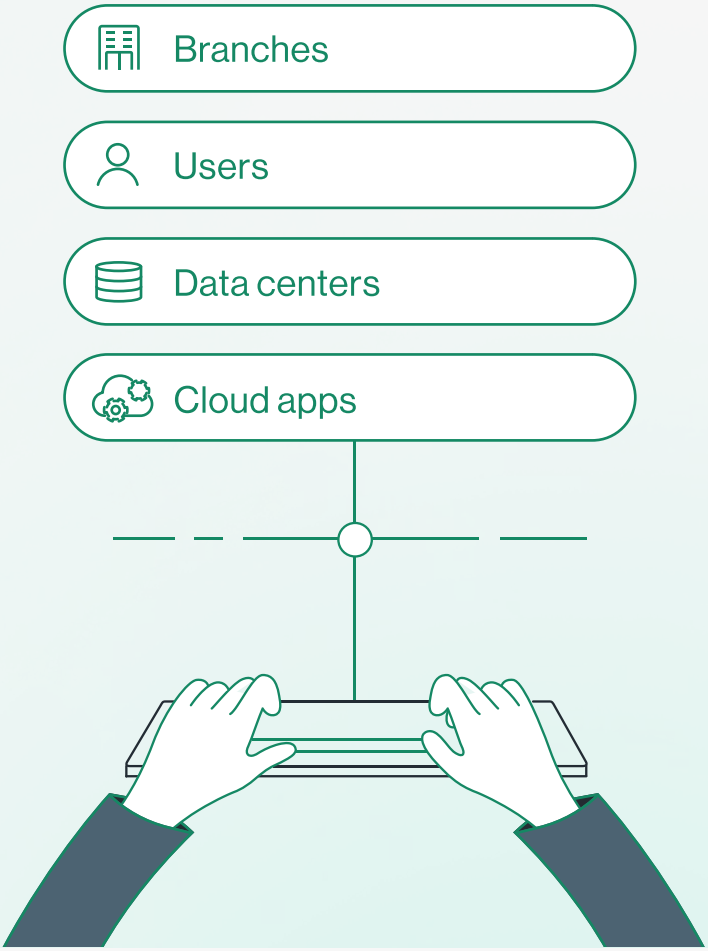

Pre-deployment

The key deliverable in this planning stage is the SASE implementation plan, which will become the deployment roadmap for your project. Identifying and involving key stakeholders early in the SASE deployment project will drive alignment, tighter collaboration, and efficient execution with fewer surprises later. These are usually representatives from departments such as information security, networking/infrastructure, Cato Professional Services, application owners/developers, partners/MSPs/VOIP vendors, and other impacted business units.

As a part of the information-gathering process, stakeholders must collaborate on the project timeline and develop key project milestones. After the kickoff meeting, other group meetings and briefings, such as architecture and network review, application review, and security auditing, should occur in parallel with other implementation plan activities. These meetings aim to identify the information and assign responsible parties to the sections of the SASE implementation plan.

Pre-deployment Key Steps:

Identify the stakeholders for your Cato SASE deployment.

- ☑ Set milestones, tasks, and critical dates for the project with all stakeholders.
- ☑ Schedule and complete a project kickoff meeting with all stakeholders.
- ☑ Engage with Cato Professional Services or your current partner or service provider organization.
- ☑ Start gathering the materials from stakeholders that will assist with planning the deployment.

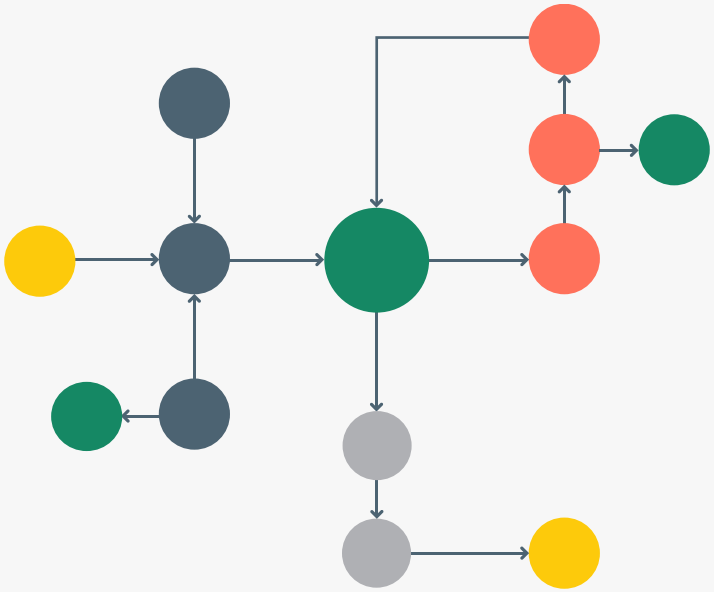

Creating Final SASE Deployment Plan

Sites, Users, Security Deployment

As the team gathers the necessary details for SASE deployment, your site deployment order will be determined based on performance, security, or other requirements in the deployment plan. Site deployment includes deciding site onboarding order, which connectivity method will be used for sites, and bandwidth licensing.

The key idea is to identify where Cato will be deployed in the network and determine the deployment ordering for the sites, data centers, and cloud environments. Having an up-to-date network diagram and current subnet assignments for your sites will be very helpful when starting to plan the deployment of your sockets to their respective sites and determine the level of effort for each site's deployment.

Finally, you will want to select which features of the Cato service you want to work with first. Several features and functions are available in the Cato SASE cloud, and you will want to start with the most important to you and your organization and then expand your Cato footprint as the project progresses.

Site Planning Key Steps:

- ☑ Determine site deployment order and precedence – work with your architecture team to build a list of sites and when you want them to be deployed into the Cato SASE Cloud.
- ☑ Submit change requests to your ISPs and MSPs for necessary IP ranges and circuit changes.
- ☑ Select which capabilities to enable first and which user groups will utilize those features.
- ☑ Complete and distribute your implementation plan to all stakeholders.

Stage 2: **Deployment**

Sites and Sockets

Cato sockets have arrived, and you're ready to start implementation. At this point in your project, you should have the network architecture to which your sockets will connect. All previous planning and preparations will pay off during this phase. The Cato socket deployment guide will provide step-by-step instructions for installation.

Based on the sites you are deploying, build and configure those sites in the Cato Management Application before connecting the corresponding socket at each location. This will ensure that security features and policies are immediately enforced as those sockets come online and sync with the Cato cloud.

Site and Socket Deployment Key Steps:

- Create and configure your sites in the Cato Management Application before connecting new sockets.
- Connect sockets to the networks and power adapters, allowing the sockets to sync with the Cato Cloud. Assign the sockets to the pre-defined sites in the Cato Management Application.

ZTNA Users

Most organizations will have remote employees, contractors, or others needing secure access to internal assets or applications. The Cato Client provides zero-trust network access (ZTNA) to those resources via the Cato global private backbone. When providing access to these remote users, you have the option to connect the Cato SASE Cloud to the identity provider (IdP) that your organization uses and have the IdP perform the user authentication. Cato supports using both SCIM and LDAP sync to import user and group membership data.

Next, you will need to configure your access control rules and policies. This can include things such as:

- Device Posture
- Client Connectivity and Access
- Always-On
- Rollout
- IP Allocation
- Proxy Configuration
- DNS Settings
- Trusted Networks

Deploying the Cato Client to your users is also simple, as there are several methods of deployment available. Users can access the self-service portal to download and install their Cato client or administrators can also use existing MDM or domain tools to deploy the client. The Cato SASE platform includes built-in update rollout functionality in the console, simplifying the management and maintenance of the deployed clients. Depending on the requirements of your organization, you can also utilize device certificates to authenticate devices to the Cato Cloud, which requires the certificate to be uploaded to the Cato Management Application to be used that way.

Another key step is to determine how your team will do testing and implement feedback from the test users. It is best practice to first deploy the Cato client to a small test group and then fine-tune the policies with that group before deploying to the broader user groups.

User Deployment Key Steps:

- ☑ Connect your IdP via the Cato Management Application for user authentication.
- ☑ Deploy Cato Client to test groups with instructions to provide feedback.
- ☑ Adjust policies and profiles based on test user feedback.
- ☑ Deploy Cato Client to broader user groups based on timeline and rollout plans determined during the project planning stage.

Stage 3: **Network**

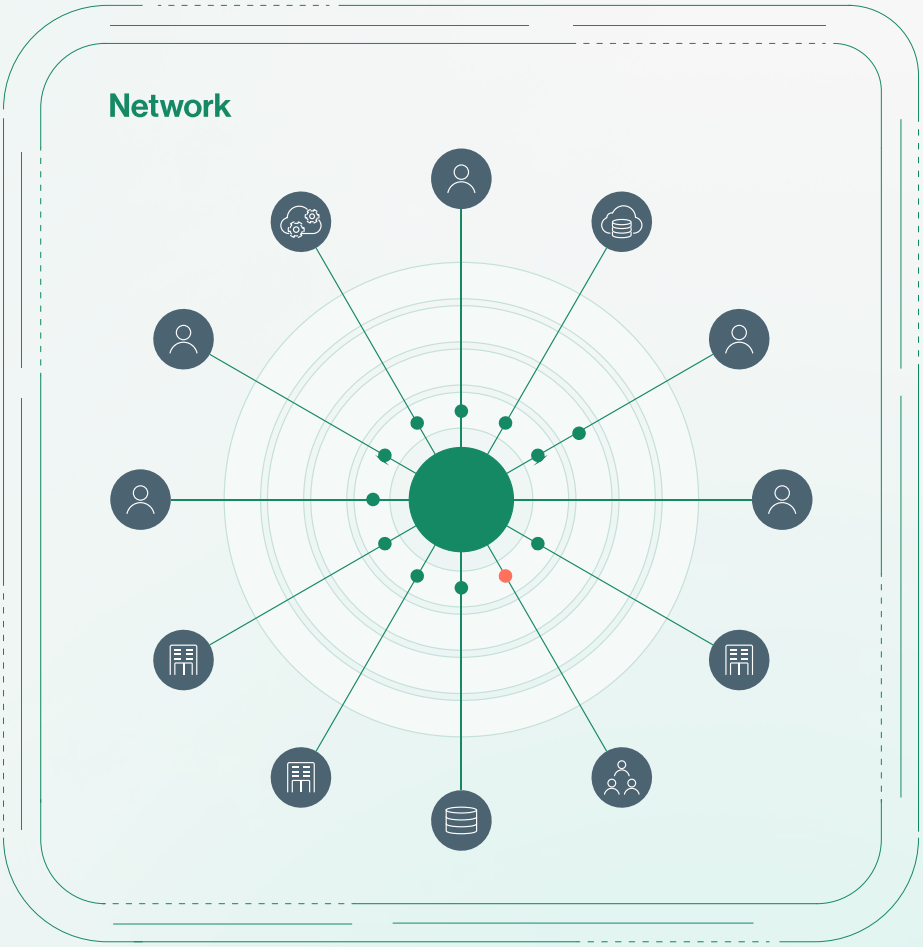

Networking in Cato SASE

After deploying your sites, which only require minimal network information, several options exist to fine-tune and expand your networking configuration in Cato. Some of these configurations include the following:

DHCP

The Cato Cloud has DHCP capabilities; after you define the DHCP ranges in the Cato Management Application, the Cato DHCP service in the PoPs assigns IP addresses to clients. You can also choose to use Cato as your organization's DHCP relay agent for a local DHCP server. All expected functions of a DHCP service are included: range, lease time, and relay settings.

DNS

By default, the Cato Cloud provides a DNS service for your account and acts as your DNS server. You can also use the Cato Management Application to configure Cato to resolve private DNS records. As a best practice, Cato recommends configuring two different DNS servers to offer the best security, performance, and redundancy.

Bandwidth Management

Bandwidth Management profiles define how you measure and control network traffic based on profiles to manage and prioritize upstream/downstream bandwidth. The profiles are determined by using a fixed speed or a percentage of the total bandwidth. By default, these profiles are applied at the account level rather than per site. Each new account includes predefined profiles based on common needs, which can be edited with great granularity. Additional profiles can be added and configured for greater granularity. Cato also supports setting networking rules, which are used to ensure that your critical applications perform as expected, no matter where you are.

Link Health

The solution provides monitoring tools to help admins understand as soon as something bad is happening on the network. They can set alerts for both blackouts and brownouts, and have them sent at a defined urgency to a customized mailing list of relevant team members.

Cato Networking Key Steps

- ☑ Set DHCP and DNS to meet your organization's requirements.
- ☑ Review and adjust bandwidth management profiles and network rules to adhere to business applications SLA requirements.
- ☑ Set monitoring alerts based on your organization's requirements.

Stage 4: **Security**

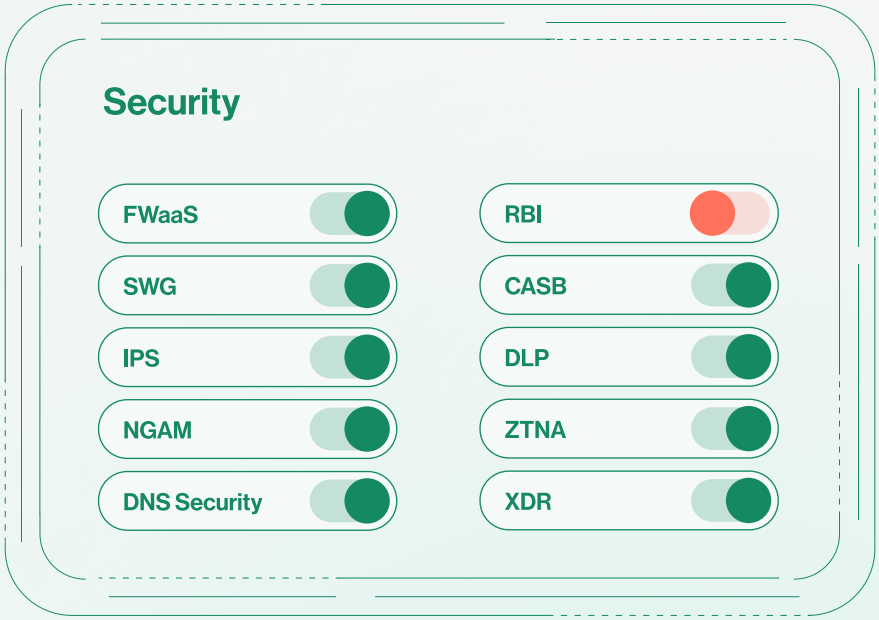

Configuring Security Settings

Cato offers a complete SASE platform, including firewall-as-a-service, secure web gateway, intrusion prevention, next-gen anti-malware, DNS security, remote browser isolation, TLS inspection, and cloud access security broker, to name a few.

Firewall as a Service

Firewall-as-a-service is broken down into two key areas: the Internet firewall and the WAN firewall.

Internet Firewall

The Internet firewall inspects traffic between the enterprise and the Internet and lets you create rules to control this traffic. The Internet firewall uses an ordered ruleset; connections are inspected according to each rule from the first rule. The Internet firewall uses a blocklist approach. This means there is an implicit ANY - ANY rule to allow any traffic and connections not explicitly blocked in the policy.

WAN Firewall

The WAN firewall in the Cato SASE Cloud controls access to objects and entities in your Wide Area Network (WAN). Configure the WAN firewall policy to create a secure access control policy and protect the network. The WAN firewall is part of the Next Gen Firewall (NGFW) integrated into the Cato SASE Cloud and lets you create rules to prevent unauthorized access to the network. The WAN firewall uses a whitelist approach, and there is a default ANY - ANY block rule to drop all connections that are not explicitly allowed in the policy.

IPS

Cato's IPS employs a multi-layered approach to security, which includes reputation analysis to safeguard against communication with compromised or malicious entities, both inbound and outbound. Additionally, it guards against known vulnerabilities by promptly integrating new CVEs into its protection protocols. The service also features anti-bot measures, which prevent outbound traffic to Command and Control (C&C) servers through reputation feeds and network behavioral analysis. Furthermore, it utilizes network behavioral analysis to detect and thwart inbound and outbound network scans. Protocol validation ensures the legitimacy of packets, thereby minimizing the risk of exploits through anomalous traffic. Moreover, custom geo-restriction policies can be enforced to block traffic from specific countries, enhancing overall security measures.

Next-Generation Anti Malware

Cato implements the SentinelOne Next-Gen Anti-Malware engine to provide a second layer of threat protection. This engine uses an AI model that detects threats in portable executable files, PDFs, and Microsoft Office documents. The AI model is developed by extracting features from millions of malware samples in the malware repository. Then, Supervised Machine Learning (SML) is used to identify and correlate different characteristics of benign and malicious files. The engine then uses this model to identify similar features in unknown files that are classified as malicious, suspicious, or benign.

TLS Inspection

Cato Networks provides Transport Layer Security (TLS) for internal and external traffic inspection. Cato inspects TLS on all ports, if the protocol is a valid TLS version and has the correct header data. When TLS inspection is enabled, the Cato PoPs decrypt the HTTPS traffic and inspects it for malicious content. Cato recommends using TLS inspection for Cato's threat protection services, such as Intrusion Prevention System (IPS) and Anti-Malware. As part of implementing TLS Inspection, you will need to install either the Cato certificate as a root certificate on the end-user hosts and devices or upload a certificate that will be used for TLS inspection to the Cato Management Application.

Security Key Steps

Based on your environment security requirements, read through some of the following articles for deploying the Cato security features:

- ☑ **Firewall as a Service** – this includes both the WAN and Internet firewall policies
- ☑ **Intrusion Protection System (IPS)**
- ☑ **Next-gen Anti-Malware Service**
- ☑ **TLS Inspection** – this is a very important feature since it underpins the CASB and DLP features as well as application visibility for analytics.

Stage 5: **Visibility and Integration**

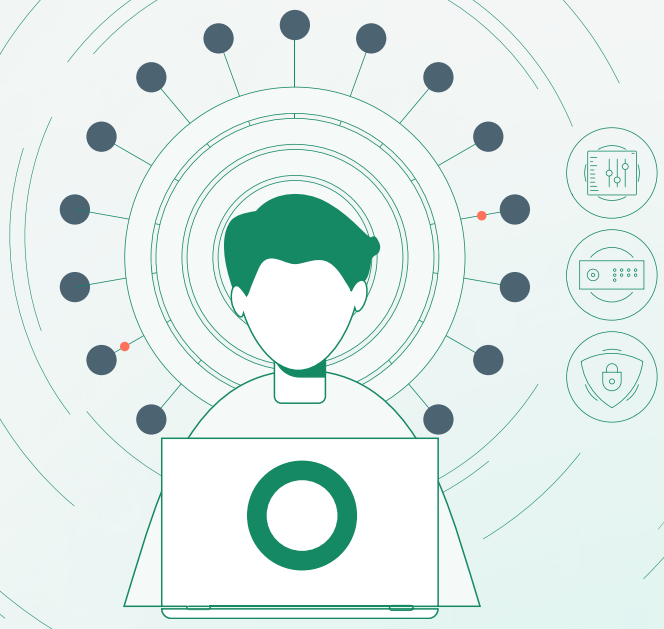

Events, Logs, and Stories

After deploying your Cato SASE solution, the platform will gather details and events in your environment. Cato will process these events into what we call stories that begin to paint the picture of what is happening in your SASE-protected network.

Start looking at the events and dashboards in the Monitoring tab of your Cato Management Application tenant. These dashboards provide you with rich, interactive, and actionable data that you and your team can use and rely on going forward. Cato's dashboards provide a simple way to view your events and stories in a single location.

Review your traffic and application data in the Cato Management Application on the following dashboards:

- ☑ App Analytics
- ☑ Stories Dashboard
- ☑ SDP Users Overview
- ☑ Security Threats
- ☑ Cloud Apps

While using the Cato XDR, you will also want to review how to use the information in your Stories Workbench, so it would be a good idea to review the documentation here.

Wrapping Up Your Deployment

It's time to take your network security and performance to the next level. Now that your infrastructure is fortified with Cato SASE capabilities, it's crucial to regularly monitor and optimize its performance to ensure maximum efficiency and security. Stay proactive by conducting regular assessments, monitoring traffic patterns, and addressing any emerging threats promptly. Remember, SASE is not a one-time implementation; it's an ongoing journey towards a more resilient and agile network infrastructure. Get started today and unleash the full potential of your Cato SASE deployment!

This checklist is not meant to be exhaustive but rather to give you and your organization a good starting point in the planning phase of adopting your Cato Networks SASE capabilities. There are also steps you will need to complete in different orders, but the idea behind this checklist is to point you in the right direction and give you links to documentation that will assist you in adopting your Cato Networks SASE platform.

Appendix

Events, Logs, and Stories

Cato Professional Services Customer Document >
Cato Hardware Socket Deployment Guides >
Cato Virtual Socket Deployment Guides >

Other Cato SASE Features

Networking

IP Ranges

The Global IP Range entity is a global object in the Cato Management Application that you define and then use in rules across multiple policies. For example, you can use the same range for servers in WAN Firewall, Network Rules, and other policies. If, at some point, the range changes, you only need to update the object once, and all the policies are automatically updated.

You can also use custom IP ranges when the IP range is only used in a specific rule. Cato also supports floating ranges, which are only applied to traffic routed via BGP and when the advertised route matches the floating range.

Security

SaaS Security API

Cato's SaaS Security API provides out-of-band visibility and control for sanctioned cloud apps. Other security features (such as DLP) can only control and monitor traffic that goes over the Cato Cloud. SaaS Security API also allows monitoring and reacting to traffic from remote users that connect directly to the cloud apps. SaaS Security API compliments Cato's inline CASB and DLP solutions to provide the best security coverage.

About Cato Networks

Cato Networks is the leader in SASE, delivering enterprise security and network access in a single cloud platform. With Cato, organizations replace costly and rigid legacy infrastructure with an open and modular SASE architecture based on SD-WAN, a purpose-built global cloud network, and an embedded cloud-native security stack.

Cato SASE Cloud Platform

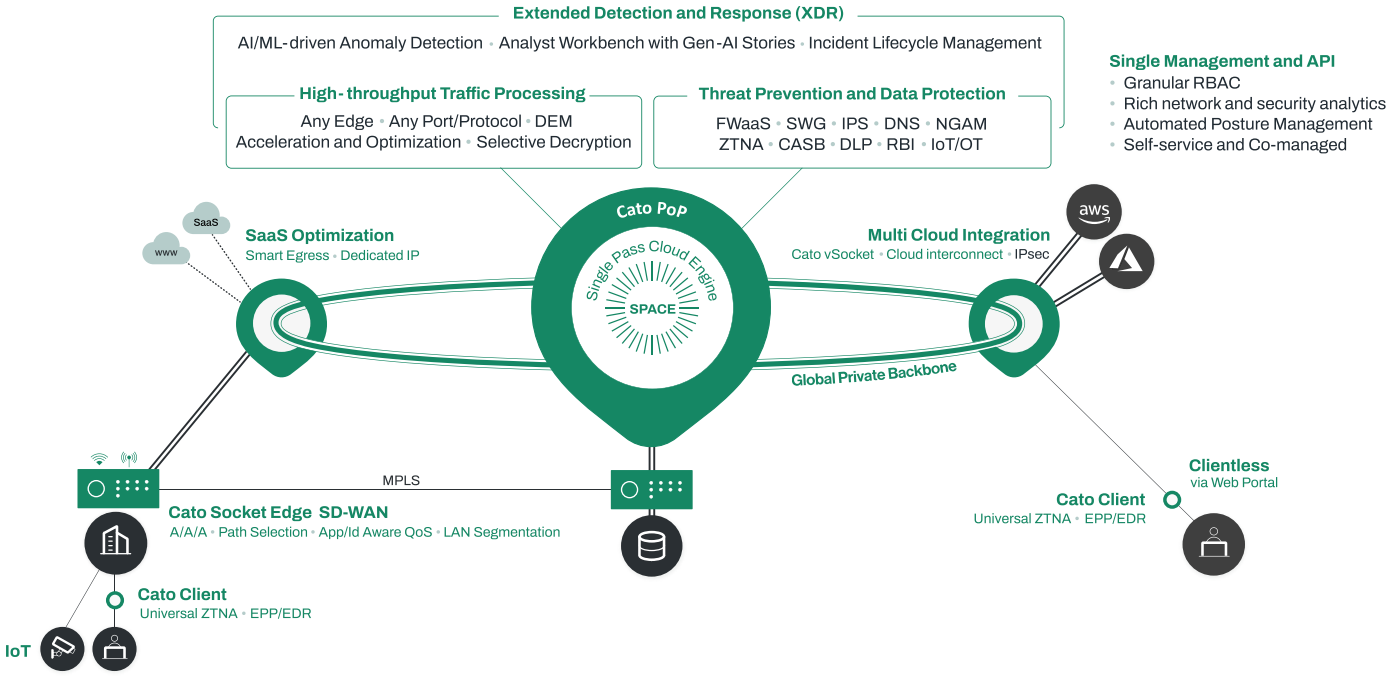

Cato. WE ARE SASE.

Cato SASE Cloud Platform

Connect

Cloud Network

Cloud On-Ramps

Protect

Network Security

Endpoint Security

Detect

Incident Life Cycle Management

Run

Unified Management and API

Use Cases

Network Transformation

MPLS to SD-WAN Migration

Global Access Optimization

Hybrid Cloud & Multi-Cloud

Business Transformation

Vendor Consolidation

Spend Optimization

M&A and Geo Expansion

Security Transformation

Secure Hybrid Work

Secure Direct Internet Access

Secure Application and Data Access

Incident Detection and Response

Contact Us