iboss®

carahsoft.

E-BOOK

Implement Zero Trust as Defined by NIST 800-207

iboss

Implement Zero Trust as Defined by NIST 800-207

E-Book

Thank you for downloading this iboss E-Book. Carahsoft serves as iboss' Master Government Aggregator® making the company's cyber solutions and services available to the Public Sector through Carahsoft's network of reseller partners, Solutions for Enterprise-Wide Procurement (SEWP) V, Information Technology Enterprise Solutions – Software 2 (ITES-SW2), NASPO ValuePoint Cooperative Purchasing Contract, and other contract vehicles.

To learn how to take the next step toward acquiring iboss' solutions, please check out the following resources and information:

For additional resources:
carah.io/ibossResources

For additional iboss solutions:
carah.io/ibossSolutions

To set up a meeting:
iboss@Carahsoft.com
(703)-889-9710

For upcoming events:
carah.io/ibossEvents

For additional iboss NetBackup solutions:
carah.io/ibossNetBackup

To purchase, check out the contract vehicles available for procurement:
carah.io/ibossContracts

For more information, contact Carahsoft or our reseller partners:
iboss@carahsoft.com | (703)-889-9710

E-BOOK

Implement Zero Trust as Defined by NIST 800-207

iboss®

Table of Contents

Overview

NIST Special Publication 800-207 provides a clear and distinct definition of a Zero Trust Architecture that can be used to transform an enterprise to a Zero Trust model. The iboss Zero Trust Edge is a direct implementation against the NIST 800-207 Zero Trust Architecture principles and guidelines to ensure the service can be used to implement Zero Trust according to the NIST 800-207 publication. This guide will specifically walk through each section of the publication providing guidance on the principles described and the method by which to implement those principles using the iboss Zero Trust Edge service.

Defining Zero Trust

NIST 800-207 begins by detailing the basics of Zero Trust (Section 2) which are critically important as everything else is built on this foundation. NIST 800-207 offers the following definitions:

> **"Zero trust (ZT) provides a collection of concepts and ideas designed to minimize uncertainty in enforcing accurate, least privilege per-request access decisions in information systems and services in the face of a network viewed as compromised. Zero trust architecture (ZTA) is an enterprise's cybersecurity plan that utilizes zero trust concepts and encompasses component relationships, workflow planning, and access policies. Therefore, a zero trust enterprise is the network infrastructure (physical and virtual) and operational policies that are in place for an enterprise as a product of a zero trust architecture plan." (NIST-800-207-SP, 2020, p. 4)**

From the definition, the term "Zero Trust" (ZT) is the collection of concepts and ideas related to "enforcing accurate, least privilege **per-request access decisions** in information systems and services," while "Zero Trust Architecture" (ZTA) is the plan that utilizes those concepts. The phrase "per-request access decisions" is highlighted to re-enforce that implementing Zero Trust goes far beyond authentication and includes the ability to analyze and make decisions on every request, after authentication occurs.

After a user authenticates against an Identity Provider (IdP), the requests to the protected resource that follow are not inspected by the IdP. The IdP provides a single access decision at login but has no visibility to the continuous requests to the resource that occur afterward. This is a critical concept to understand and it's the difference between "authentication" and "authorization". Authentication focuses on determining who the user is. Authorization is focused on the permission that is granted to access a system resource and is defined by NIST as:

> **"The right or a permission that is granted to a system entity to access a system resource."**
>
> https://csrc.nist.gov/glossary/term/authorization

From the NIST 800-207 definition of Zero Trust, authorization must be given for every request which is not possible within the authentication system which does not have access to every request.

The definition is further summarized in the following sentence within NIST 800-207:

> **"This definition focuses on the crux of the issue, which is the goal to prevent unauthorized access to data and services coupled with making the access control enforcement as granular as possible. That is, authorized and approved subjects (combination of user, application (or service), and device) can access the data to the exclusion of all other subjects (i.e., attackers). To take this one step further, the word "resource" can be substituted for "data" so that ZT and ZTA are about resource access (e.g., printers, compute resources, Internet of Things [IoT] actuators) and not just data access." (NIST-800-207-SP, 2020, p. 4)**

To summarize, Zero Trust, according to the NIST 800-207, is centered around protecting resource access so that only approved subjects are allowed access to the resource while all others are automatically denied.

Zero Trust in Correlation to the NIST Risk Management Framework (RMF)

The NIST Risk Management Framework is a cybersecurity framework that "provides a process that integrates security, privacy, and cyber supply chain risk management activities into the system development life cycle" **(https://csrc.nist.gov/projects/risk-management/about-rmf)**.

The NIST 800-207 Zero Trust Architecture SP follows the principles of the NIST RMF and is a set of controls and process that are implemented to manage and mitigate risk. Particularly, Zero Trust focuses on reducing and managing risk related to enterprise-owned resource access. Particularly, Zero Trust focuses on protecting access to all enterprise-owned resources which include data, application and services so that only approved users and devices can access those resources while automatically denying access to everyone else.

The process of implementing Zero Trust will follow the general process flow of the NIST RMF which include preparing to implement Zero Trust, categorizing the resources that need to be protected by Zero Trust, selecting and implementing the Zero Trust controls, assessing that the Zero Trust service is operating as expected, authorizing the Zero Trust service and continuously monitoring the Zero Trust service which ensures that only authorized subjects are interacting with protected resources.

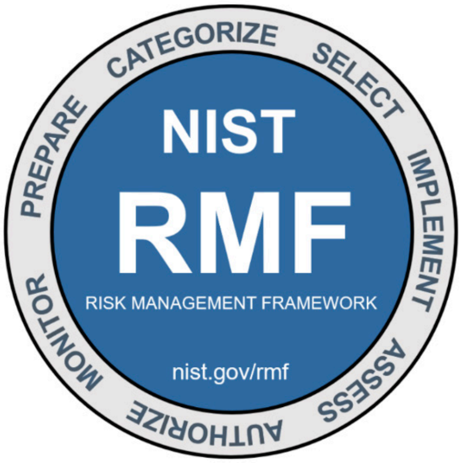

For **more information on each RMF Step,** including **Resources for Implementers** and **Supporting NIST Publications**, select the Step below.

Prepare	Essential activities to prepare the organization to manage security and privacy risks
Categorize	Categorize the system and information processed, stored, and transmitted based on an impact analysis
Select	Select the set of NIST SP 800-53 controls to protect the system based on risk assessment(s)
Implement	Implement the controls and document how controls are deployed
Assess	Assess to determine if the controls are in place, operating as intended, and producing the desired results
Authorize	Senior official makes a risk-based decision to authorize the system (to operate)
Monitor	Continuously monitor control implementation and risks to the system

Figure 1- The NIST Risk Management Framework Process, https://csrc.nist.gov/Projects/risk-management

Preparing to Implement the Zero Trust Architecture

Understanding the core principles of the Zero Trust Architecture are foundational before beginning the Zero Trust transformation. Without this understanding, a proper Zero Trust implementation cannot be achieved.

Understanding the Policy Decision/Enforcement Point and the Implicit Trust Zone

Understanding the Policy Decision/Enforcement Point (PDP/PEP) is central to implementing Zero Trust as it is the heart of the Zero Trust Architecture. The PDP/PEP is the gatekeeper that is responsible for authorizing "per-request" access decisions to protected resources. The most basic diagram of Zero Trust Architecture that can be created is shown by the NIST 800-207 and displayed below:

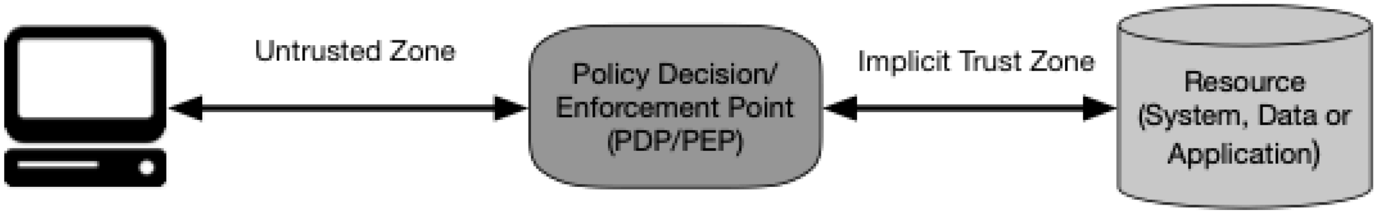
Untrusted Zone
Policy Decision/ Enforcement Point (PDP/PEP)
Implicit Trust Zone
Resource (System, Data or Application)

Figure 1: Zero Trust Access

Figure 2 - NIST 800-207 SP, Page 5

The subject on the left needs access to the protected resource on the right. The Policy Enforcement Point (PEP) which authorizes that connection, sits in between the subject and the protected resource for each and every request to the resource. To the left of the PEP is the Untrusted Zone, which represents completely untrusted users, network, assets, services or anything else that needs access to the resource. To be more specific, a user on the left side of the diagram is a user requesting access from any location or a device (such as a laptop) that is located inside or outside of an enterprise network. To the right of the PEP is what is called the “Implicit Trust Zone” and represents data that should only contain traffic from the PEP or the protected Resource. This is the area after authentication and authorization occurs and allows the PEP to transfer the approved request to the protected resource. The protected resource should be “anchored” to the PEP, so it only accepts traffic from the PEP as that traffic represents connections that have been authenticated and authorized.

The NIST summarizes this with:

> **“The PDP/PEP applies a set of controls so that all traffic beyond the PEP has a common level of trust. The PDP/PEP cannot apply additional policies beyond its location in the flow of traffic. To allow the PDP/PEP to be as specific as possible, the implicit trust zone must be as small as possible.” (NIST-800-207-SP, 2020, p. 5)**

This is one of the most critical aspects of building a robust Zero Trust implementation. If the protected resource (data, application or service) can be accessed by bypassing the PEP, the principles of Zero Trust immediately fall apart. This is because the PEP is responsible for “enforcing accurate, least privilege per-request access decisions in information systems and services” (NIST-800-207-SP, 2020, p. 4) and it cannot enforce “per-request access decisions” if access bypasses the PEP. This is a very nuanced aspect of Zero Trust but it is not trivial. If this foundational principle is not met, anything else built on top of Zero Trust falls apart.

Remember, Zero Trust is designed to protect enterprise-owned resources that include data, applications and services that can be located on-prem or in the cloud as a SaaS service by ensuring that every request is authorized by a PEP. Enforcing this principle on SaaS resources can be challenging because of their multi-tenant nature. With SaaS applications and data, many unrelated enterprises log into the same shared service. However, it’s important to ensure that access to the protected resources is completely restricted unless it’s authorized via the specific Policy Enforcement Points for the specific organization to ensure the request has been authorized by the organization. The iboss Zero Trust Edge makes this possible due to its containerized architecture which allows Access Control Lists to anchor the resource to the Policy Enforcement Point so that no traffic flows into the SaaS application or data without being authorized by the PEPs authorizing requests for the specific enterprise.

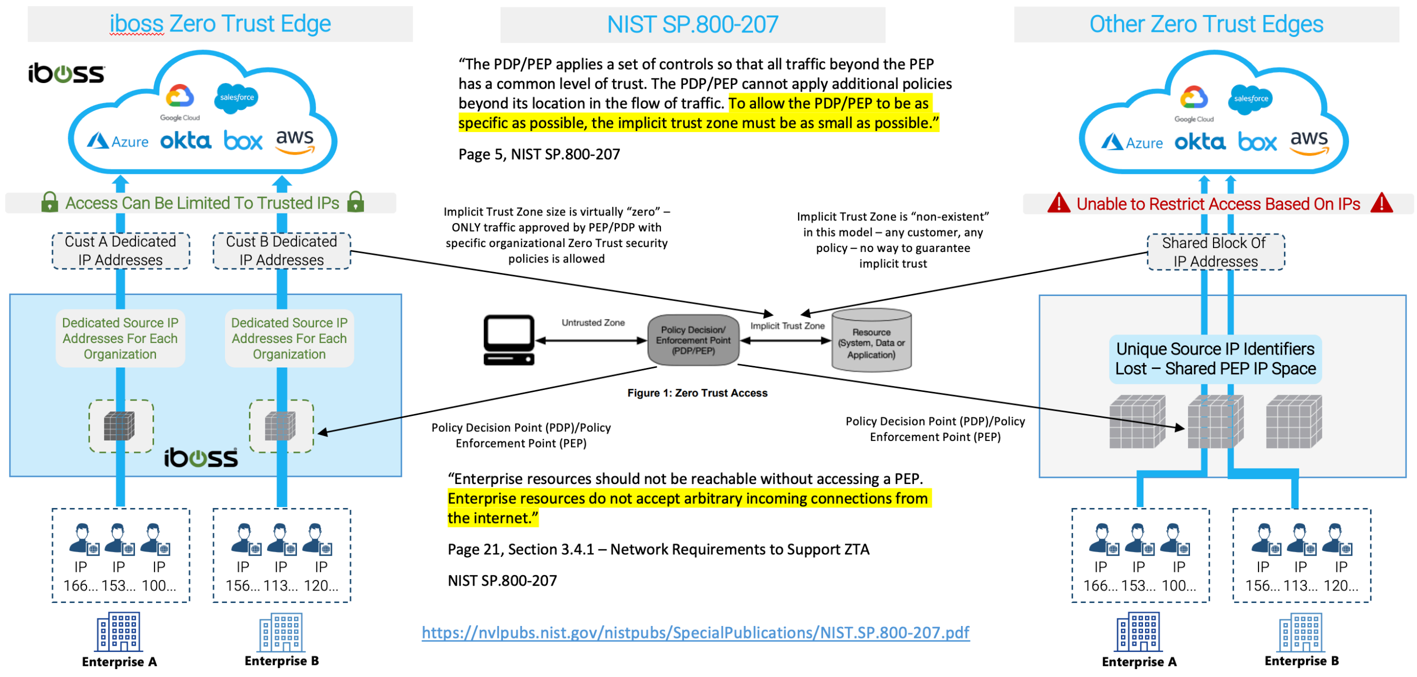
iboss Zero Trust Edge
iboss
Google Cloud
salesforce
Azure
okta
box
aws
Access Can Be Limited To Trusted IPs
Cust A Dedicated IP Addresses
Cust B Dedicated IP Addresses
Dedicated Source IP Addresses For Each Organization
Dedicated Source IP Addresses For Each Organization
iboss
IP IP IP
166... 153... 100...
IP IP IP
156... 113... 120...
Enterprise A
Enterprise B
NIST SP.800-207
Implicit Trust Zone size is virtually "zero" – ONLY traffic approved by PEP/PDP with specific organizational Zero Trust security policies is allowed
Implicit Trust Zone is "non-existent" in this model – any customer, any policy – no way to guarantee implicit trust
Untrusted Zone
Policy Decision/ Enforcement Point (PDP/PEP)
Implicit Trust Zone
Resource (System, Data or Application)
Figure 1: Zero Trust Access
Policy Decision Point (PDP)/Policy Enforcement Point (PEP)
Policy Decision Point (PDP)/Policy Enforcement Point (PEP)
Other Zero Trust Edges
Google Cloud
salesforce
Azure
okta
box
aws
Unable to Restrict Access Based On IPs
Shared Block Of IP Addresses
Unique Source IP Identifiers Lost – Shared PEP IP Space
IP IP IP
166... 153... 100...
IP IP IP
156... 113... 120...
Enterprise A
Enterprise B

"The PDP/PEP applies a set of controls so that all traffic beyond the PEP has a common level of trust. The PDP/PEP cannot apply additional policies beyond its location in the flow of traffic. To allow the PDP/PEP to be as specific as possible, the implicit trust zone must be as small as possible."

Page 5, NIST SP.800-207

"Enterprise resources should not be reachable without accessing a PEP. Enterprise resources do not accept arbitrary incoming connections from the internet."

Page 21, Section 3.4.1 – Network Requirements to Support ZTA

NIST SP.800-207

https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-207.pdf

Figure 3 - The iboss Zero Trust Edge is able to anchor cloud resources to the Policy Enforcement Point (PEP) so that only authorized traffic for the organization is allowed to flow to the protected resource

In the figure above using the iboss Zero Trust Edge on the left, resources reject any connection that does not originate from the Policy Enforcement Point by creating Access Control Lists (ACLs) that only accept traffic from the PEP. The ACLs use the unique and dedicated IP space that only source traffic from one organization which ensures that the policies that authorize the transaction are from that specific organization. With shared Policy Enforcement Point IP Address space, as with the diagram on the right, the resource is accepting arbitrarily authorized connections from any enterprise that has traversed the PEP.

Using an example clarifies this principle. Suppose a highly sensitive document within a SaaS application share drive is made inadvertently public to the world by publishing it with world access permissions (no login required). With an architecture that cannot restrict access to that protected resource to only authorized and approved subjects of that organization, any user that uses the same service will be able to click the link and open the document. With iboss, because every request needs to be authorized through the Policy Enforcement Point which authorized connections only for subjects for the specific organization, all other users that are not approved, including those using the iboss Zero Trust Edge from a different organization, will be automatically rejected. This provides instant zero CASB protection to all sensitive content and applications.

When properly implemented, the protected resources are invisible and inaccessibly to attackers as there is no direct access to any protected resources and can only be achieved by being specifically authorized by a PEP. Choosing the right foundational architecture is critical when implementing the NIST 800-207 architecture and ensuring that the architecture supports preventing direct access to resources is likely one of the most foundational criteria.

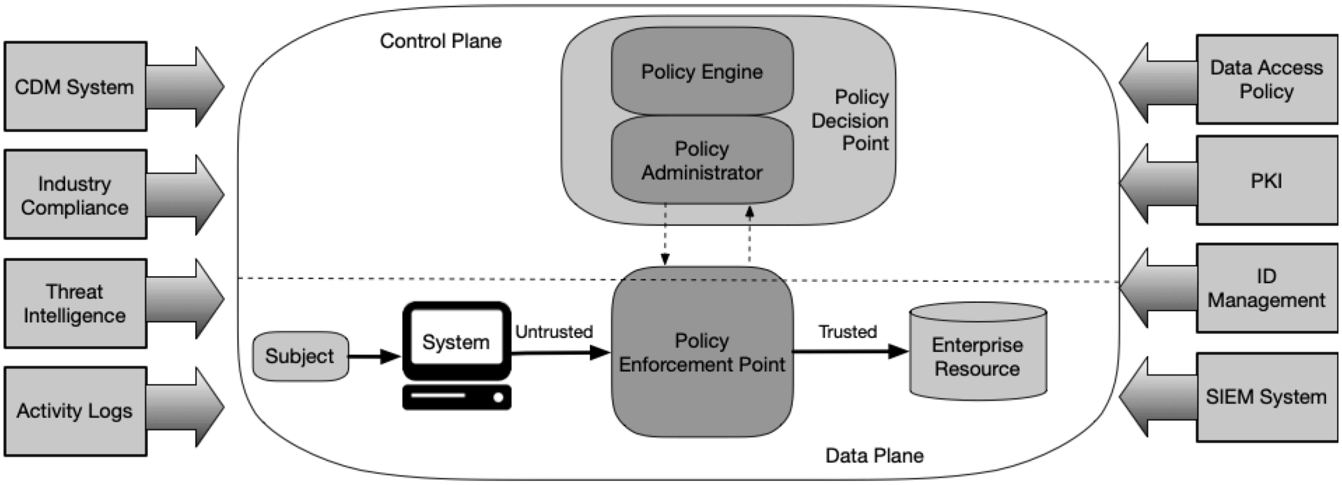

Figure 2: Core Zero Trust Logical Components

Figure 4 - NIST 800-207 SP, Page 9

Logical Components of Zero Trust Architecture

The NIST 800-207 shows the following diagram for the core components of Zero Trust Architecture:

The Policy Enforcement Point is the heart of the Zero Trust Architecture and sits at the center of the diagram. The iboss Zero Trust Edge serves all of the core components at the center of the Zero Trust Architecture diagram, including the Policy Engine, the Policy Administrator and the Policy Enforcement Point which authorizes "per-request" access to protected enterprise resources.

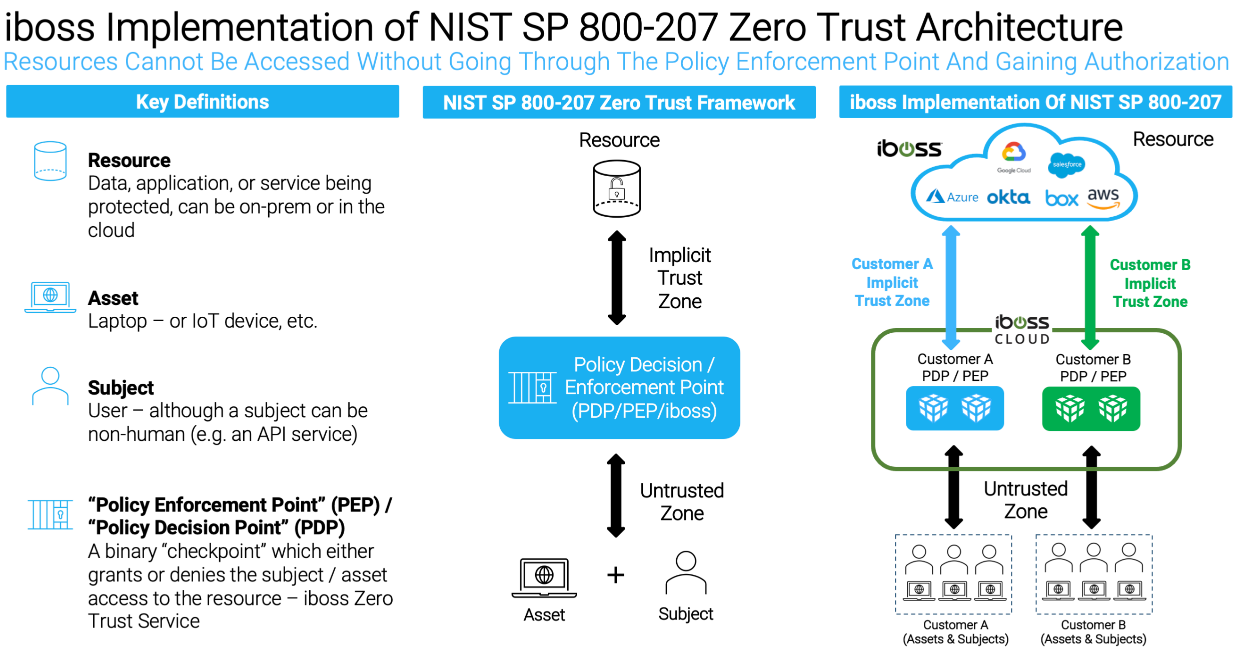

iboss Implementation of NIST SP 800-207 Zero Trust Architecture

Resources Cannot Be Accessed Without Going Through The Policy Enforcement Point And Gaining Authorization

Key Definitions

Resource
Data, application, or service being protected, can be on-prem or in the cloud

Asset
Laptop – or IoT device, etc.

Subject
User – although a subject can be non-human (e.g. an API service)

"Policy Enforcement Point" (PEP) / "Policy Decision Point" (PDP)
A binary "checkpoint" which either grants or denies the subject / asset access to the resource – iboss Zero Trust Service

NIST SP 800-207 Zero Trust Framework
Resource
Implicit Trust Zone
Policy Decision / Enforcement Point (PDP/PEP/iboss)
Untrusted Zone
Asset
+
Subject
iboss Implementation Of NIST SP 800-207
iboss
Resource
Google Cloud
salesforce
Azure
okta
box
aws
Customer A Implicit Trust Zone
Customer B Implicit Trust Zone
iboss CLOUD
Customer A PDP / PEP
Customer B PDP / PEP
Untrusted Zone
Customer A (Assets & Subjects)
Customer B (Assets & Subjects)

Figure 5 - The iboss Zero Trust Edge serves as the heart of the Zero Trust Architecture model

Ultimately, it is the responsibility of the PEP to allow or deny a request to a protected resource. This decision uses signals coming from a variety of sources, including Identity Providers (IdPs) which provide authentication, asset and device management systems, threat feeds and other inputs to ultimately make the most confident decision of whether a request should be authorized or denied to the protected resource.

Categorizing and Classifying Enterprise-Owned Resources

Following along the NIST Risk Management Framework, categorizing and classifying the enterprise-owned resources that need to be protected is the next critical step in implementing the Zero Trust Architecture. Zero Trust starts with a clear understanding of what needs to be protected and a clear understanding of what the security objectives are for each resource as well as the security impact an attacker could have if those resources were affected. Resources include anything the enterprise wants to protect, including data, applications, services or any other technology which can be protected by a Policy Enforcement Point.

Tagging and Labeling Enterprise-Owned Resources

Tagging enterprise-owned resources allows an enterprise to organize and understand what resources it has and what needs to be protected. With the proliferation of SaaS and cloud resources, it has become exponentially more difficult for most organizations to get a clear understanding of the data and applications that are being utilized by its workforce. Zero Trust helps with this as one of the initial steps is to classify and categorize the resources being used so they are clearly understood.

The iboss Zero Trust Edge allows tags to be applied to data, applications and services. The tags represent labels and descriptions of what the resource is and what category the resource belongs to. For example, a SaaS financial ERP system might be tagged with “financial-erp-system” and placed under a category of “Finance”.

With iboss, security objectives and impact levels are associated with each tag. This allows automatic association of security and impact levels while data and applications are classified.

Assigning Security Objectives and Impact Levels

The iboss Zero Trust Edge implements the NIST FIPS 199 model for security objective and impact levels (**https://nvlpubs.nist.gov/nistpubs/fips/nist.fips.199.pdf**). As defined by NIST FIPS 199, for resources needing protection, there are three security objectives an enterprise looks to achieve:

Confidentiality	“Preserving authorized restrictions on information access and disclosure, including means for protecting personal privacy and proprietary information…” [44 U.S.C., Sec. 3542] A loss of confidentiality is the unauthorized disclosure of information.
Integrity	“Guarding against improper information modification or destruction, and includes ensuring information non-repudiation and authenticity…” [44 U.S.C., Sec. 3542] A loss of integrity is the unauthorized modification or destruction of information.
Availability	“Ensuring timely and reliable access to and use of information…” [44 U.S.C., SEC. 3542] A loss of availability is the disruption of access to or use of information or an information system.

Table 1 - NIST FIPS 199 Security Objectives

Confidentiality deals with data loss. Integrity deals with data destruction, as could occur when ransomware encrypts sensitive data for ransom. Availability is related to ensure workers can access the sensitive data needed to complete their tasks. For example, if data is encrypted by ransomware and held hostage for a ransom, the data would not be available to workers which could have a high impact level on the organization's ability to function.

FIPS 199 then outlines three impact levels for each of the above security objectives:

Low	"The loss of confidentiality, integrity, or availability could be expected to have a **limited adverse effect on organizational operations**, organizational assets, or individuals."
Moderate	"The loss of confidentiality, integrity, or availability could be expected to have a **serious adverse effect on organizational operations**, organizational assets, or individuals."
High	"The loss of confidentiality, integrity, or availability could be expected to have a **severe or catastrophic adverse effect on organizational operations**, organizational assets, or individuals."

Table 2 - FIPS 199 Security Objective Impact Levels

For each of the security objectives, they are associated with a low, moderate or high impact level. The goal is to answer questions such as:

- What would happen to the organization if this data was lost and inappropriately disclosed?
- What would happen to the organization if this data was crypto-locked and inaccessible due to a ransomware attack?
- What would happen if users could not access the data due to an ongoing cyberattack?

Ultimately, a low, moderate or high impact level score is associated with confidentiality, then again to integrity and finally to availability.

Each iboss resource tag allows setting a low, moderate or high setting for each security objective including confidentiality, integrity and availability. As each resource is tagged within the iboss Zero Trust Edge, they automatically inherit the security impact levels for each of these security objectives because the tag is already associated with the security objective and impact level information.

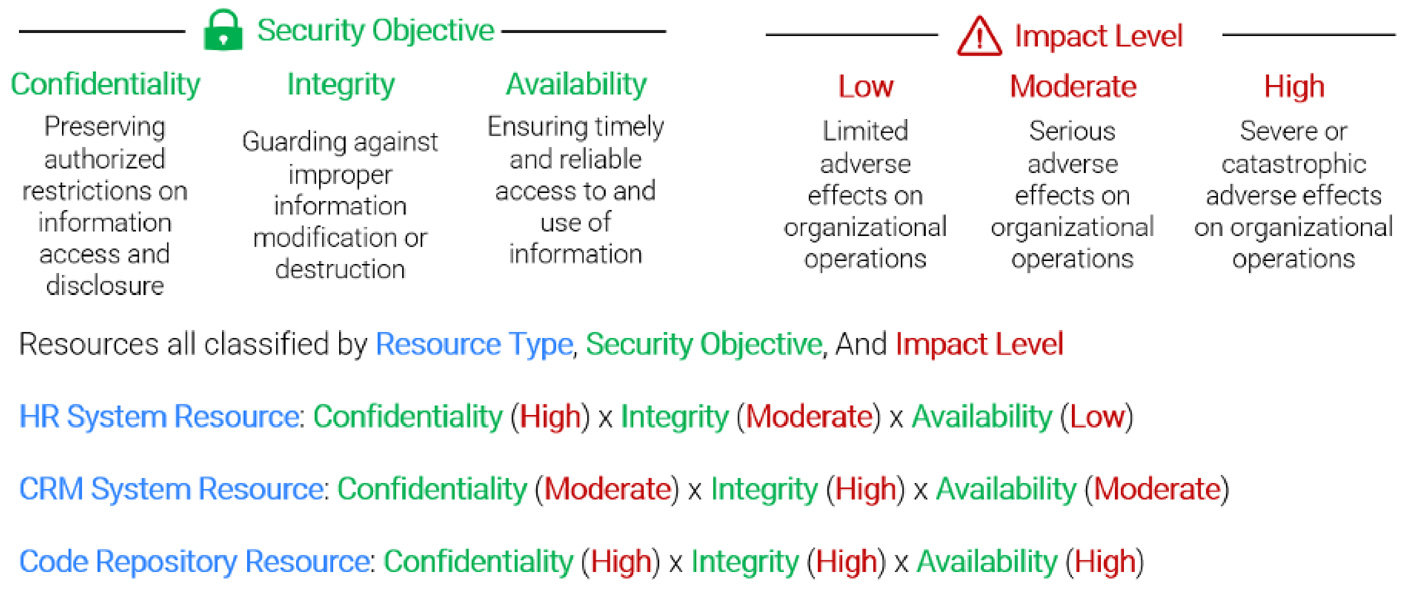

Security Objective

Confidentiality: Preserving authorized restrictions on information access and disclosure

Integrity: Guarding against improper information modification or destruction

Availability: Ensuring timely and reliable access to and use of information

Impact Level

Low: Limited adverse effects on organizational operations

Moderate: Serious adverse effects on organizational operations

High: Severe or catastrophic adverse effects on organizational operations

Resources all classified by Resource Type, Security Objective, And Impact Level

HR System Resource: Confidentiality (High) x Integrity (Moderate) x Availability (Low)

CRM System Resource: Confidentiality (Moderate) x Integrity (High) x Availability (Moderate)

Code Repository Resource: Confidentiality (High) x Integrity (High) x Availability (High)

Figure 6 - Tagging resources within the iboss Zero Trust Edge automatically assigns security objectives and impact levels

Once the resources are identified and tagged, the organization has taken a major step to reducing cyber-risk by providing a clear understanding of what data and applications are in use, where they are located and the security impact a compromise to those resources would have on the organization.

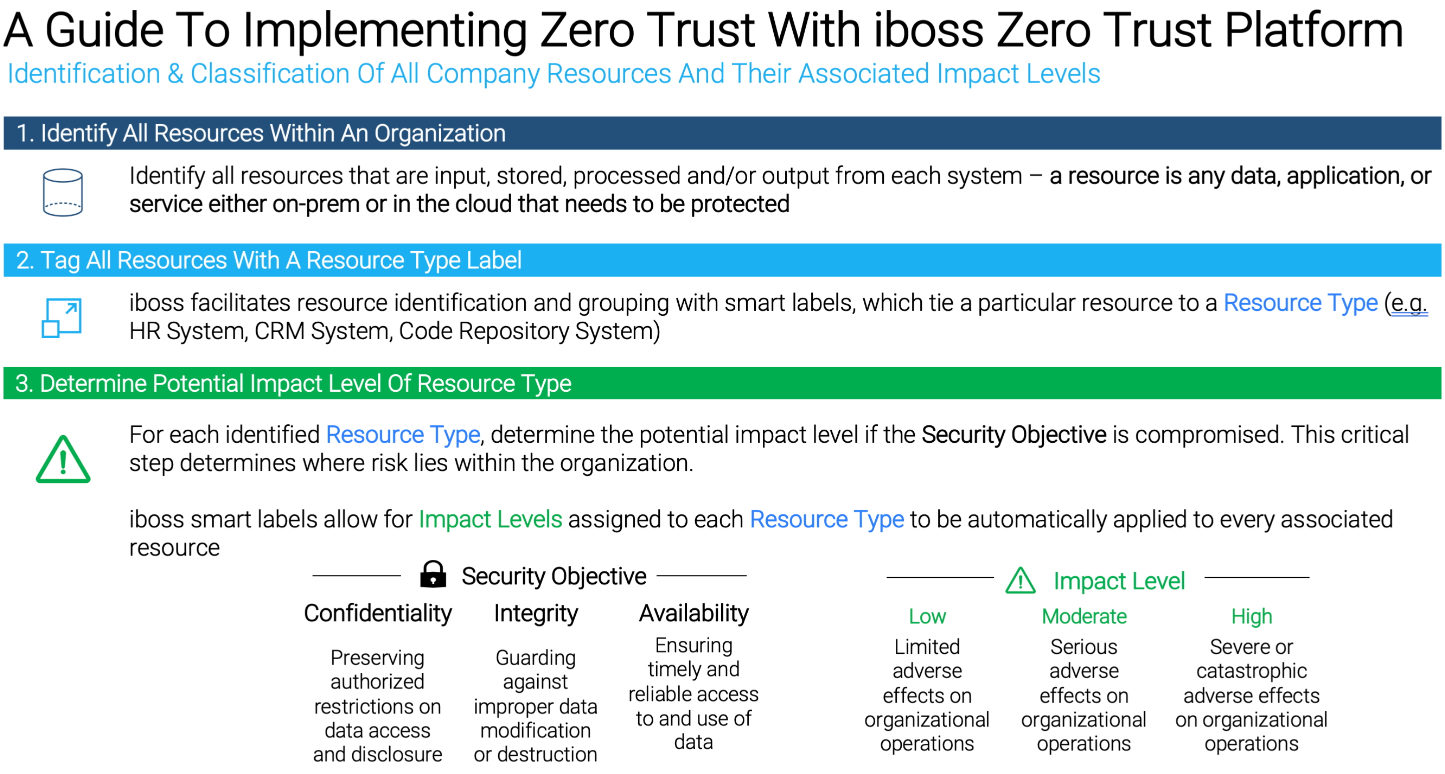

A Guide To Implementing Zero Trust With iboss Zero Trust Platform

Identification & Classification Of All Company Resources And Their Associated Impact Levels

1. Identify All Resources Within An Organization

Identify all resources that are input, stored, processed and/or output from each system – **a resource is any data, application, or service either on-prem or in the cloud that needs to be protected**

2. Tag All Resources With A Resource Type Label

iboss facilitates resource identification and grouping with smart labels, which tie a particular resource to a Resource Type (e.g. HR System, CRM System, Code Repository System)

3. Determine Potential Impact Level Of Resource Type

For each identified Resource Type, determine the potential impact level if the **Security Objective** is compromised. This critical step determines where risk lies within the organization.

iboss smart labels allow for Impact Levels assigned to each Resource Type to be automatically applied to every associated resource

Security Objective

Confidentiality: Preserving authorized restrictions on data access and disclosure

Integrity: Guarding against improper data modification or destruction

Availability: Ensuring timely and reliable access to and use of data

Impact Level

Low: Limited adverse effects on organizational operations

Moderate: Serious adverse effects on organizational operations

High: Severe or catastrophic adverse effects on organizational operations

Figure 7 - The iboss Zero Trust Edge allows enterprises to identify, tag and assign security impact levels to protected resources which greatly reduces cyber-risk and provides visibility to sensitive resources

Creating the Trust Algorithm and iboss Resource Policies

The next step in the process is to create and assign the Trust Algorithms for the protected resources. This is performed within the iboss Zero Trust service which serves as Policy Administrator and the Policy Engine. The Trust Algorithm combines resource policies and dynamic, adaptive scoring which will provide the information necessary for the iboss Policy Enforcement Points (PEPs) to authorize or deny connections to the protected resources that were classified in the previous step.

There NIST 800-207 describes the Trust Algorithm with this example:

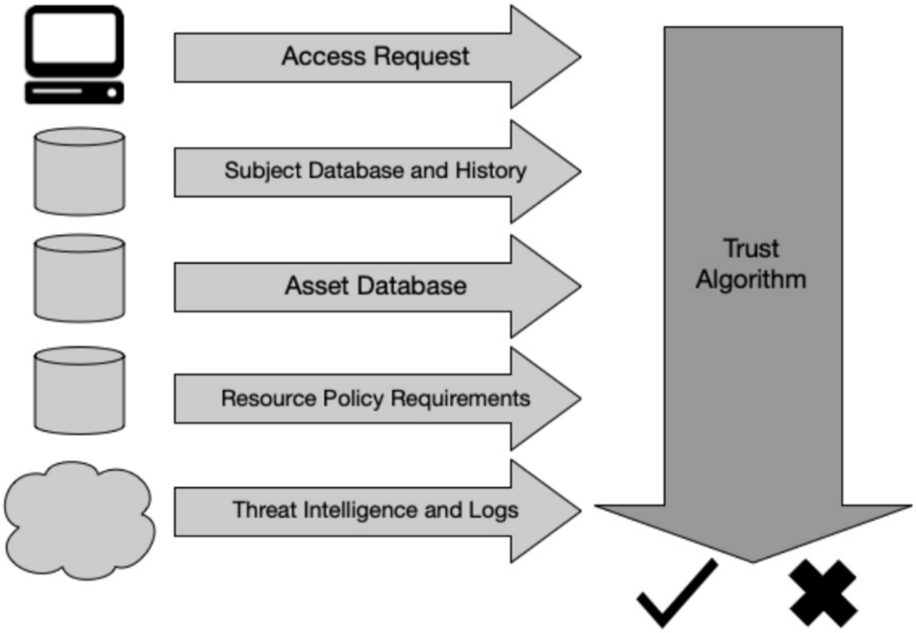
Access Request
Subject Database and History
Asset Database
Resource Policy Requirements
Threat Intelligence and Logs
Trust Algorithm

Figure 7: Trust Algorithm Input

Figure 8 - NIST 800-207 SP, Page 18

The Trust Algorithm takes a set of inputs, such as who needs access to the resource, which asset is being used to access the resource, resource policy requirements and other information to ultimately authorize or deny access to a protected resource. This is a key point. The inputs can vary greatly and can be analog in nature. The output of the Trust Algorithm is binary. It results in the Policy Enforcement Point either granting access to the resource or denying it.

There are key types of Trust Algorithms defined in the NIST 800-207, such as criteria-based versus score-based algorithms. The criteria-based algorithm allows or denies access based on a set of criteria. The score-based algorithm uses a "point system" to grant access to the resource when the score is high enough to ensure high confidence that access should be authorized.

The iboss Zero Trust Edge leverages resource policies to associate the Trust Algorithm to the resources defined previously. Policies can be created for groups of resources or individual resources which provide the inputs necessary for the PEPs to authorize or deny access requests on a per-request basis.

The iboss Zero Trust Edge supports criteria-based and score-based Trust Algorithms within the resource policies. For example, a set of criteria can be created that indicate that users with a certain attribute are allowed to access a resource, including restricting access to the resource to a certain geo. Additionally, the score-based algorithm provides additional context which dynamically applies a score to each and every access request representing the confidence level for allowing and authorizing access to the protected resource. For example, if a user meets the criteria-based access policy for a resource, but the confidence score for that access is too low, the user is automatically denied access to the resource.

iboss Trust Scoring

The iboss Zero Trust Edge goes beyond applying risk scores to users and extends scoring to all actors involved within the interaction for a protected resource. There are three actors typically involved in an exchange with a protected resource: 1) The user, 2) The asset/device, 3) The Resource itself. The iboss Zero Trust Edge dynamically and adaptively applies a score, at the time of access, to the user, to the asset and to the device. It then combines all three scores to produce an overall access confidence score. Administrators can deny access if the score doesn't meet the minimum score threshold level for the user, the asset, the resource or the overall transaction.

In addition, every transaction to protected resources is logged with the associated dynamic transaction confidence score. This allows visibility into changing confidence score for each resource. For example, if a set of assets need critical patches and are accessing a highly sensitive resource, the overall access score for that resource will drop and be made visible to the administrator.

Confidence Level Algorithms Configured For Each Resource Policy To Allow For Real-Time Decision Making

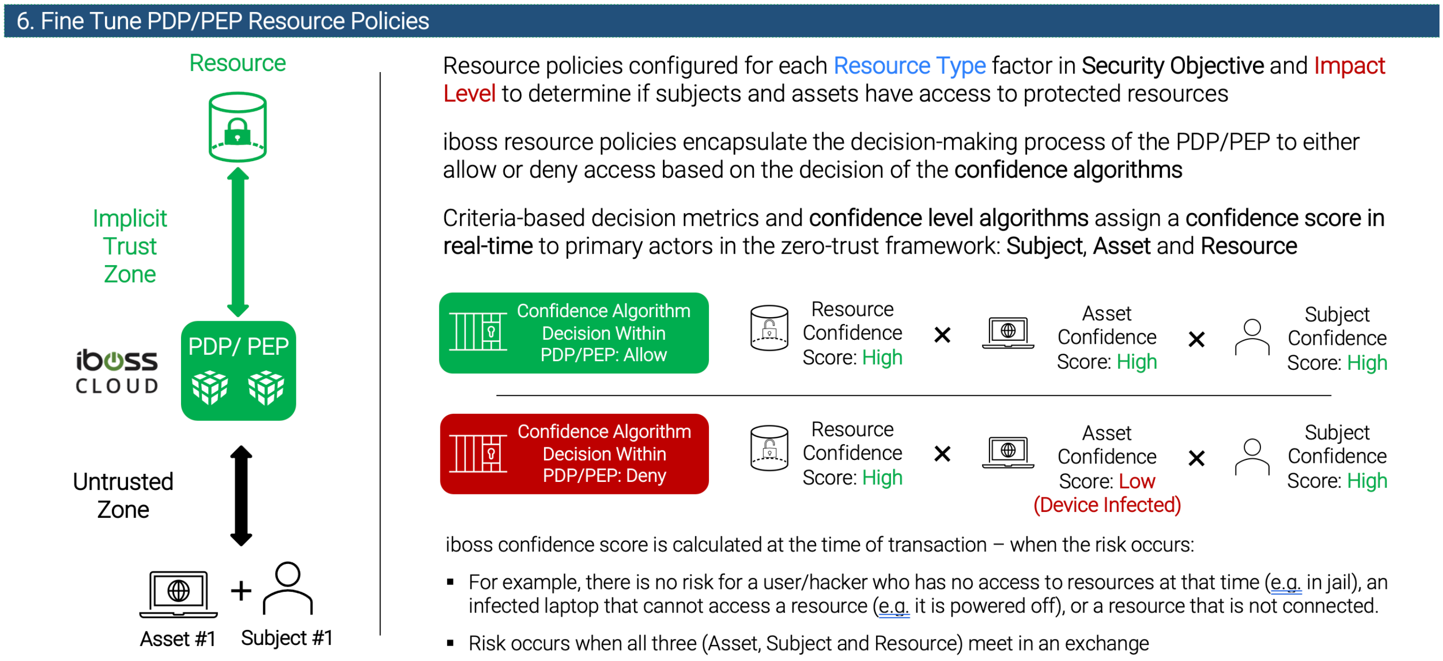

6. Fine Tune PDP/PEP Resource Policies

Resource
Implicit Trust Zone
iboss CLOUD
PDP/ PEP
Untrusted Zone
Asset #1
Subject #1
Confidence Algorithm Decision Within PDP/PEP: Allow
Resource Confidence Score: High
Asset Confidence Score: High
Subject Confidence Score: High
Confidence Algorithm Decision Within PDP/PEP: Deny
Resource Confidence Score: High
Asset Confidence Score: Low (Device Infected)
Subject Confidence Score: High

Resource policies configured for each Resource Type factor in Security Objective and Impact Level to determine if subjects and assets have access to protected resources

iboss resource policies encapsulate the decision-making process of the PDP/PEP to either allow or deny access based on the decision of the confidence algorithms

Criteria-based decision metrics and confidence level algorithms assign a confidence score in real-time to primary actors in the zero-trust framework: Subject, Asset and Resource

iboss confidence score is calculated at the time of transaction – when the risk occurs:

- For example, there is no risk for a user/hacker who has no access to resources at that time (e.g. in jail), an infected laptop that cannot access a resource (e.g. it is powered off), or a resource that is not connected.
- Risk occurs when all three (Asset, Subject and Resource) meet in an exchange

Figure 9 - iboss Zero Trust Edge Trust Algorithm Confidence Scoring

Create Resource Policies for Each Protected Resource

The goal in this step is to create resource policies for each resource. The iboss Zero Trust Edge uses the NIST Zero Trust model for resource policies. The users must be explicitly granted access to a resource while all others are automatically denied. This is an important part of the core definition of the NIST 800-207 Zero Trust Architecture. As a resource policy is created for each resource, the organization is reducing risk by understanding who and what should be interacting with the resources. Ultimately, every resource will be part of a resource policy.

For this step, the resource policies can be set to a monitor only mode so that users can be onboarded without causing major disruptions due to misconfigured resource policies. Transactions in monitor only mode will always be allowed, but the transaction will be recorded for visibility.

Connecting Users through Zero Trust Edge

Now that resources have been classified and resource policies have been created, the next step is to connect users through the iboss Zero Trust Edge so that they are flowing through the Policy Enforcement Points whenever they access the protected resources. The NIST 800-207 shows four primary methods of connecting users. The iboss Zero Trust Edge supports every method described in the NIST 800-207 for connecting assets and users through the Zero Trust Edge.

The four methods of connecting users and assets through the Zero Trust Edge according to the NIST 800-207 are summarized below:

Device Agent/Gateway-Based Deployment	This method is focused on enterprise-owned devices and involves installing an agent on the device which connects through the gateway (PEP) to the resource. Agents are required so that posture checks can be performed on the asset.
Enclave-Based Deployment	This method is focused on enterprise-owned devices and involves installing an agent on the device which connects to a gateway that provides access to a set of resources that sit behind the PEP. This is an example of a hybrid model where the PEP may exist on-prem and authorize connections to protected resources.
Resource Portal-Based Deployment	This is for assets which may not be enterprise-owned but need access to resources. With this method, care needs to be taken so that data does not end up on non-enterprise owned devices so browser isolation can be used which creates a barrier (VDI-like) between the asset and the resource.
Device Application Sandboxing	This is for Operating Systems that support sending data only for applications within the sandbox through the gateway/PEP for access to protected resources. Because the sandbox on the device can be wiped, this can also be used for personal devices. However, this is likely the least popular method due to the need for the Operating System to support the sandboxing feature.

Figure 11 - Methods for connecting assets and users through the Zero Trust Edge according to the NIST 800-207 SP

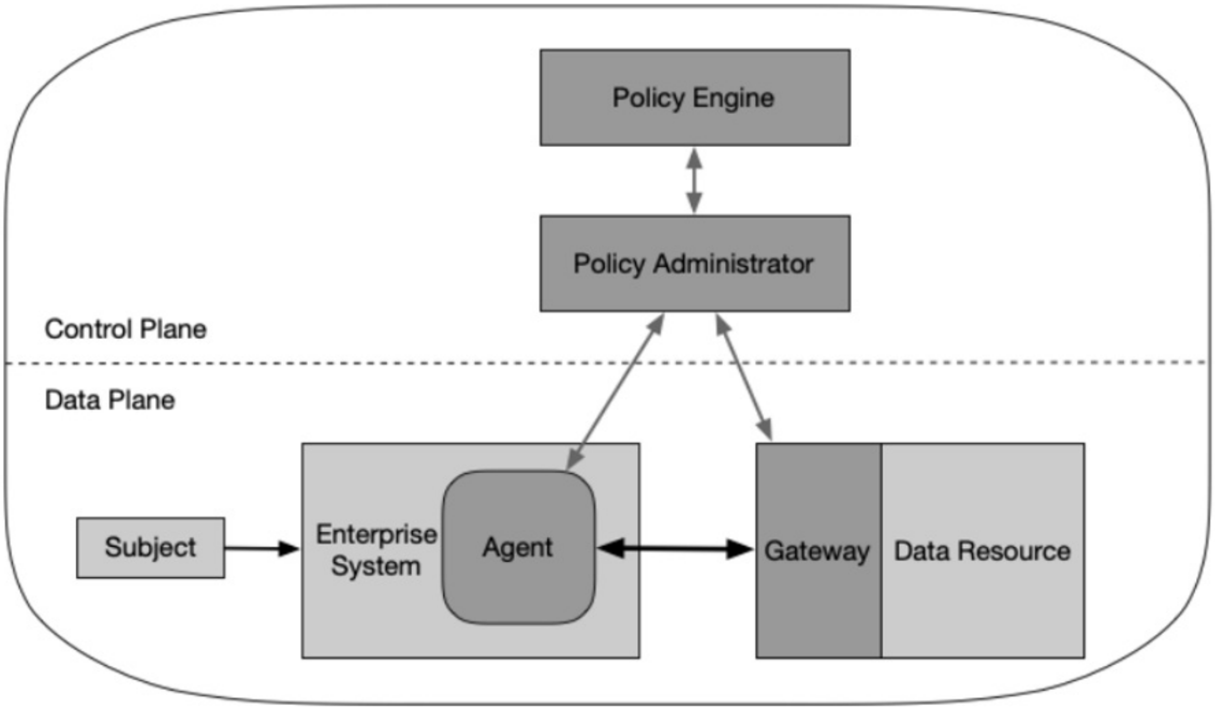

Figure 3: Device Agent/Gateway Model

Figure 12 – The Device Agent/Gateway Model, NIST 800-207 SP, Page 14

Using the Device/Agent Gateway Model to Connect Users

This method is likely the most popular for enterprise-owned assets. An agent is pushed to devices and allows those devices and users to connect to protected resources automatically through the gateway PEP.

The agent is important as it's able to run local device posture checks to ensure the device meets minimum requirements to access protected resources. For example, the agent can check to see that the device has its local firewall enabled and that anti-malware is enabled and running on the asset. This information is fed into the score-based Trust Algorithm which either authorizes or denies access to protected resources.

The iboss Zero Trust Edge has agents for all popular Operating Systems including Windows, Mac, iOS, Linux, Chromebooks and Android. The agents are referred to as "cloud connectors" within the iboss Zero Trust Edge and serve the purpose of connecting users and assets to protected enterprise-owned resources automatically regardless of where the resource is located. The cloud connectors are typically pushed in mass and install silently and automatically. They support being pushed by Mobile Device Management (MDM) or being pushed through SCCM.

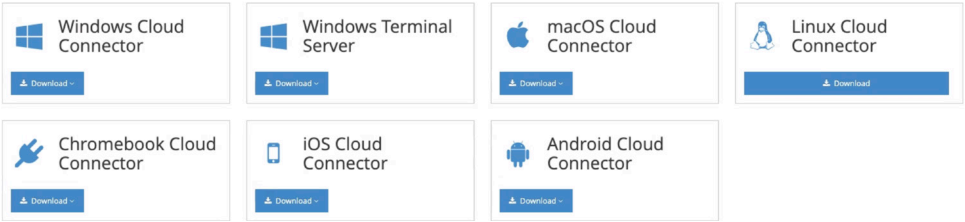

Figure 13 - The iboss Zero Trust Edge has cloud connectors for every popular Operating System

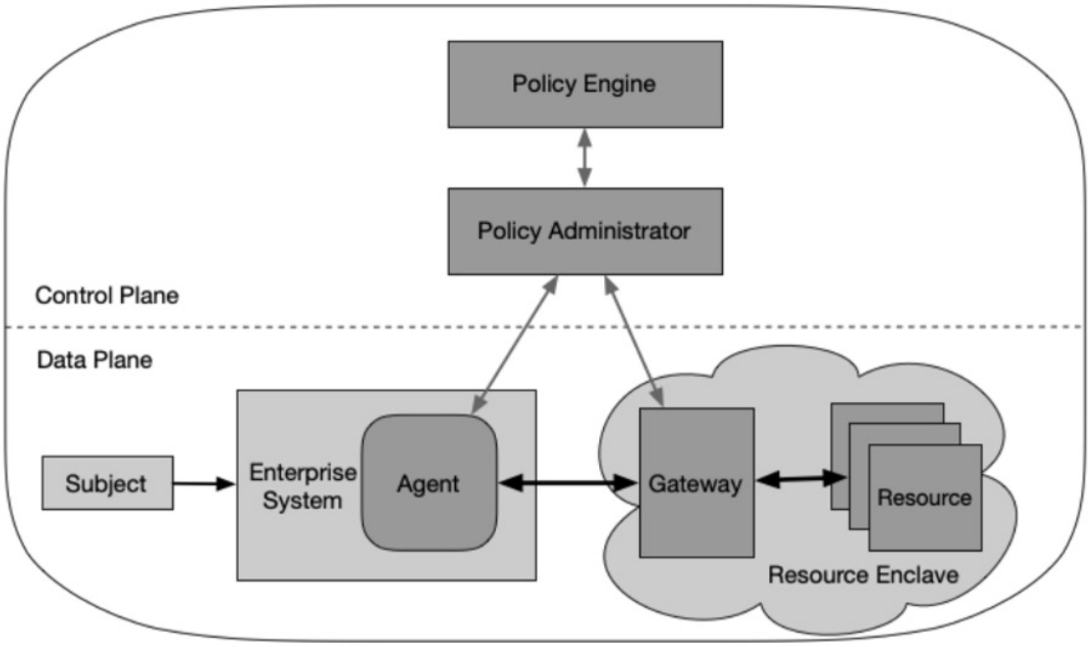

Figure 4: Enclave Gateway Model

Figure 14 - The Enclave Gateway Model, NIST 800-207 SP, Page 15

Using the Enclave Gateway Model to Connect Users

The NIST 800-207 Enclave Gateway Model allows the Zero Trust Edge to extend into a private cloud, such as a data center, to provide access to protected resources.

This model is similar to the Device Agent/Gateway model but supports the gateway being inside of a private data center, for example. The agent connects to the gateway which runs near the protected resources in the private cloud data center which provides access to the protected resources.

The iboss Zero Trust Edge, with its unique containerized architecture, allows the PEPs to extend naturally into the data center to provide access to on-prem protected resources. The on-prem gateways/PEPs extend the PEPs that run within the iboss cloud global service so that security (CASB, malware defense, data loss prevention) and visibility are exactly the same regardless of where the resource lives or the user connects.

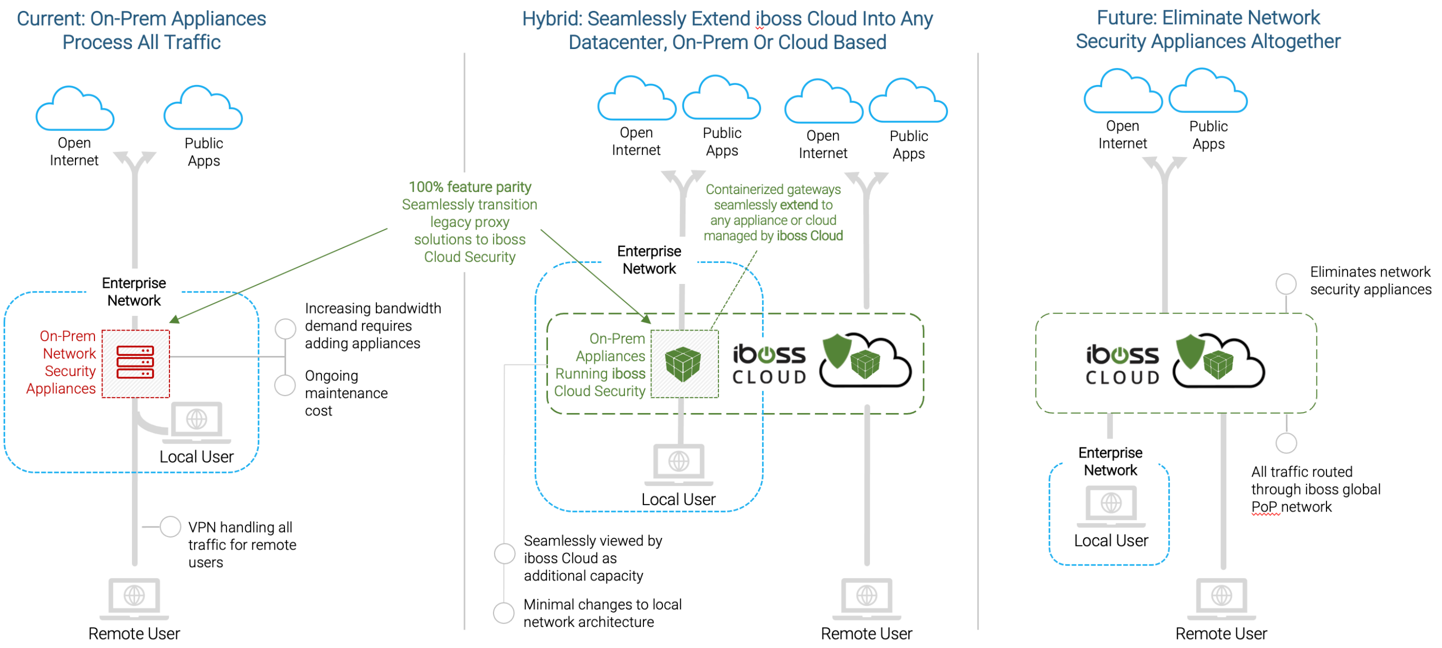

Current: On-Prem Appliances Process All Traffic
Open Internet
Public Apps
Enterprise Network
On-Prem Network Security Appliances
Local User
Increasing bandwidth demand requires adding appliances
Ongoing maintenance cost
VPN handling all traffic for remote users
Remote User
100% feature parity Seamlessly transition legacy proxy solutions to iboss Cloud Security
Hybrid: Seamlessly Extend iboss Cloud Into Any Datacenter, On-Prem Or Cloud Based
Open Internet
Public Apps
Open Internet
Public Apps
Containerized gateways seamlessly extend to any appliance or cloud managed by iboss Cloud
Enterprise Network
On-Prem Appliances Running iboss Cloud Security
iboss CLOUD
Local User
Seamlessly viewed by iboss Cloud as additional capacity
Minimal changes to local network architecture
Remote User
Future: Eliminate Network Security Appliances Altogether
Open Internet
Public Apps
Eliminates network security appliances
iboss CLOUD
Enterprise Network
Local User
All traffic routed through iboss global PoP network
Remote User

Figure 15 - The iboss Zero Trust Edge supports the Gateway Enclave Model by extending the containerized gateways into the data center via physical form containerized gateways

This model is also a great model for enterprises that have legacy appliance-based approaches to security which leverages proxies on-prem to protect resources. The iboss containerized cloud gateways are drop-in replacements to those legacy appliances but are then managed and driven by the global cloud Zero Trust Service. The gateways also broker access to on-prem resources which are protected by the Zero Trust Architecture principles.

Using the Resource Portal Model to Connect BYOD Users

The NIST 800-207 Resource Portal model can be used to connect users that are leveraging personal devices to protected resources. The data from the protected resources must never touch the personal device as that will result in data loss since there is no control of the personal device itself. To mitigate this, the resource portal model can levigate Browser Isolation which puts a pane of glass in front of the resource, allowing the user to interact with the resource while preventing the data from touching the end-user's device. This is the modern equivalent of VDI, but instead uses the end-user's browser to provide the isolated interface. It's scalable and light weight and works on any device with a modern browser.

Within the iboss Zero Trust Edge, Browser Isolation is native and provides a VDI-like interface to the end-user leveraging the user's browser. Because Browser Isolation is a native component of the iboss Zero Trust Edge, it is able to send the data through the same Policy Enforcement Points (PEPs) that the cloud connectors are connected to which ensures:

- The same level of security
- The same level of visibility
- The ability to connect those users with resources on-prem or in the cloud

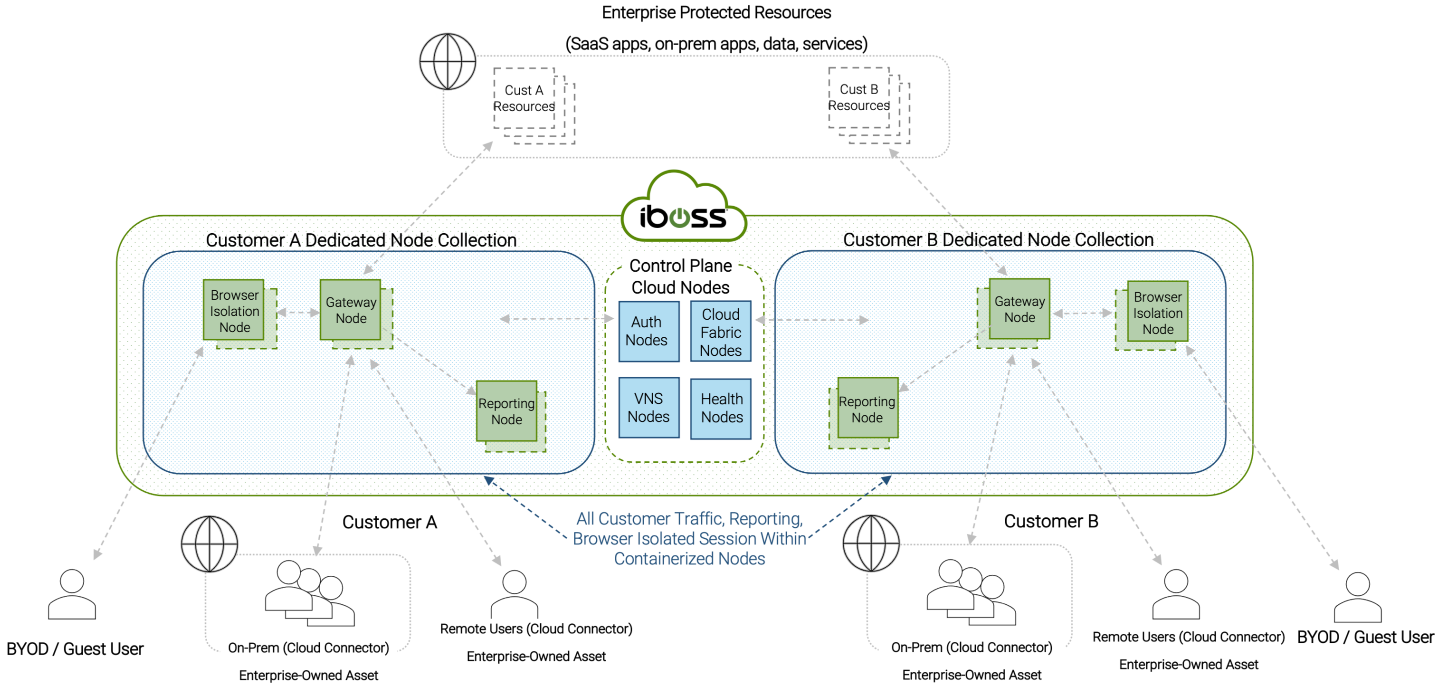
Enterprise Protected Resources
(SaaS apps, on-prem apps, data, services)
Cust A Resources
Cust B Resources
iboss
Customer A Dedicated Node Collection
Customer B Dedicated Node Collection
Browser Isolation Node
Gateway Node
Reporting Node
Control Plane Cloud Nodes
Auth Nodes
Cloud Fabric Nodes
VNS Nodes
Health Nodes
Gateway Node
Browser Isolation Node
Reporting Node
Customer A
All Customer Traffic, Reporting, Browser Isolated Session Within Containerized Nodes
Customer B
BYOD / Guest User
On-Prem (Cloud Connector)
Enterprise-Owned Asset
Remote Users (Cloud Connector)
Enterprise-Owned Asset
On-Prem (Cloud Connector)
Enterprise-Owned Asset
Remote Users (Cloud Connector)
Enterprise-Owned Asset
BYOD / Guest User

Figure 16 - Browser Isolation allows BYOD devices to connect to protected resources with the security and visibility as those connected with cloud connectors on enterprise-owned assets

Leveraging the Device Application Sandboxing Method to Connect Users

For devices that support application sandboxing, such as Android, data can be isolated to the sandbox on the OS to prevent leakage to the general storage of the personal device.

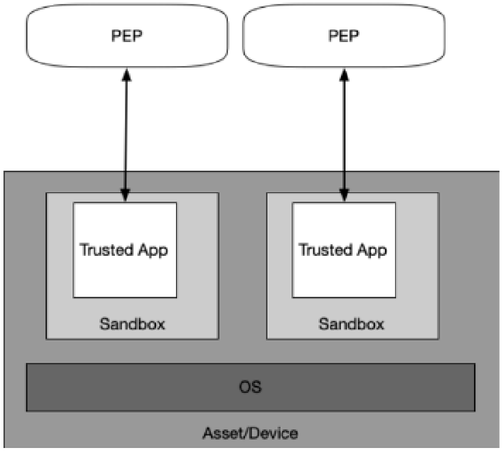

Figure 6: Application Sandboxes

Figure 17 - NIST 800-207 SP Device Application Sandboxing, Page 16

The iboss Zero Trust Edge supports this model and will automatically redirect traffic from the sandbox through the PEPs. This is likely the least popular of the three prior NIST 800-207 methods because the Operating System must support the sandboxing capability. Personal devices must also be enrolled into MDM so that the data and applications within the sandbox can be wiped when the personal device is unenrolled. Given Browser Isolation can completely isolate data from personal devices, this method provides better protection because access to the data is automatically removed at the moment the Browser Isolated session ends. In the sandbox model, if the device is not connected to the Internet, it is possible the data will still reside on the device well after access has been removed. The iboss Zero Trust Edge connectors support sending only data from the sandbox through the gateway PEPs to authorize access to protected resources.

Ensuring All Devices are Connected through the Zero Trust Edge

The goal in this step is to migrate all devices so that they go through the iboss Zero Trust Edge whenever accessing a protected resource. This will likely be performed in phases with a few devices being moved through the Zero Trust Edge first, followed by waves of larger devices being migrated next. As each device is migrated through the iboss Zero Trust Edge, organizational risk is decreased as security and visibility into protected resource access is achieved.

Locking Protected Resource to the Policy Enforcement Points (PEPs)

The final step is to lock resources so that they only interact with gateway Policy Enforcement Points and do not interact directly with any other user or asset. This is the final critical step to achieve a true Zero Trust architecture as this will bring to realization the core definition of the NIST 800-207 Zero Trust architecture which requires authorization for every request accessing a protected resource.

Company Resources Only Accessible To Traffic From iboss PDP/PEP

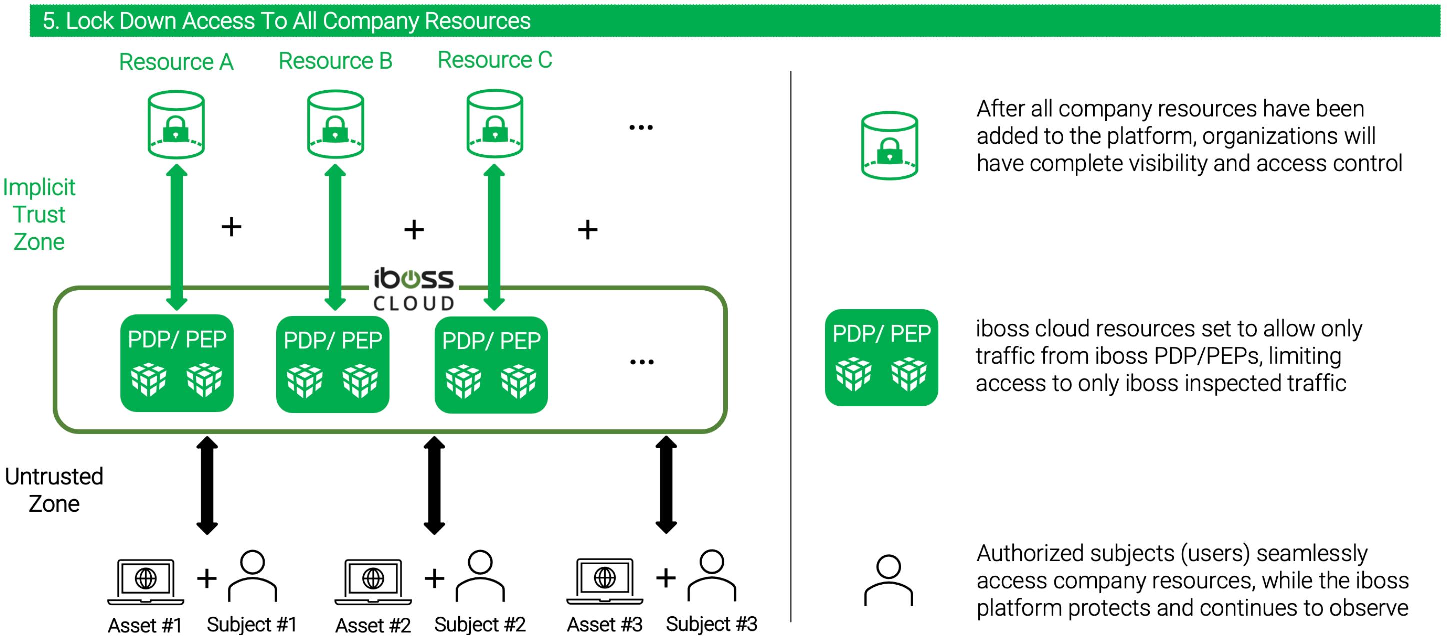

After all company resources have been added to the platform, organizations will have complete visibility and access control

iboss cloud resources set to allow only traffic from iboss PDP/PEPs, limiting access to only iboss inspected traffic

Authorized subjects (users) seamlessly access company resources, while the iboss platform protects and continues to observe

Figure 18 - Protected resources will be locked to the Policy Enforcement Points of the iboss Zero Trust Edge making them completely invisible and inaccessible by attackers and only available to approved and trusted subjects

With the iboss containerized architecture, it is possible to lock every resource, including SaaS applications and data, to the PEP. This is because the iboss Zero Trust Edge will assign unique, unshared IP addresses for the PEPs which only source traffic for a single organization.

Locking the resource can be accomplished using two methods. The first involves configuring the Identity Provider to only allow connections for authentication from the iboss dedicated IP space that represents the Policy Enforcement Points. IdPs support this with the concept of “zoning”. The zone that will accept logins includes traffic from the PEPs.

A second method, if the resource supports it, is to add the IP ACL restrictions on the resource itself so that regardless of the login, the resource only ever accepts requests that come from the PEP. This is very effective in cases such as the “impossible traveler” where a user authenticates to the IdP but their session gets compromised and an attempt to access the resource is made from a distance from the original login that would not be possible. Each and every request is ensured to be authorized regardless of authentication location or method.

For each resource that is locked, the organization increases their security posture and reduces organizational risk. Eventually, all resources are locked to the PEPs and organizational risk is greatly reduced.

Removing Monitor-Only Mode to Enforce Zero Trust Access

In the final step, if the iboss resource policies were placed in “monitor-only” mode, this can be removed so that the PEPs move into a fully enforcing mode. In this mode, only approved subjects are allowed while all others are denied for the resource. Once all of the resources have been moved to an enforcing policy, the organization has transitioned fully to a Zero Trust model.

Monitoring the Zero Trust Implementation

With each and every access being authorized through the PEP, an enterprise now has full visibility about which users, assets and resources are interacting. Additionally, dynamic risk scoring is being generated on each of the users, assets and resources to see how this is trending over time. The iboss Zero Trust Edge provides continuous real-time visibility to all users, assets and resources. It also dynamically adapts to respond to threats in real-time so that resources are protected automatically when the risk changes for a user, subject or resource.

Summary

The iboss Zero Trust Edge provides the ability to implement the NIST 800-207 Zero Trust Architecture Special Publication easily and in a straightforward manner. This transition not only reduces cyber-risk and improves the enterprise’s security posture, but it also changes the way users gain access to protected resources. In legacy models, users are required to enable and disable VPNs to gain access to resources that are located on-prem. With Zero Trust, users are simply connected to the iboss Zero Trust Edge which automatically gives them access to authorized and approved resources, regardless of whether the user is on-prem or in the cloud. The location of the resource is abstracted from the user making the end-user experience better.

Transforming to a Zero Trust model also ensures there is no east-west traffic movement and that all resources are completely invisible and inaccessible by attackers. The Policy Enforcement Points are the only gateway to the protected resource, regardless of whether the resource is in the cloud or on-prem, and will deny access by default while only authorizing approved subjects that meet a high level of confidence that need access to a resource. The journey is one that will revolutionize the way organizations think about protecting resources that are located everywhere, while allowing users and assets to connect to those resources from anywhere.

+1 877.742.6832
sales@iboss.com

101 Federal St
Boston,MA 02110