WHITE PAPER

Zero Trust Network Access (ZTNA) Technical

Implementing BlastShield to Comply with Zero Trust Architecture Requirements

BlastWave

TL:DR

This document provides a technical overview of BlastWave's BlastShield solution that enables Zero-Trust Network Access (ZTNA). The white paper explores the zero trust concepts of "never trust, always verify" and "always assume breach." We break down the principle tenets and approaches of zero trust architecture (ZTA). The document identifies the challenges of meeting the requirements of ZTNA and provides a detailed description of the BlastShield solution, including its components, benefits, technical specifications, implementation guidance, orchestration, compliance, use cases, and pricing structure.

The intended audience for this document is chief information security officers (CISOs); chief information officers (CIOs); IT and OT security directors; governance, risk, and compliance managers; security architects; information security managers; identity and access management administrators; product security managers; network administrators; and DevOps security managers. This document will help stakeholders develop zero trust strategies to protect enterprise resources across information technology (IT), operational technology (OT), and the Internet of Things (IoT).

Business unit leaders, cybersecurity policy makers, and insurers will also benefit from this information to better evolve their existing environment towards zero trust using software-defined perimeter (SDP) approaches.

BlastWave

CONTENTS

Zero Trust Network Access:
What is Zero trust?

Zero Trust (ZT) is a set of cybersecurity concepts and approaches that move defenses from network-centric, perimeter-based security (e.g. "castle and moat") to a model where the trust of users, assets, and resources is not implied.

In today's ubiquitous enterprise perimeter-based security models, users are provided access to assets and resources based on their location and network connectivity. For example, employees connecting to a corporate network at their physical office may have broad access to devices such as servers, printers, and workstations across the network. In this perimeter-centric model, access to applications and data is controlled by managing a complex system of configurations and policies that include network subnetting, firewall access control lists (ACLs), authentication, virtual private networks (VPNs), digital certificates, identity management, and privileged access management (PAM). User permissions differ based on whether they are inside or outside the perimeter.

While the concepts of Zero Trust have been around for a long time, John Kindervag of Forrester is credited with coining the term. (ZT) was born out of necessity with the movement of resources and users outside of the traditional enterprise perimeter. Assets and data have moved from private data centers to cloud applications, such as Amazon Web Services, Microsoft Azure, Google Cloud Platform, Salesforce.com, ServiceNow, Hubspot, Workday, Monday, Asana, SAP Cloud, Oracle Cloud, and thousands of other SaaS applications. Users, including employees and contractors, are now spending more time working remotely than from the office. Finally, the explosion of the Internet of things (IoT) has added billions of new devices connected to the network that may connect to industrial operational technology (OT) systems and private and public cloud applications.

In 2020, the United States National Institute of Standards and Technology (NIST) published "NIST SP 800-207: Zero Trust Architecture." A ZT approach focuses on protecting all enterprise assets (devices, infrastructure components, applications, virtual, and cloud components) and subjects (end users, applications, and other non-human entities that request information from resources).

A Definition of Zero Trust and Zero Trust Architecture[2]:

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.

The Tenets of Zero Trust Architecture

The goal of ZTA is to prevent unauthorized access to data and services while making the enforcement of access controls as granular as possible. This involves using authentication and authorization to shrink the "implicit trust zone" as much as possible while maintaining broad visibility.

The principle tenets of zero trust architecture include:

1. All data sources and computer services are considered resources
2. All communication channels are secure regardless of network location
3. Access to individual enterprise resources is granted on a per-session basis
4. Access to resources is determined by dynamic policy – including the observable state of client identity, application/service, and the requesting asset – and may include other behavioral and environmental attributes
5. The enterprise monitors and measures the integrity of security posture of all owned and associated assets
6. All resource authentication and authorization are dynamic and strictly enforced before access is allowed
7. The enterprise collects as much information as possible about the current state of assets, network infrastructure, and communications and uses it to improve its security posture

In a nutshell, zero trust means moving to a model that grants access on a granular, per-session basis to all data sources and computer services while monitoring the integrity of security posture and collecting as much information as possible.

Figure 1: Core Zero Trust Logical Components

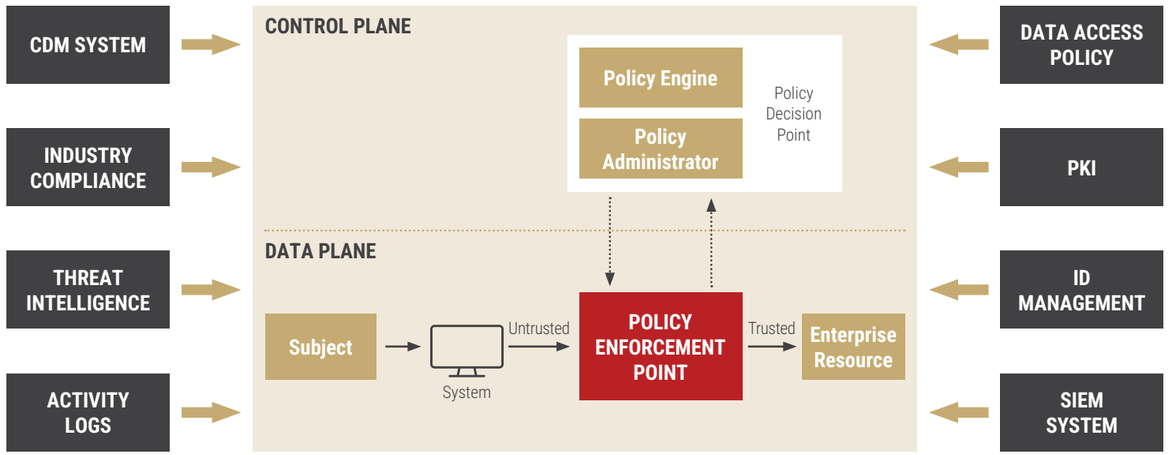

BlastWave

Zero Trust Architecture **Deployment Models**

To understand zero trust, it is helpful to understand the components of zero trust architecture according to NIST SP 800-207:

Policy engine (PE)

This component is responsible for the decision to grant access to a resource for a given subject. The PE uses policies and input from external sources as input to a trust algorithm to grant, deny, or revoke access to the resource. The PE is paired with the policy administrator component. The policy engine makes and logs the decision (as approved, or denied), and the policy administrator executes the decision.

Policy administrator (PA)

This component is responsible for establishing and/or shutting down the communication path between a subject and a resource (via commands to relevant PEPs). It generates any session-specific authentication and authentication token or credential used by a client to access an enterprise resource. It is closely tied to the PE and relies on its decision to ultimately allow or deny a session.

Policy enforcement point (PEP)

This system is responsible for enabling, monitoring, and eventually terminating connections between a subject and an enterprise resource. The PEP communicates with the PA to forward requests and/or receive policy updates from the PA. While this is represented as a single logical component in ZTA, it may be broken into two different components: the client (e.g., software agent on a laptop) and resource side component (e.g., software agent on a host device, gateway that controls access, or a single portal component that acts as a gatekeeper to downstream devices). Beyond the PEP is the trust zone hosting the enterprise resource.

There are several variations in ZTA approaches, including:

ZTA using enhanced identity governance (EIG)

This model uses the identity of actors as the key component for policy creation. Access policies are based on identity and assigned attributes. EIG approaches are often implemented using an open network model whereby subjects are provided network access but access to resources is restructured based on identity, roles, and permissions. **The downside in an EIG approach is that malicious actors could still attempt network reconnaissance and/or launch a denial of service (DOS) attack internally or against a third party.**

ZTA using microsegmentation

This model uses a familiar approach to using switches, firewalls, special purpose gateways, and host-based endpoint protection as PEPs to protect resources or a small group of resources. **While used extensively, this model is complex, challenging to manage, and prone to human error.**

ZTA using cloud-based portal or device/gateway

This model is similar to the ZTA using microsegmentation, except the PEP is moved into the cloud in the form of a gateway or portal to provide access to resources. **While this model simplifies management, it allows** attackers to discover and attempt to access the gateway or portal and attempt a (DOS) attack or otherwise disrupt the network.

ZTA using software-defined perimeters (SDPs)

This approach is often implemented as an overlay to an existing network and provides a level of protection that transcends network location or underlying configuration. This model protects every resource and subject by implementing a PEP on each enterprise asset, gateway, and end user device.

Gateways protect downstream assets that cannot host a PEP, such as a legacy OT system or IoT device. In this model, the agent on the end user device makes a request to the PA to access the resource. The PA forwards the request to the PE and, if the request is approved, the PA sets up the connection over the control plane, and the end user device and resource communicate over the data plane. In this model, the perimeter is not defined by the location of the network - it is defined by the presence of the PEP software on an asset or end user device. **While this model is easier to manage and provides granular and precise access controls, the organization must have a robust device management system and the ability to install software broadly across assets and end user systems.**

Figure 2: ZTA Device Agent/Gateway Model

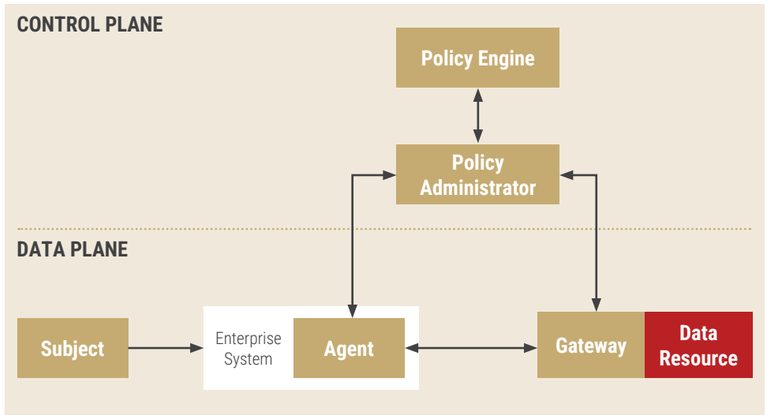

BlastWave

Zero Trust Compliance

Zero Trust is gaining momentum across the public and private sectors. In addition to the NIST SP 800-207 that was published to better define ZTA, technology analysts Gartner and Forrester have continued to expand their coverage and guidance in this area. Additionally, in May 2021, the U.S. White House issued Executive Order 140283 on "Improving the Nation's Cybersecurity," ushering in sweeping changes to migrate the Federal Government to a zero trust architecture. The next month, the Cybersecurity and Infrastructure Security Agency (CISA) published the "Zero Trust Maturity Model."

Recognizing the increase in sophisticated and persistent threats impacting U.S. organizations and the importance of improving the nation's ability to protect critical systems and detect and respond to incidents, the executive order directs the Cybersecurity and Infrastructure Security Agency (CISA), the Federal Bureau of Investigation (FBI), the Intelligence Community (IC), and Office of Management and Budget (OMB) to act to:

- Remove barriers to sharing threat information
- Modernize the Federal Government's cybersecurity
- Enhance the software supply chain
- Establish a cyber safety review board
- Standardize the Federal Government's playbook for responding to cybersecurity vulnerabilities and incidents
- Improve detection of cybersecurity vulnerabilities and incidents on Federal Government networks
- Adopt national security systems cybersecurity requirements that are equivalent to or better than the requirements set forth in the executive order

Following the EO 14028, the Office of Management and Budget (OMB) issued Memorandum M-22-09 that sets forth a Federal zero trust architecture strategy5 , requiring Federal Government agencies to meet specific cybersecurity goals by the end of fiscal year (FY) 2024. The memorandum cybersecurity goals, aligned with the five pillars in the CISA Zero Trust Maturity Model, include:

Identity

Agency staff use enterprise-managed identities to access the applications they use in their work. Phishing-resistant multi-factor authentication (MFA) protects those personnel from sophisticated online attacks.

Devices

The Federal Government has a complete inventory of every device it operates and authorizes for Government use, and can prevent, detect, and respond to incidents on those devices.

Network

Agencies encrypt all DNS requests and HTTP traffic within their environment and begin executing a plan to break down their perimeters into isolated environments.

Applications and Workloads

Agencies treat all applications as internet connected, routinely subject their applications to rigorous empirical testing, and welcome external vulnerability reports.

Data

Agencies are on a clear, shared path to deploy protections that make use of thorough data categorization. Agencies are taking advantage of cloud security services to monitor access to their sensitive data and have implemented enterprise-wide logging and information sharing.

Notably, the memorandum instructs users to stop using MFA that relies on SMS, voice calls, one-time codes, or push notifications by 2024. By January 2023, agencies must offer a phishing-resistant MFA option (PR-MFA). PR-MFA options include using:

Federal Government Personal Identity Verification (PIV) such as phishing-resistant tokens, including PIV cards

Derived PIV authenticators such as FIDO2 security keys and WebAuthn

Non-PIV authenticators that use phishing-resistant tokens such as biometrics and public key cryptography

Agencies are also encouraged to pursue greater use of passwordless authentication methods. They should not rely on privileged access management (PAM) solutions that authenticate through single-factor credentials or passwords that require special characters and regular password rotation because these methods cannot withstand phishing attacks and credential theft.

The adoption of zero trust architecture will help enterprise organizations meet a variety of digital identity, access control, authentication, and confidentiality cybersecurity standards and guidance, including:

IEC 62443-3-3 System Security Requirements and Security Levels

- **FR 1:** Identification and Authentication Control
- **FR 4:** Data Confidentiality
- **FR 5:** Restricted Data Flow

IEC 62443-4-2 Technical Security Requirements for IACS Components

- **FR 1:** Identification and Authentication Control
- **FR 4:** Data Confidentiality
- **FR 5:** Restricted Data Flow

ISO/IEC 27001 Information Security Standard

- **A.9** Access Control
- **A.10** Cryptography
- **A.13** Communications Security

NIST SP 800-53 r5

- **3.1** Access Control
- **3.7** Identification and Authentication
- **3.20** Supply Chain Risk Management

NIST SP 800-63A, 63B, 63C, and 63-3 Digital Identity Guidelines

- Authenticator Assurance Levels
- Authenticator and Verifier Requirements
- Authenticator Lifecycle Requirements
- Session Management
- Derived Credentials
- Privacy Considerations

NIST SP 800-171 Protecting Controlled Unclassified Information in Non-federal Systems and Organizations

- **3.1** Access Control
- **3.5** Identification and Authentication

NIST SP 800-172 Protecting Controlled Unclassified Information in Non-federal Systems and Organizations

- **3.1** Access Control
- **3.5** Identification and Authentication

NIST SP 800-207: Zero Trust Architecture

CMMC 2.0 Cybersecurity Framework

- **Level 2 Advanced:** aligned with SP 800-171
- **Level 3 Expert:** aligned with SP 800-172

U.S. White House EO 14028: Executive Order on Improving the Nation's Cybersecurity

OMB Memorandum M-22-09: Moving the U.S. Government Toward Zero Trust Cybersecurity Principles

[5] "M-22-09 Memorandum for the Heads of Executive Departments and Agencies." U.S. Office of Management and Budget. January 26, 2022.

https://www.whitehouse.gov/wp-content uploads/2022/01/M-22-09.pdf

BlastShield Overview

What is BlastShield?

BlastShield is a zero trust network access (ZTNA) solution that helps organizations meet the requirements of a zero trust architecture (ZTA). Unlike other ZTA approaches that rely primarily on enhanced identity governance (EIG), complex layers of microsegmentation, or cloud-based gateways, BlastShield leverages a software-defined perimeter (SDP) approach that provides granular access controls while avoiding the risks of stolen credentials and complex management.

BlastShield combines phishing-resistant MFA, simple orchestration, granular access controls, peer-to-peer full-mesh networking, and device cloaking to provide a level of security that addresses the ZTA goals of preventing unauthorized access and making access control enforcement as granular as possible. Unlike any other ZT solution provider, BlastWave radically simplifies the security stack and enables zero trust without sacrificing performance.

Figure 3: Software-Defined Perimeter Components based on the CSA SDP Specification 1.0

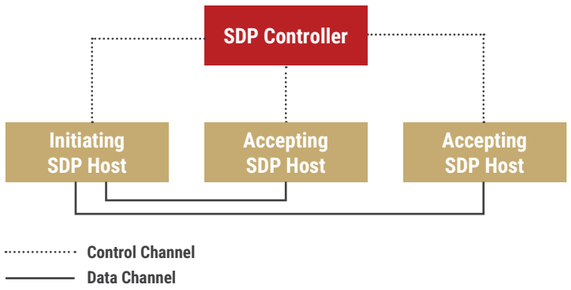

What is a Software-Defined Perimeter?

A software-defined perimeter (SDP) is a security technique that controls access to resources based on identity and forms a virtual boundary around networked resources. Based on work done by the U.S. Department of Defense, the Cloud Security Alliance further developed the SDP framework in 2014.

In the CSA SDP model, an Initiating Hosts (IH) communicates with the SDP Controller to request a list of Accepting Hosts (AH) to which they can connect. The AH accepts connections from an IH only after being instructed to do so by the Controller.

NIST SP 800-207, published in 2020, goes into more detail in describing how a SDP can be used to meet the requirements of a zero trust architecture (ZTA). In this context, a Policy Enforcement Point (PEP) is installed on a User's system and a Resource, such as a client and gateway. The PEP communicates with a Policy Administrator (PA) on the control plane to forward connection requests and receive policy updates. A Policy Engine (PE) is responsible for making the decision to grant access. The PEPs communicate with each other over the data plane.

In both of these SDP definitions, a client makes a request to connect to a resource, and a controller enforces a policy, making the controller a single-point of failure in the real-time establishment of a connection. BlastShield has improved upon these basic SDP models by pushing policies and permissions to the protected endpoint devices so that they can connect with each other without connecting to the controller in real-time, provided they are authenticated and have the granular access permissions in the latest policy update.

Figure 4: Software-Defined Perimeter Components ZTA Software Defined Perimeter with Resource Enclave

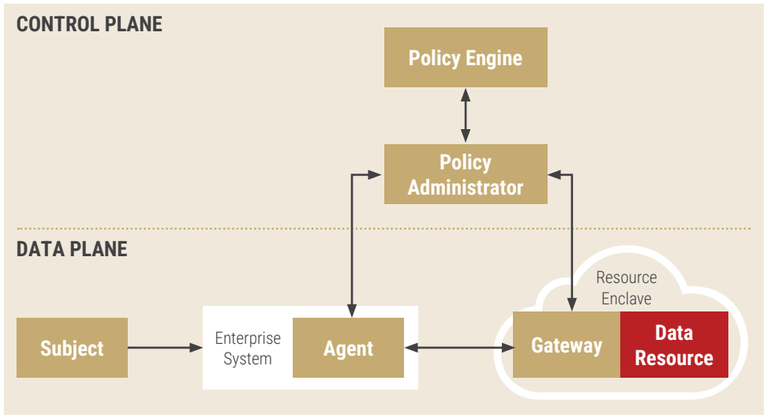

While IT focuses on information security, OT prioritizes operational continuity and safety. Because when OT networks stop, the world stops.

BlastShield Solution Components

BlastShield simplifies security by combining many security controls into a single solution. These security controls (phishing resistant MFA, encryption of data-in-motion, microsegmentation, granular access controls, device invisibility, and application proxy) are enabled by deploying software agents (aka policy enforcement point or PEP) on end user devices, host machines, and gateway appliances. The agents and security controls are managed using the BlastShield Orchestrator. BlastShield components include:

Client

The BlastShield Client is downloadable software for Microsoft Windows, macOS iOS, Linux, and Android. The Client is deployed on end user devices that initiate requests to resources protected by BlastShield. Available for download via the BlastWave website, Apple App Store, and Google Play store, the Client is considered a ZTA Policy Enforcement Point (PEP) for user devices.

Authenticator

The BlastShield Authenticator is downloadable software for iOS and Android mobile devices. The Authenticator is used to facilitate phishing-resistant passwordless authentication. The user registers the Authenticator with the Client when the Client is installed on the user device. Subsequently, when logging into the Client, a user can authenticate without a password using the Authenticator or a FIDO2 security key.

Host Agent

The BlastShield Host Agent is a software agent that is installed on any IP-connected physical or virtual machine running Linux, Microsoft Windows, or macOS. The Host Agent Software is considered a ZTA PEP for resources. When the Host Agent is installed on a target device, the administrator must also install a special file generated by the Orchestrator that initiates an authentication process that validates the identity agent and on-boards the device by having it generate a new public-private key pair used for authentication and encryption.

Gateway Agent

The BlastShield Gateway Agent provides protection of endpoints that are not protected by a Host Agent. A Gateway is created by installing the BlastShield Gateway Agent software on a physical or virtual machine. Gateways can identify and connect to Endpoints by using three Addressing Modes: MAC address, VLAN, or NAT. The Gateway Agent software is considered a ZTA PEP installed on a gateway in front of resources. The gateway can be configured as Active or Passive, depending upon the use case.

- **Active:** The gateway is set up inline to protect downstream Endpoints that are registered with the gateway. To reach the Endpoints, traffic must flow inline through the gateway. This model is effective at protecting Endpoints from internal attackers.
- **Passive:** The gateway is set up on the network and not inline. Clients can only connect to Endpoints that are registered with the gateway. This model is effective for secure remote access to legacy infrastructure without impacting other devices communicating on the network.

Orchestrator

The BlastShield Orchestrator is a cloud-based application that provides a single pane of glass to manage Users, Agents, Groups, Policies, Services, and Proxies. The Orchestrator generates special files called BlastShield Invitations (.bsi file) that are used during the onboarding of a device with a Host or Gateway Agent. The Orchestrator uses simple concepts to organize Users and Agents into Groups. Policies can be created that allow Groups of Users and Agents to communicate with each other using granular access controls. Furthermore, communication can be filtered by IP protocol(e.g. TCP, UDP, HTTPS, etc.). Finally, the Orchestrator can be used to set up Proxies that allow administrators to proxy traffic to specifically configured domains enabling conditional access to cloud applications. The Orchestrator participates in registration and session establishment. The Orchestrator is not an in-line gateway that proxies all traffic like many other SDPs and cloud-based SASE solutions.

The Orchestrator is cloud-based; however, BlastWave enables customers to deploy and self-manage the Orchestrator on-premise to support air-gapped networks and highly confidential data enclaves. The Orchestrator performs the functions of the ZTA Policy Engine (PE) and Policy Administrator (PA).

Together the BlastShield Client, Authenticator, Host Agent, Gateway Agent, and Orchestrator enable security controls that make it easy to set up explicit access between users that have been authenticated using phishing-resistant MFA and agents that have been registered using public key cryptography that meets the highest levels of authentication assurance as defined by NIST SP 800-63.

BlastShield is suitable for implementation on a variety of target devices in IT, OT, and IoT environments. Devices that cannot be installed with a BlastShield Agent can sit behind a BlastShield Gateway, enabling organizations to protect IoT devices, IP cameras, legacy infrastructure, and other constrained devices.

Figure 5: BlastShield Solution Components

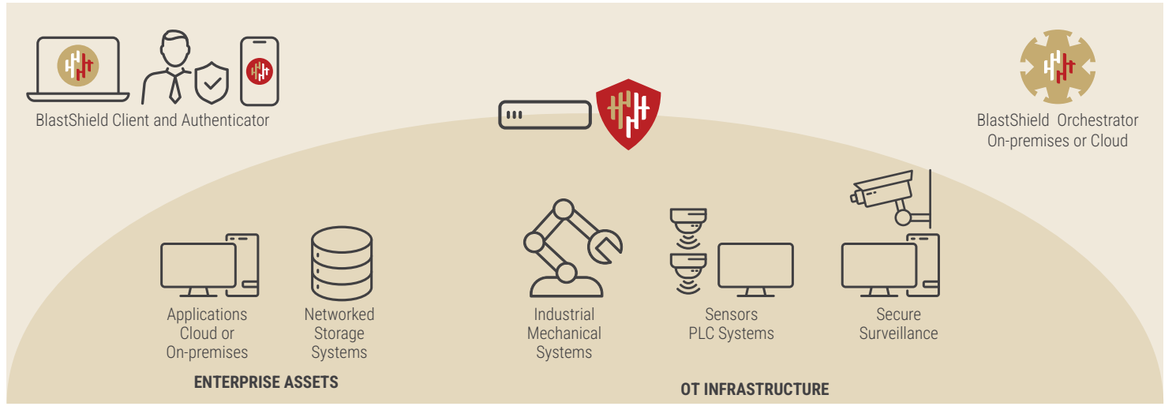

Figure 6: BlastShield Zero Trust Architecture

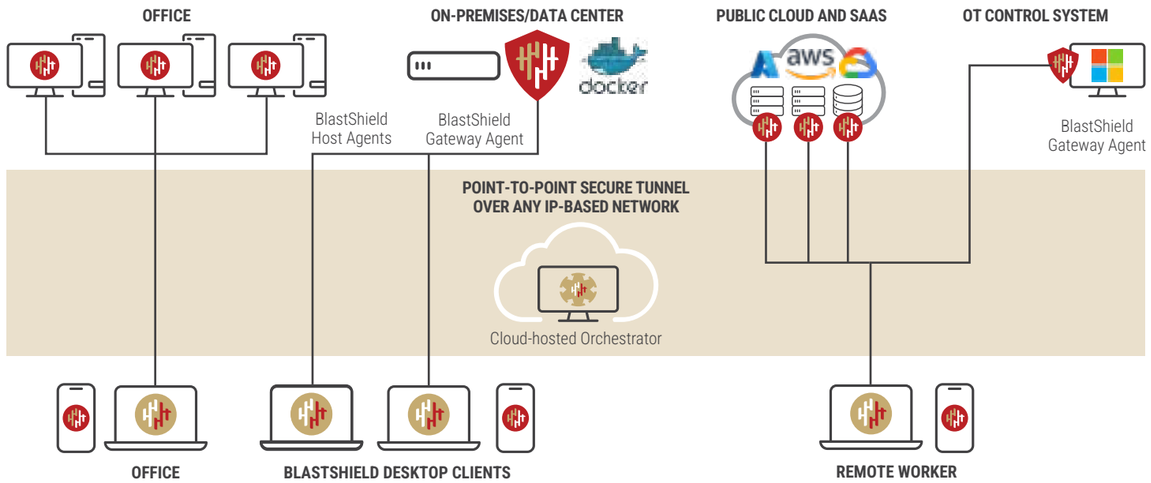

Key Features and Benefits

The BlastWave BlastShield solution resolves the IP transport problem at the IP layer itself by hiding the network assets using a software-defined passwordless solution that cannot be tampered with. This solution is the only meshed overlay IP-based solution in existence today that covers the IP network end-to-end, from one asset to another, across the open IP network.

Key features include phishing-resistant MFA, encryption of data-in-motion, microsegmentation, granular access controls, device invisibility, and application proxy.

Phishing-Resistant MFA

BlastShield enforces phishing-resistant MFA for users logging into the BlastShield network. BlastShield supports two methods of passwordless MFA: 1) BlastShield Authenticator plus a biometric (something you have and something you are) and 2) FIDO2 security key plus a passcode (something you have and something you are) and 2) FIDO2 security key plus a passcode (something you have and something you know).

When a user installs the BlastShield Client on their user device, they confirm their identity using one of the passwordless MFA methods. A public key generated by the Authenticator App or FIDO2 security key is registered with the Orchestrator to confirm the identity of the user each time they log in. Future logins use a challenge-response method that uses the public key of the user's Authenticator or FIDO security key, so that only that device can attest to its identity. What makes BlastShield's MFA method phishing-resistant is that both factors of authentication are unique to the user and cannot be stolen, used, or derived remotely.

Figure 7: Phishing-Resistant MFA

Encryption of Data-in-Motion

When installing a new BlastShield Client or Agent onto a machine, a public key associated with the specific device is sent encrypted to the Orchestrator to be stored in its database. The Orchestrator sends relevant Policies and the public keys of the Agents to which the device has the permission to connect.

When a Client makes a request to connect to an Agent for which it has the permission to connect, the Client checks to make sure that the Client is authorized to connect to the Agent. If it is not Authorized, the packet is dropped. If it is authorized to connect to the Agent, the Agent must verify that the Client is authorized to connect, otherwise the session is terminated. Once they are mutually authenticated, the Client and the Agent will establish an encrypted tunnel.

BlastShield uses AES-256-GCM for stream encryption. The GCM block mode adds authentication to the encryption to ensure that the ciphertext hasn't been tampered with. The initial handshake and key negotiation signs all messages using ECDSA (Elliptic Curve Digital Signature Algorithm), with the NIST approved curve P-256. The key negotiation for the stream encryption is done using ECDHE (Elliptic Curve Diffie-Hellman Ephemeral) and HKDF (HMAC-based Extract-and-Expand Key Derivation Function). All generated symmetric keys are completely independent of the node's master key, a concept known as Perfect Forward Secrecy.

BlastShield uses wolfSSL's wolfCrypt FIPS Ready cryptographic engine which is FIPS 140-2 Level 1 validated. BlastShield also follows the FIPS enforced best practices of default entry point, and power on self test. wolfCrypt uses a hardware root of trust, if present, to generate public-private key pairs. Leveraging a hardware root of trust allows BlastShield to meet NIST SP 800-63 Authentication Assurance Level 3, the highest level of authentication assurance.

Figure 8: Secure Encrypted Tunnels

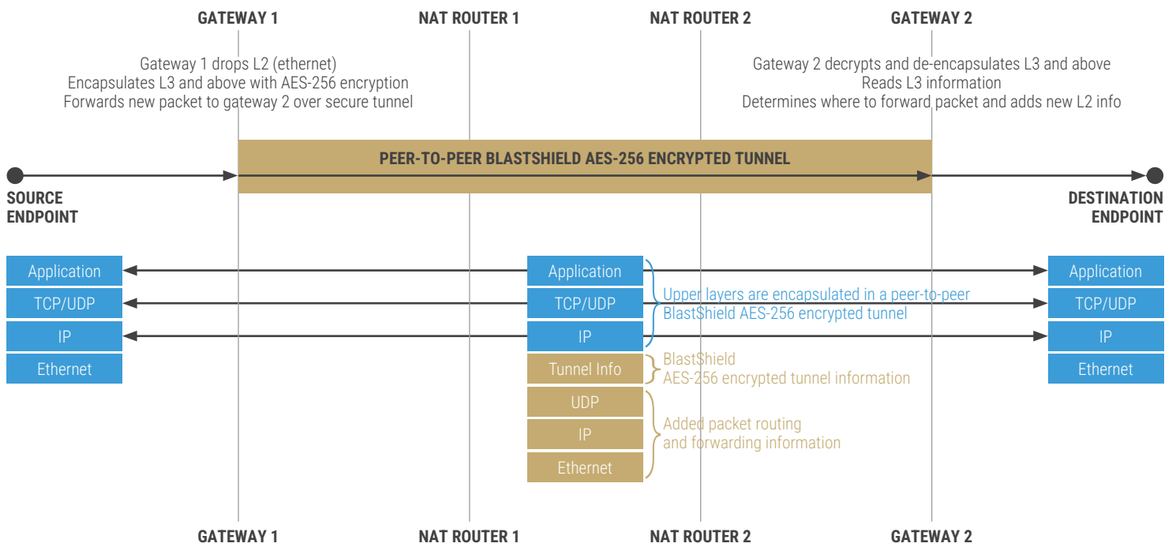

Granular Access Controls and Microsegmentation

The BlastShield Orchestrator simplifies the configuration of access controls and microsegmentation. BlastShield eliminates the need to design subnets and set up firewall access control lists (ACLs), VPN gateways, and VPN clients. Setting up access controls and microsegmentation is simple and intuitive:

Create Users and Agents
Using the Orchestrator, create Users and Agents and install the Client and Agents onto devices.

Organize Users and Agents into Groups
For example, you could add Users to Groups for developers, HR, sales, marketing, etc. You could also add Host Agents and Gateway Agents to Groups for Dev, QA, ERP, sales-apps, etc.

Set up Policies to provide granular access
Create Policies in the Orchestrator that allows Groups to communicate to or from another group. For example, you could set up a policy that allows the HR Group of users to communicate to and from the ERP Group of applications and machines. You could set up a different policy that allows a developers Group to communicate with the Dev and QA Groups of resources.

Tune microsegmentation by setting up Services
The Orchestrator allows you to set up Services that can be associated with Policies. Services in Orchestrator refers to application layer protocols or ports that can be customized to further limit access defined by Policies.

Configure Proxies to limit access to cloud applications
You could also set up Proxies to limit access to cloud applications such as Azure, Google Workspace, Salesforce.com, Hubspot, etc. To do this, you simply create a Proxy with domains to which you want to limit access (e.g. outlook. office.com, app.hubspot.com, lightening. force. com, etc.). BlastShield will allow you to set up a Proxy that presents a single public IP address to the domain that can be used for conditional access. Unlike other proxy solutions, the BlastShield public IP address is not visible to the public Internet using scanning tools, making the Proxy invisible to the Internet except when a User initiates a connection to an application associated with the domain configured in the Proxy.

Figure 9: BlastShield Orchestrator

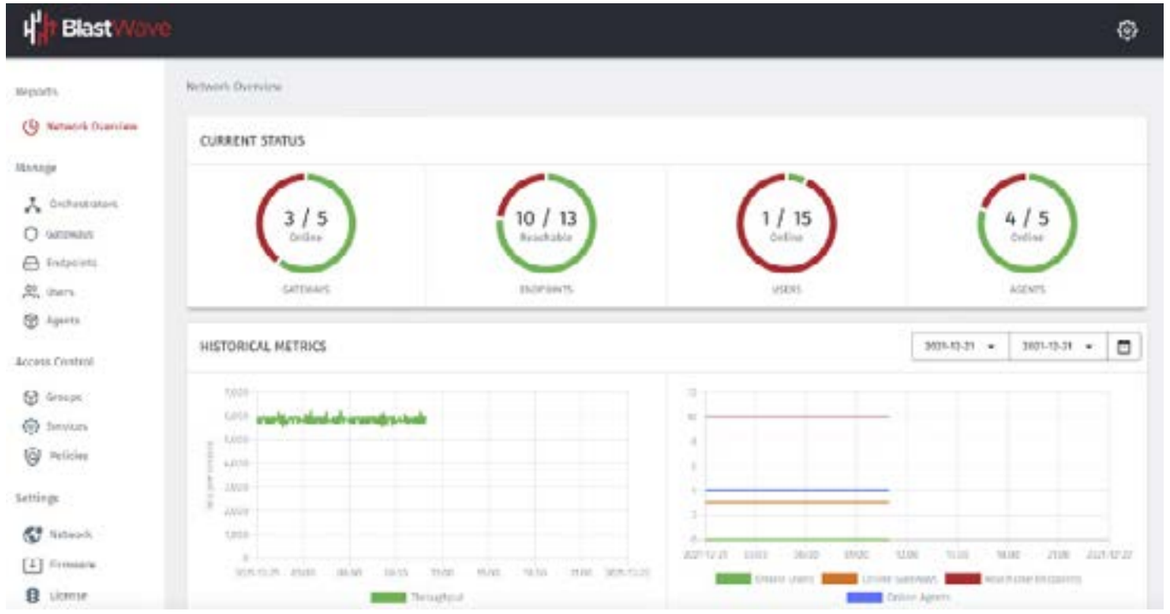

Device Invisibility

A device installed with a BlastShield Host or Gateway Agent does not offer any publicly available IP addresses, making those devices and gateways invisible to unauthenticated bad actors. IP scanning of a BlastShield network by unauthenticated users will show no devices to attack.

Hardware-based VPNs, cloud-based proxy servers, and secure access service edge (SASE) solutions all expose the public IP addresses, making them easy targets for DDoS and surveillance attacks.

Additionally, SASE solutions decrypt the traffic because payloads need to be in plaintext so that the data can be scanned before being re-encrypted and sent on to its destination.

Device invisibility allows BlastShield to protect against both internal and external attackers. Zero trust assumes breach and that an attacker is already in your network. An attacker that compromises an asset on your network will be unable to move laterally to compromise assets that are protected by BlastShield.

Figure 10: ZTA BlastShield Security and Invisibility

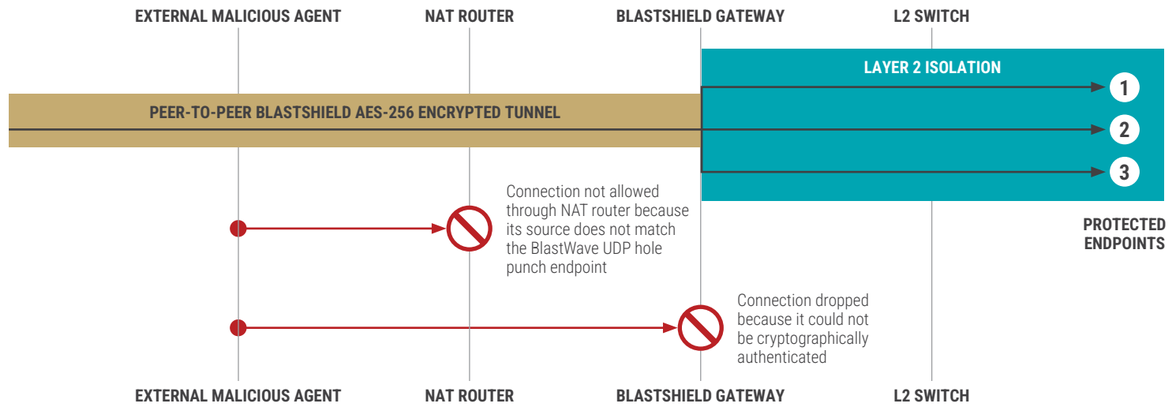

Peer-to-Peer Full Mesh Connectivity

BlastShield supports peer-to-peer, full-mesh connectivity, allowing Clients to communicate directly to any Agent. Additionally, any Agent can communicate directly to other Agents to create a full-mesh site-to-site network.

There are many benefits of BlastShield's peer-to-peer model vs alternative VPN and cloudbased SASE models. BlastShield's model is simpler to manage and is higher performance. In a recent report by the Tolly Group, BlastShield was found to be up to 34x faster than other providers.

Figure 11: Client-to-Application ZTNA Performance (according to Tolly Report, September 2022)

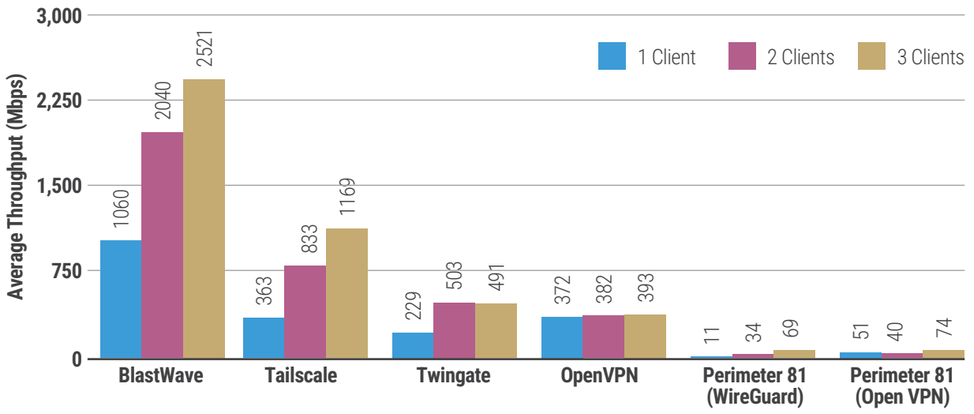

BlastShield Maximum Throughput
- 34x faster than Perimeter 81
- 6.6x faster than OpenVPN
- 5.1x faster than Twingate
- 2.2x faster than Tailscale

Cloud Application Proxy

BlastShield can provide secure access to SaaS applications such as Microsoft 365 and Salesforce.com by using a cloud-hosted BlastShield Host Agent acting as a proxy server.

The domains to be secured are configurable in the Orchestrator, and it is possible to proxy different sets of domains for different users for micro-segmentation purposes. User access attempts are logged and may be exported to syslog for monitoring. Access to the SaaS applications is secured using BlastShield's phishing-resistant MFA.

Figure 12: SaaS Proxy Agent

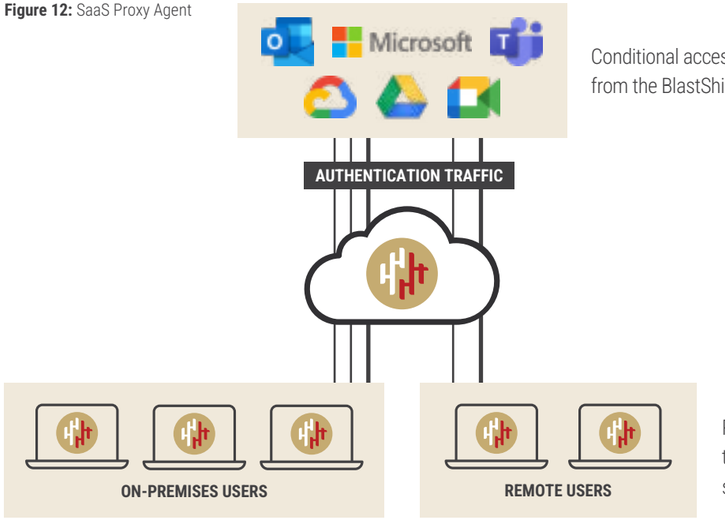

Conditional access controls only allow traffic from the BlastShield protected network

Policy controls only allow authorized users from the BlastShield protected network to access the secure SaaS proxy

Cloud and On-Premises Orchestration Platform

BlastShield's cloud-based Orchestrator enables administrators to add or edit Users, Agents, Groups, Policies, Services, and Proxies. The Orchestrator can also integrate with identity management platforms such as Azure AD, Okta, Ping Identity, and SIEMs.

Figure 13: BlastShield ZTNA with On-Premises and Cloud Orchestrators

BlastShield provides Zero Trust Network Access
Ensures that only authenticated users can access specific servers and resources

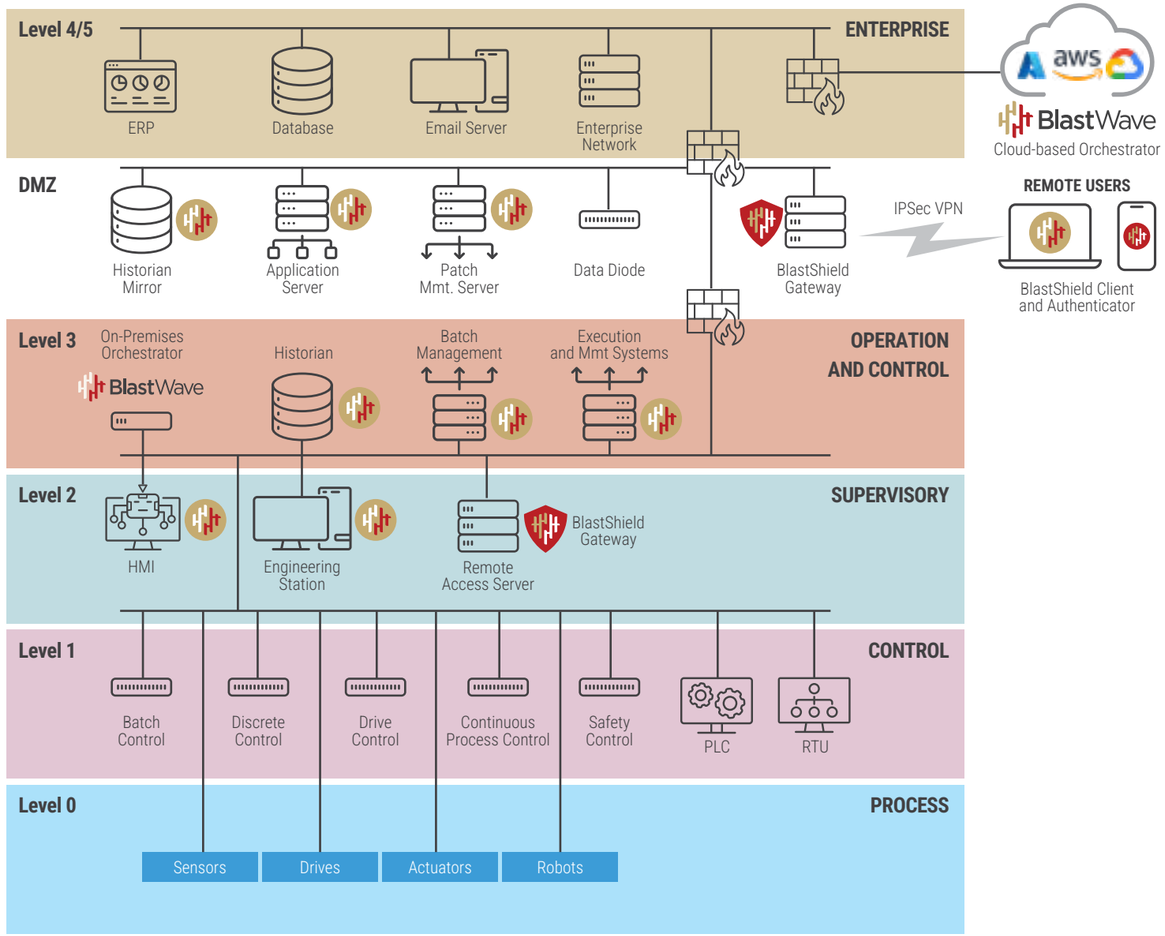

Technical Specifications

The BlastShield ZTNA solution is an advanced implementation of a software-defined perimeter (SDP). Here is a summary of the technical specifications of the solution.

Table 1: BlastShield Technical Specifications

Technical Specification	Description	
BlastShield Supported Operating Systems (as of November 3, 2022)	**Client:**	MacOS 10.13 or later; Windows 10 or later; most modern Linux distributions running on either a 64-bit Intel CPU or 32-bit or 64-bit ARM CPU; iOS version 13 or later; Android version 7 or later
	Authenticator:	iOS 13 or later; Android 7 or later
	Host Agent:	Ubuntu 18.04.6 LTS, 20.04.4 LTS, 20.04.5 LTS, 22.04 LTS, and 22.04.1 LTS; Debian 10 (Buster) and 11 (Bullseye); CentOS 7 and 8; Amazon Linux 2; Raspbian GNU/Linux 10 (Buster); Fedora 35 and 36; Windows Server 2012, 2016, 2019, 10; MacOS 10.13 or later; 64-bit
	Gateway Agent:	x86 and 32/64-bit Arm CPUs are supported
BlastShield Storage and Memory Requirements	**Client:**	100 MB storage, 50 MB RAM6 MB
	Authenticator:	6 MB
	Host Agent:	3 MB (Linux) to 100 MB (Windows), 50 MB
	Gateway Agent:	2 GB storage, 1 GB RAM
Cryptographic Engine	WolfSSL WolfCrypt (FIPS 140-2 FIPS-ready)	
Supported Algorithms	Symmetric (encryption): AES-256-GCM stream encryption. ECDHE and HKDF for key negotiation Asymmetric (authentication): ECDSA with the NIST-approved curve curve P-256, ECDHE, HKDF	
Supported Phishing-Resistant MFA Approaches	1) Passwordless authentication using the BlastShield Authenticator Mobile App with a biometric 2) FIDO2 security key with a passcode	
NAT Traversal Scheme	UDP hole-punch	
Topology Supported	Peer-to-peer full-mesh	
Connection Pairings	• User to application • User to machine • Machine to machine (site-to-site) • Machine to user	• Machine to application • User to user • User to cloud application (BlastShield SaaS Proxy)
Performance (Max Throughput)	**Client-to-application:** • 1,060 Mbps (1 client) • 2,040 Mbps (2 clients) • 2,521 Mbps (3 clients)	**Site-to-Site:** • 2,670 Mbps
Device Invisibility Enablement	Delete inbound security groups on devices installed with BlastShield agents. Public IP addresses will not be visible to the public Internet and therefore not scannable.	

Implementation and Support

Implementation of BlastShield consists of installing the BlastShield Client on user devices and installing the agents on other machines or devices that need protection.

Get Started with BlastShield

Download the BlastShield Authenticator The BlastShield™ Authenticator is a mobile app that verifies your unique identity. Download the app from your phone using the links below. You will need to give the app permission to access your camera to scan QR codes.

Download the BlastShield Client

The BlastShield Client allows you to access a BlastShield Network and launch the Orchestrator to manage Users, Agents, Groups, and Policies. Download the BlastShield Client onto your computer or smart device.

Sign-up for BlastShield

Once you've installed the BlastShield Authenticator App and Client, it's time to set up a BlastShield account at https://blastwave.com.

Authenticate and Connect to BlastShield

Now that you have registered, you can connect to the BlastShield network using the Desktop Client on your computer and the mobile Authenticator App on your phone. You will be asked to authenticate your biometrics on your phone to complete the authentication process.

Use the following steps to connect:

a. Launch the Client, select Mobile App as the Authentication Method.
b. Launch your BlastShield mobile authenticator app. Scan the QR code
c. Select Add New in the Desktop Client and click OK
d. Apply your face ID or fingerprint on your phone when prompted to authenticate.
e. You are now connected to the BlastShield network.

Launch the Orchestrator

To launch the Orchestrator, simply click on the 'Launch Orchestrator' button in the Desktop Client. This will trigger an additional authentication step with a QR code scan and biometric check using the Mobile Authenticator app. Please remember that you must be connected to BlastShield before you can launch the Orchestrator. See the videobelow to learn how to launch the Orchestrator.

The Orchestrator will launch in your web browser at https://orchestrator.blastshield.com. The Dashboard will be displayed as shown below.

Add BlastShield Agents for Secure Remote Access (instructions for Linux) The BlastShield™ Authenticator is a mobile app that verifies your unique identity. Download the app from your phone using the links below. You will need to give the app permission to access your camera to scan QR codes.

Step 1: Add a New Agent in the Orchestrator

- Click on "Agents" in the "Manage" menu in thethe left sidebar, then click the red "Add New Agent" button at the top right.
- The New Agent dialogue opens. Add a name for the Agent and a DNS Hostname. The DNS Hostname is optional and can be used to identify the Agent in the BlastShield network as BlastShield runs its own DNS.
- Then click on the red "Save and Download Invitation" button and choose the option for "Save and copy Linux/macOS installation command to the clipboard". Click on that option to copy the command.

Step 2: Install and register the agent

- Open a terminal session on the Linux server where you are going to install the Agent
- Paste the command you just copied to the terminal and hit enter. This will start the software download.
- The software will automatically install and run. The Agent will then automatically register with the Orchestrator. When the process is complete, you will see the following message in the terminal window: "Installation successful, the agent IP address is <Agent IP address>."

Step 3: View the status of the Agent

- Now that the installation and registration processes are complete, your Agent is up and running.
- You can check the status of the Agent by typing the following: sudo systemctl status blastshield
- The logs may be viewed as follows: a. sudo journalctl -u blastshield.service
- The status of the new Agent on your server should appear as "Online" in the Orchestrator as shown in the image below

For more detailed instructions on implementation for Linux, Windows and MacOS, visit support.blastwave.com.

BlastWave also provides professional services and phone support.

Orchestration and Management

The BlastShield Orchestrator is used to create, modify, and remove all Users, Agents, and Policies within the BlastShield Network. Only users with authorized privileges can access and use the Orchestrator. As discussed earlier in this document, the Orchestrator acts as a ZTA policy engine (PE) and policy administrator (PA). Rather than configuring subnets, ACLs, and VPNs, BlastShield makes it easy to set up granular access by creating and managing Users, Agents, Groups, Policies, and Proxies.

To learn more about orchestration and management, visit support.blastwave.com.

Figure 13: BlastShield ZTNA with On-Premises and Cloud Orchestrators

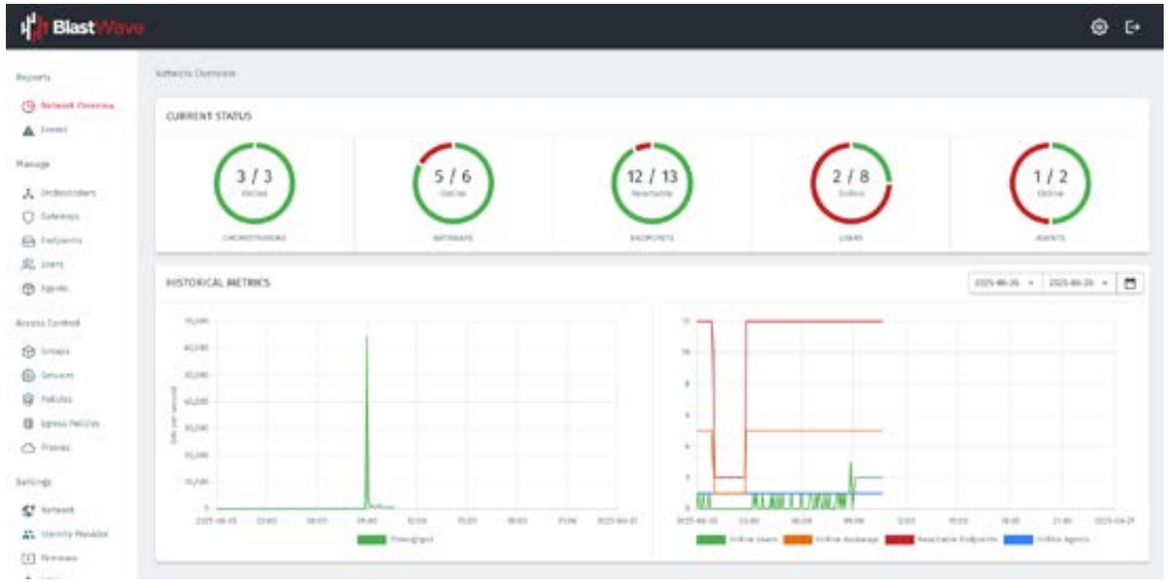

Use Cases

BlastShield is deployed in a number of use cases to replace VPNs and simplify secure access. BlastShield is flexible and can be deployed across IT, OT, and IoT environments. Because our solution does not rely on network location or network-based perimeters, we enable the creation of SDPs that address the need for better security controls, faster networks, and simpler management.

Table 2: BlastShield Use Cases

Use Case	Description
Secure Remote Access	BlastShield improves the level security of remote access by making devices invisible to unauthenticated bad actors and enforcing phishing-resistant MFA to authenticate users. BlastShield integrates with identity management platforms to enable SSO and simplify onboarding of users. Based on a recent zero trust performance study by the Tolly Group, BlastShield is ranked as the fastest ZTNA solution on the market, demonstrating max throughput 34x faster than other competitors. BlastShield can replace antiquated VPN hardware with a solution that supports direct full-mesh connectivity.
Site-to-Site Communications	Site-to-site communication implemented using IPsec tunnels is complex to manage due to complexity of the subnetting, ACL, and VPN configuration management. At the same time, IT assets that used to reside in an enterprise data center within the enterprise network perimeter are no longer constrained. The convergence of IT, OT, and IoT is dismantling the notion of a location-based perimeter. BlastShield enables the creation of a software-defined perimeter (SDP) network that can support direct, peer-to-peer, full-mesh connectivity. Rather than leaving VPN gateways and firewalls exposed to the Internet, BlastShield makes devices invisible to the public Internet.
Protecting Legacy OT Infrastructure	Industrial asset owners, operators, and OEMs are struggling to protect legacy systems, including engineering workstations, and unsupported Windows systems. BlastShield Gateway Agents can be deployed on gateway appliances in OT environments to protect downstream legacy systems that sit in the DMZ, supervisory, and operations levels without impacting the control and process levels.
Privacy of IoT Data-in-Motion	IoT devices such as IP cameras, drones, and a variety of OT control systems are being connected to the network at an increasing rate. As businesses aim to generate new revenue or reduce their cost of operations, BlastShield can help by supporting the installation of secure agents on Windows or Linux-based IoT devices or in front of the devices in the form of a gateway. BlastShield is easier to manage than other systems that require broad digital certificates and PKI management. Our approach simplifies the typical User, Agent, and Policy management by allowing admins to manage Groups, Policies, and Proxies
Privileged Access Management	Administrative access to business-critical systems often requires privileged access that is more strict than access controls designed for standard users. Privileged access management (PAM) may be based on network location and attribute-based information. The movement to zero trust will make PAM solutions irrelevant because they do not comply with the requirements of zero trust architecture. BlastShield provides granular access controls and authenticates the user with phishing-resistant MFA. Because our SDP approach does not rely on network location or user attributes to make user access decisions, BlastShield is an available replacement for PAM.
Accessing Cloud Applications and B2B SaaS	Cloud Access Security Brokerage (CASB) services can be overkill for simply enabling conditional access to third-party SaaS applications. Rife with monitoring, reporting, access controls, and confusing configuration options, CASB is also prone to misconfiguration and human error. Additionally, CASB approaches still leave the public IP addresses of the proxies and firewalls exposed, making them easy targets for attackers. BlastShield's SaaS Proxy Agent allows organizations to easily support conditional access and hides the public IP address of the proxy, making it unscannable unless you are authenticated using phishing-resistant passwordless authentication.

BlastShield Licensing and Pricing

BlastWave licenses the use of its downloadable BlastShield software and access to its cloudbased Orchestrator on a per agent basis. Pricing is based on the number and type of devices protected by BlastShield, represented by the number of installed Active Clients, Agents, and Gateways. BlastWave charges an annual licensing fee for each device depending upon the type of agent installed.

Our licensing includes access to features such as ZTNA, phishing-resistant MFA, single sign-on (SSO) support, microsegmentation, cloud orchestration, gateways, REST API, and optional on-prem orchestration.

For more information on our pricing,contact us at **www.blastwave.com**.

Conclusion and Recommendations

Industry trends and Federal Government mandates are driving adoption of zero trust architectures (ZTA). Federal agencies, enterprises, and industrial companies would be wise to learn more about zero trust concepts and develop strategies to migrate to a ZTA. Recognizing the importance of improving cybersecurity controls, the U.S. has directed many security-related agencies to facilitate sweeping changes in the way they manage cybersecurity. The U.S. Office of Management and Budget issued a memorandum that set a deadline of FY 2024 for agencies to comply with sweeping new cybersecurity guidelines based on the NIST SP 800-207: Zero Trust Architecture and CISA Zero Trust Security Model Guidance.

BlastWave's BlastShield Zero-Trust Network Access (ZTNA) solution can help organizations accelerate the migration to ZTA. BlastShield uses a software-defined perimeter (SDP) approach to ZTA that combines phishing resistant MFA, simple orchestration, granular access controls, peer-to-peer full-mesh networking, and device invisibility to provide a level of security that goes beyond zero trust standards. BlastWave radically simplifies the security stack and enables zero trust without sacrificing performance or cost.

To learn more about BlastWave and BlastShield, visit our website **www.blastwave.com**.

BlastWave's OT Protection Solution

BlastWave delivers a comprehensive Zero Trust Network Protection solution to provide the best possible outcome for OT environments. With a unique combination of network cloaking, secure remote access, and software-defined microsegmentation, we minimize the attack surface, eliminate passwords, and enable segmentation without network downtime.

To learn more, come to www.blastwave.com

v20250626

About BlastWave

BlastWave securely connects Industrial Control Systems, Operational Technology, and Critical Infrastructure networks with Zero Trust Protection and delivers industrial-grade cybersecurity with consumer-grade ease-of-use. Visit **www.blastwave.com** to learn more.

BlastWave

1045 Hutchinson Ave.
Palo Alto, CA 94301 USA
T: +1 650 206 8499