RedShield Platform Technical Overview

Introduction

Modern web applications and APIs exposed to the internet are subjected to continuous cyber attacks, targeting a wide range of vulnerabilities and often impacting data confidentiality, system integrity, and service availability.

RedShield's platform is designed to mitigate cyber risk with a dual philosophy in mind: to specifically address vulnerabilities so that they cannot be exploited, whilst also detecting threat actors so that they cannot perform attacks.

By using this dual-pronged approach, RedShield is able to provide the world's most effective web security shielding service.

Underpinning this, RedShield provides a fully configured and managed platform for shielding applications at scale; delivering increased resilience, performance, and security.

Customer application traffic is transmitted between client and server through the RedShield dataplane; which comprises the networks and security systems which accelerate customer applications, mitigate vulnerabilities, and prevent malicious activity.

This paper describes the major components and capabilities of the RedShield platform.

Global Architecture for Resilience and Performance

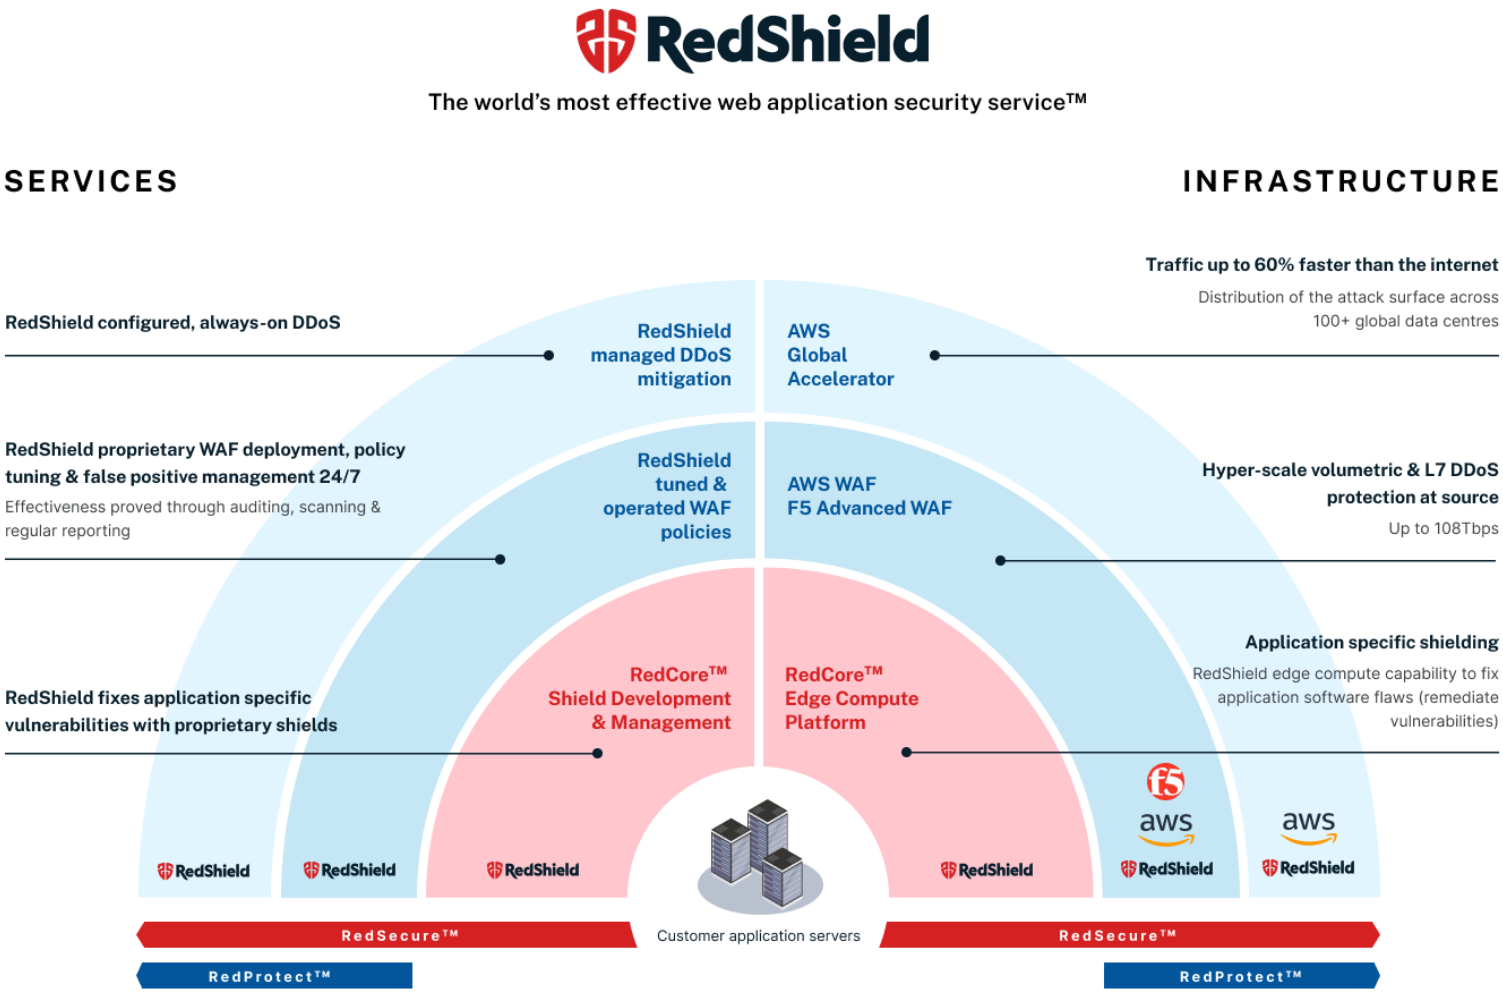

AWS Global Accelerator

RedShield services include a fully managed global DDoS mitigation and traffic distribution network provided in partnership with AWS.

Global Accelerator is a networking service that improves the performance of application traffic by up to 60% using Amazon Web Services' global network infrastructure. When the internet is congested, AWS Global Accelerator optimizes the path between clients and servers to keep packet loss, jitter, and latency consistently low.

Global Accelerator is an anycast network which spans global end points, and allows RedShield to block volumetric DDoS threats at source. Using a massively distributed network of over 100 datacenter edge locations globally, attacks are mitigated close to the threat source, across an elastic defensive edge, providing maximum uptime and reliability in the face of rapidly changing global threats.

Without AWS Global Accelerator

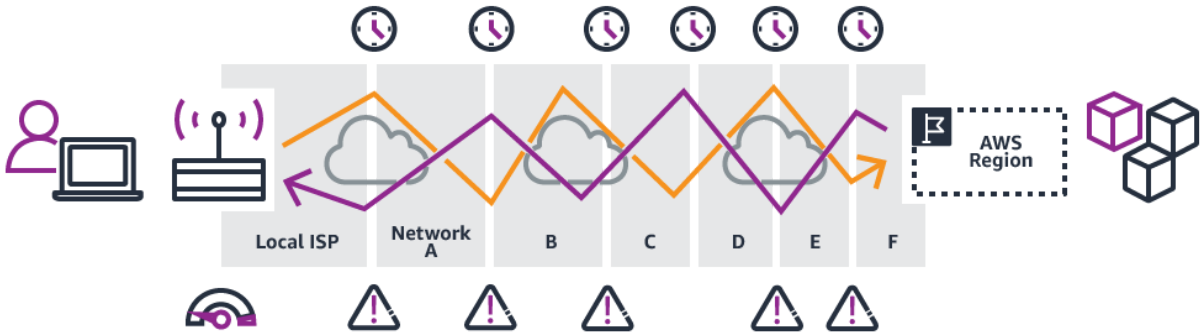

With AWS Global Accelerator

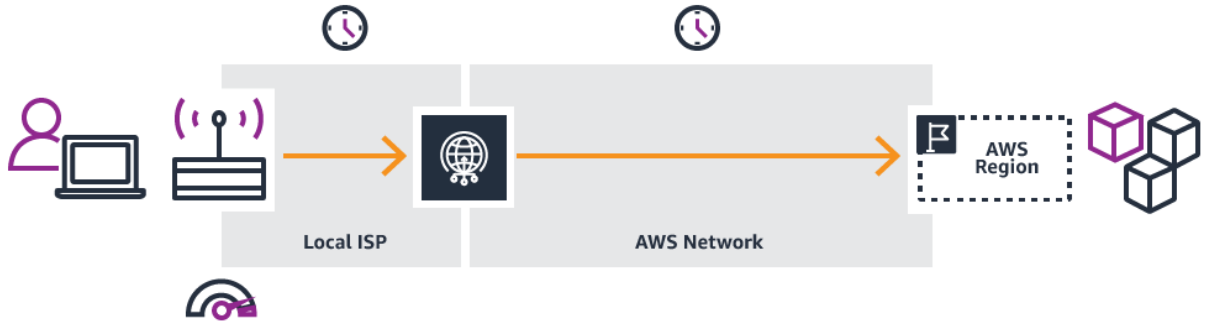

Multi-Vendor Layered WAF

Configured, tuned and operated WAF policies

Historically, WAFs have been difficult to configure and operate; and have suffered from annoying false positives; leading many organizations to disable blocking, and use ever more simplistic policies, which provide less and less protection for vulnerable applications.

RedShield services include high security policies deployed by security experts; tuned, updated and fully operated on an ongoing basis; customized to each application. Services

include a false positive helpdesk to ensure that customer applications benefit from maximum security levels, free from end-user experience issues.

WAFs are sourced from market leading vendors and deployed in layers to maximize each application's defensive posture.

Amazon Web Services WAF

AWS WAF is serverless and scales automatically, and is used in the RedShield dataplane to mitigate threats at ultra-large scale including layer 2-7 denial of service, and other types of attacks from very large botnets.

RedShield's service includes a fully configured and tuned AWS WAF policy in-path for all customer applications, as well as serverless proxy layers which provide scalable and resilient TLS encryption; to shield applications from encrypted attacks, and attacks targeting the TLS layer itself.

F5 Networks Advanced WAF

When fully configured, tuned and operated by RedShield, F5 Networks' Advanced WAF provides extremely comprehensive and detailed policies, and is used in the RedShield dataplane to provide protection for applications which have known or discoverable vulnerabilities, and preventing breaches by hackers and advanced persistent threats (APTs).

RedShield's automated systems and service for tuning and managing F5 Advanced WAF policies for each application, mean that a very high level of security is achieved at scale, whilst preventing false positives and responding quickly to new and emerging threats.

RedShield provides a highly responsive support service including a direct-to-end-user false positive helpdesk, to ensure that our customer's legitimate users are never impacted by overactive security controls.

Compliance and Security

End-to-end encryption

TLS encrypted application traffic is processed by a serverless, scalable TLS encryption proxy layer, which handles both attacks targeting the TLS layer itself, and encrypted attacks which attempt to bypass traditional defenses.

Geo-constraint of TLS keys and service data

Trusted proxying of encrypted traffic within the dataplane is performed with controlled geographic isolation of key storage, in compliance with common customer security requirements.

Compliance frameworks

RedShield assists customers in achieving compliance with major frameworks and regulations including:

- ISO/IEC 27001:2013
- PCI DSS
- SOC 2
- GDPR
- CCPA
- HIPAA

Other major frameworks are in progress, and some are listed on RedShield's website.

Customers requiring assistance with specific compliance frameworks and regulations can check https://www.redshield.co/security for further information on RedShield's compliance program, or contact security@redshield.co for further information.

Resilience and Availability

Application traffic is handled by AWS' globally distributed platform, with continuous monitoring, automated scaling and deep redundancy within key components.

RedShield dataplane can also assist with application server redundancy requirements, providing global and local failover and load balancing, with server and network path monitoring, maintenance and automated error pages.

RedPipe - Private Datacenter Connectivity

RedShield also offers RedPipe, a secure DDoS-resistant way to publish your web applications to the internet. RedPipe achieves this by leveraging peering with Kordia's communications infrastructure to offer a protected, private channel for communications between the RedShield cloud and your web application servers.

This ensures requests routed through the RedShield service can reach your web application servers, even if your internet facing links are saturated or down - keeping your web applications available in the event of a DDoS attack targeting the back end datacenter.

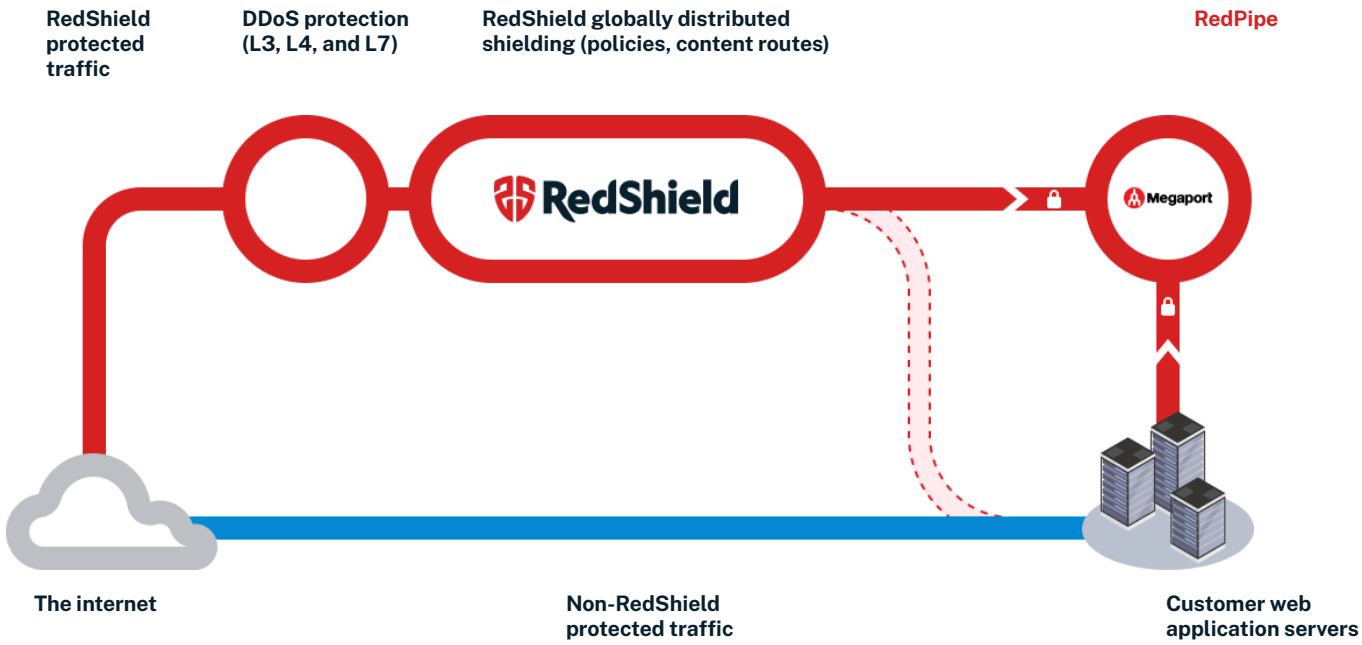

A key point to note is that while the aim of the RedPipe product is to deprecate the need for customers to expose their web application services directly to the internet, the customer may choose to either:

- Fully replace connectivity between protected web application servers and end users on the internet with RedShield RedPipe - however, other service not ; or
- Retain connectivity over the Internet as a backup in case of a disruption on Kordia's service.

Where the customer has other non-HTTP/HTTPS services that must be exposed to the internet, RedShield may also be able to route this via the RedShield service to provide DDoS and IP layer protections. Examples of this may include DNS, SSH and FTP protocols.

Application Vulnerability and Threat Mitigation

Layer 2-7 DDoS Mitigation at Hyper Scale

RedShield provides mitigation of DDoS attacks automatically, using always-on defensive measures to keep applications safe from even the very largest attacks at a global level.

Blocked Category
Volumetric DDoS • Globally distributed anycast edge network • Serverless, elastic TLS proxy layers • Traffic Flood detection • High capacity attack absorption • IP based attack mitigation • Always-on defense, blocking at source • Low-latency inspection of encrypted attacks
Layer 4-7 DDoS • TCP middlebox reflection • Request floods - ultra-large scale • Half open L4 • L4 Connection limits per IP • L4 Connection limits per URL • SSL renegotiation limits • Slowloris, slow POST, slow HTTP variants detection • Large objects, random/cache bypass attacks • Nuisance clients - account registration, message posts • Application behavior under server congestion

Web Application and API Vulnerability and Threat Mitigation

RedShield provides shielding for web applications and APIs using a dual-pronged strategy: firstly by eliminating known or discoverable vulnerabilities, using specific shield configuration to ensure that vulnerabilities are no longer exploitable, and secondly by detecting malicious actors when attacks occur, and preventing threats from reaching web application servers.

Blocked Category
RFC violations • TCP protocol compliance checks • TLS/SSL protocol checks • General HTTP RFC compliance checks • Cookie not RFC-compliant
WAF evasion techniques detected • Directory traversals • %u decoding

- IIS backslashes
- IIS Unicode codepoints
- Bare byte decoding
- Apache whitespace
- Bad unescape
- Parameter splitting *(on request)*
- Application logic attacks *(advanced shielding on request - exploits which typically bypass all WAFs)*

HTTP protocol compliance failure

- Header name with no header value
- Several Content-Length headers
- Chunked request with Content-Length header
- Bad multipart parameters parsing
- No Host header in HTTP/1.1 request
- CRLF characters before request start
- Content length should be a positive number
- Bad HTTP version
- Null in request
- Check maximum number of headers
- Bad host header value
- Check maximum number of parameters
- Mandatory HTTP header is missing

Input violations

- Brute Force: Maximum login attempts are exceeded
- Illegal method
- Illegal redirection attempt
- Request length exceeds defined buffer size
- Failed to convert character
- Illegal static parameter value

Additional blocking elements

- Illegal URL
- Modified WAF cookie

Negative security signatures

Over 5,000 signature checks are performed at line-rate against each HTTP request and server response; with per-application custom policy tuning and ongoing signature management included in every shielding service.

Signatures are frequently updated with support for emerging threats and high priority CVEs.

- Abuse of Functionality
- Authentication/Authorisation Attacks

• Buffer Overflow • Command Execution • Cross Site Scripting • Denial of Service • Detection Evasion • Directory Indexing • HTTP Response Splitting • Information Leakage • LDAP Injection • Non-browser Client • Other Application Attacks • Path Traversal • Predictable Resource Location • Remote File Include • SQL Injection • Server Side Code Injection • Trojan/Backdoor/Spyware • Vulnerability Scan • Xpath Injection
Server response transformation • Insertion of HSTS headers • Insertion of Same Origin Headers • Insertion of secure cookie flag • Insertion of HTTP ONLY flag • Removal of Server Banner

Advanced Shields

RedShield provides services to mitigate specific vulnerabilities which may be discovered during penetration testing or vulnerability scanning.

These shields are code objects which run in the RedShield dataplane, providing customers with a serverless edge compute capability, together with the engineering services to deploy, support and manage shields at scale in production.

RedShield has many thousands of shields available in the library, to be used for specific customer applications.
Some brief examples of shield objects are listed below, of the large number we have available, and the virtually endless possibilities for new shield creation on-demand.

Shield example: Credit card anti-fraud check

At the time when a user submits a credit card number for payment, this shield checks the credit card number against Accertify to determine whether the card is lost or stolen, or being used in a fraudulent manner; and blocks the payment if they don't accept it.

It is simple to integrate with most web applications, requiring payment path details, and the name of the parameter which contains the credit card number.

Shield example: SAP CVE-2020-6308

This shield prevents the Server Side Request Forgery associated with SAP CVE-2020-6308. The shield was developed using the PoC found here: https://github.com/InitRoot/CVE-2020-6308-PoC . Customer paths and parameter names are supported in cases where the server is configured in a special way.

Shield example: Add Google v2 reCAPTCHA to web page

Google v2 reCAPTCHA may be added to any web page, in order to prevent automated traffic from interacting with the application.

Shield example: Files may be uploaded containing malware

This shield uses multiple third party commercial antivirus engines to check if an uploading file is malicious. This shield is configured to apply to specific upload paths within the application.

Bot Mitigation

RedShield dataplane includes security controls which target bots and automated attackers. Default measures are applied to all customer sites, and more intensive controls are activated and tuned by RedShield on request for applications which experience particular threats, such as web scraping, and other undesirable automated client activity.

Blocked Category
Standard bot mitigation • Trusted Bot IP, reverse lookup & device fingerprint • Untrusted Bots IP, reverse lookup & device fingerprint • Non-browser client detection • Aggressive IP Reputation • Requiring UA header as fair use • TLS step up

Advanced bot mitigation

Available by request

- Automated behavioral analysis and fair usage enforcement
- CAPTCHA and javascript responses to suspicious requests
- Rapid surfing and non-human browser detection
- Account takeover, credential stuffing, scraping and brute force mitigation
- Tripwires and tarpits
- Multi factor authentication

Banned Malicious IP Addresses

RedShield's dataplane automatically bans IP addresses which are known to belong to malicious threat sources; effectively preventing the worst offending bots and hackers on the internet from connecting to vulnerable systems.

Continuously updated IP address lists are sourced from RedShield internal threat logs, as well as open source and commercial IP address list providers, and applied to all customer applications, to ban high risk IPs from all services.

Blocking of banned IP addresses is active by default across all customer applications.