

Introduction

As the leader in SaaS web security, Zscaler’s focus is to provide policy-based secure web access for any device, anywhere. Through ground-breaking innovations in its massively scalable cloud architecture, Zscaler Web Security Cloud provides an ultra-low latency SaaS security solution that needs no hardware and no software. Three major factors are driving organizations around the world to adopt Zscaler:

- **Mobility:** Laptops, road warriors and home offices are impossible to protect with appliances in the enterprise DMZ. Rather than having a perimeter around the office, Zscaler creates a global perimeter around the Internet.
- **Complex Threat Environment:** Employees cannot be protected from Web threats with just a URL filter. Active content, Social Networks, and Ad Networks have resulted in every major site hosting malware at some point in time. Web 2.0 risk mitigation requires a technology that can do inline scanning of every page for complex threats, and blocks sensitive content from leaving the organization. It is impossible for IT to keep up with constantly evolving threats. A managed security service that can shield vulnerabilities and block zero day exploits is essential.
- **Consolidation, Simplification, and Reduced Cost:** Powered by unique technologies like NanoLog, PageRisk, ByteScan and a true multi-tenant architecture, Zscaler’s comprehensive cloud service offers advanced security, Web 2.0 controls and data leakage prevention (DLP). Enterprises can simply point their traffic to any of Zscaler’s 40+ data centers and be risk free. Administrators have a single console to manage policy and analyze any transactions in real-time for all offices as well as the mobile work force. There is no hardware or software to deploy, and no agents to install. Simplification and economies of scale enable IT to mitigate risk at half the price.



Figure 1: Zscaler Architecture for Comprehensive Web-security and Management

This document provides a technical overview of the Zscaler infrastructure, enabling technologies, Zscaler services, unified policy and reporting capabilities, as well as traffic forwarding and authentication methods.

Zscaler Infrastructure

Legacy vendors offer security as a cloud service by creating data centers that host racks of appliances traditionally used within enterprise perimeters. This is not only costly, but the enterprise is bound to only a few data centers of the provider where their policy is hosted. With each appliance capable of handling only one enterprise, latency and performance in a multi-tenanted environment depend on how many organizations share the appliance. In addition, transaction logs are spread across many data centers making it difficult to see details in real-time.

Zscaler's greatest achievement is the architecture that was created from scratch to take advantage of being a pure cloud provider, while delivering a truly multi-tenant and highly scalable platform for deep security. The fundamental innovation is in functionally distributing components of a standard proxy to create a giant global network that acts as a single virtual proxy so that any user can go to any gateway at any time for policy-based secure internet access. Zscaler infrastructure comprises three key components; Zscaler Enforcement Node (ZEN), Central Authority (CA) and Nanolog™ (Transaction Log) Servers as shown below.

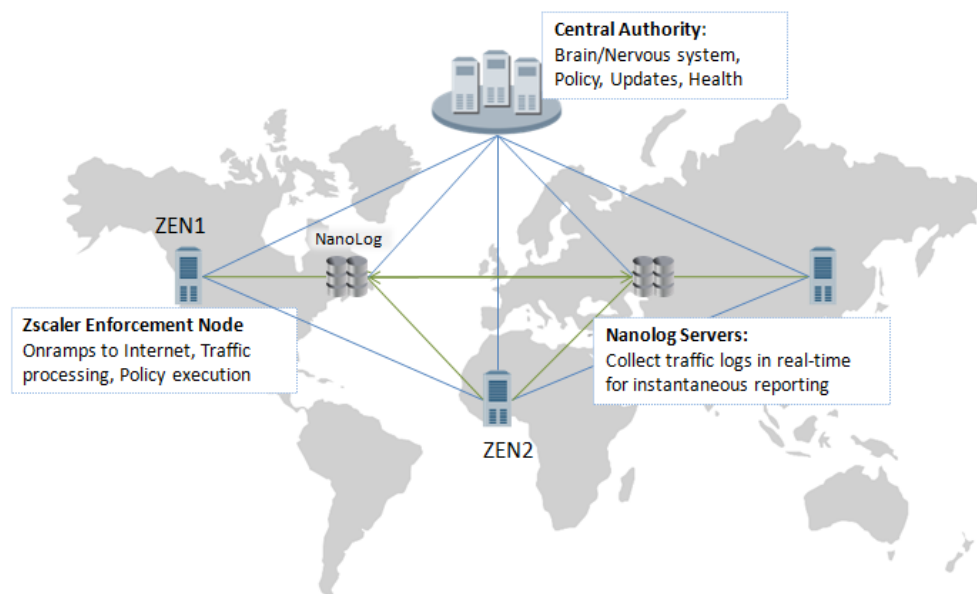


Figure 2: User-policies stored in the CA are enforced on ZEN with logs stored on Nanolog™

An enterprise forwards all web traffic to the nearest enforcement node (ZEN). Policies governing the user's access to any website are served by the CA, and enforced on the ZEN. All transaction logs are stored in a centralized Nanolog™ server for real-time retrieval. Any user can access any ZEN, and all components have multiple levels of redundancy to ensure high availability.

Zscaler Enforcement Node (ZEN)

Powered by over thirty patents, each ZEN is a fully featured inline proxy that enforces corporate secure, manage and comply policies at user-level granularity. The ZEN incorporates a hardened custom-built OS and a custom TCP/IP stack to deliver 90% of transactions in less than 90 microseconds. Zscaler's ByteScan™ technology enables each ZEN to scan every byte of the web-request, content, responses and all related data for inline blocking of threats like viruses, cross site scripting (XSS), and botnets. This capability also enables Dynamic Content Classification (DCC) of unknown sites. By scanning each page the ZEN computes a PageRisk™ index for every page loaded and enables administrators to control content served to their users based on acceptable risk.

The ZEN also incorporates Zscaler's unique authentication and policy distribution mechanism that enables any user to connect to any ZEN at any time. This enables enterprise to simply point traffic to any ZEN and ensure full policy enforcement while getting all reports back in real-time.

Central Authority (CA)

The central authority complex is the brain of the Zscaler cloud. The CA manages and monitors all nodes and ensures that they are always up-to-date with the latest real-time feeds, software and synchronized for cloud wide intelligence of threats. The CA directs users to the closest ZEN through DNS and PAC file resolution to ensure minimum latency when users are on the road. Through its multitenant architecture, the CA provides each organization with its own portal to self manage their policy. Any change to the policy is communicated to the ZENs within seconds. The CA provides an end user authentication framework through integration with Secure LDAP or ID Federation systems.

The central authorities are a globally distributed peer-to-peer cluster with an automatically elected master. This ensures all cloud components can always talk to a central authority even if there are major internet outages that isolate an entire region.

Nanolog Severs

Backed with multiple patents, Zscaler's Nanolog technology enables administrators to access any transaction log almost instantly. Each ZEN uses the Nanolog technology to perform lossless compression of logs by a factor of 50:1. These logs are transmitted every second to the Nanolog servers over secure connections. Logs are multicast to multiple Nanolog servers for redundancy. Through an innovative reporting and database framework created specifically for web logs, the Nanolog server can support 15 million logs per second. This technology provides an administrator with real-time reports and capability to query (within seconds) complete transaction level details for any user, department or location at any time. Each server has over 16 Terabytes of capacity enabling Zscaler to provide multi-year data retention.

Zscaler Cloud Services

Zscaler offers four cloud-based subscription services that help customers secure and manage every aspect of their employees' web usage. Zscaler Secure, Manage, Comply and Analyze are offered in Basic, Advanced and Premium price-bundles, and can be deployed without any new hardware or software. Zscaler delivers wire-rate performance with microsecond latencies with all four services enabled concurrently. Zscaler customers can focus on creating their security and web usage policies, while Zscaler executes their policies at the finest level of granularity across the enterprise.

Secure

Zscaler Secure services consist of Zscaler Anti-virus and Anti-spyware, Advanced Security, and Web Access Control.

Zscaler Anti-virus and Anti-spyware (AV/AS)

Zscaler provides an inline, ultralow latency AV/AS solution that protects users from file based attacks. Files of any size including multi-level archives are scanned in real-time. Blocking malware in the cloud is instantaneous and universal, saves bandwidth costs, and obviates the need to patch endpoints or multiple appliances to effectively protect users. Zscaler's enforcement nodes also scan every web page for embedded viruses, malicious javascript or advertisements that may lure users to download fake AV software. Compared to an inline UTM, IPS or firewall, Zscaler – being a true proxy – provides superior protection because the entire file is downloaded, assembled, uncompressed and scanned for millions of viruses before it is delivered to the end user. By sharing intelligence of infected as well as clean content across the entire cloud, the ZENs are able to deliver fully scanned files without the user noticing any latency. In addition to its inline AV/AS, Zscaler also uses multiple commercial AV/AS engines concurrently in an offline mode, to uncover any additional threats.

Web Access Control

Old browsers and plug-ins with known vulnerabilities are the easiest attack vector. Zscaler allows administrators to be proactive with security by enforcing users to use only acceptable browser types and versions. Zscaler also warns users when they are using browsers or plug-ins with older patch levels if the versions they are using have known exploits in the wild.

Advanced Security (Zero-Day, Botnets, Phishing, P2P, Browser Vulnerabilities)

Web has become dynamic with active content on every site, and well reputed sites sourcing content from various sites including numerous advertisement networks. Hackers now use SQL injection attacks not to deface websites, but to surreptitiously insert hidden iframes that infect unsuspecting users through drive by downloads. It is impossible for administrators to block "bad sites", because most of the reputed sites have served malware at some point in the recent past either by getting compromised or using an infected advertisement network. Most websites serve dynamic content that changes based on the user's login or surfing history. The only way to protect the user is to really inspect exactly the content they are being served EVERY TIME, and not rely on content scanned by a crawler in the cloud. Asking another appliance to scan content doesn't work because the appliance does not have the user's password.

Zscaler's advanced security uses the ByteScan technology to scan every byte of all requests and responses. This enables Zscaler to detect hidden iframes, cross site scripts, signs of phishing attempts, cookie stealing, and botnet command and control. This unique capability also enables Zscaler to block anonymizers hosted in Facebook even if Facebook itself is allowed.

For each page served, Zscaler computes a PageRisk Index that takes into account use of suspicious techniques like javascript obfuscation and zero pixel images. This information is correlated with other factors such as GeoIP-based location of the website and its reputation to compute a dynamic risk index. Administrators can then define a policy to block content beyond their acceptable risk level.

Manage

Zscaler's Manage services consist of URL Filtering, Web 2.0 Application Control and Bandwidth Control.

URL Filtering

Zscaler's URL filtering leverages multiple databases to ensure global coverage. You have complete flexibility to override any classification done by Zscaler. You can add/remove URLs from Zscaler's predefined 90 categories, 30 super categories and 6 classes to keep reporting and rules consistent. The three level hierarchy makes it easy to create rules and look at reports at the appropriate granularity. Web access policy can be set for specific users, groups, and locations or any combination. Actions include the ability to block, caution, allow, or provide time-based access and quotas in terms of time or volume. Zscaler's inline Dynamic Content Classification (DCC) engine ensures pages are classified by their content if the URL is not sufficient.

Web 2.0 Application Control

Web 2.0 is changing how enterprises do business. The challenge with Web 2.0 is to ensure it helps business, but does not hinder productivity. For example, it is desirable that marketing can post a video on YouTube, while others can only view the posts. Zscaler provides granular control over Web 2.0 applications like webmail, streaming media, social networking and instant messaging. Zscaler understands the protocols used by individual Web 2.0 applications to provide granular control over actions such as posting an update to LinkedIn. The administrators can configure policies by users or groups to enable organizations to take advantage of the latest Web 2.0 platforms without compromising productivity or security of critical data.

Bandwidth Control

90% of enterprise traffic is on HTTP/HTTPS. The web is used for business as well as recreation. A few users watching a movie excerpt on Hulu can completely disrupt a customer presentation on Webex. IT administrators need the ability to control bandwidth by specific web applications rather than ports and protocols.

Zscaler supports bandwidth control by application "class". Administrators can define a transaction to fall in a bandwidth class based on a diverse set of criteria. Large file transfers from any fast website can clog bandwidth. Depending on the size of file being transferred, the transaction can be put in a bandwidth constrained bucket. The same policy can be applied to streaming media sites. On the other hand, transactions to business applications such as Webex, LiveMeeting, Salesforce or Netsuite can have a minimum reserved bandwidth regardless of other web traffic. Bandwidth control policies can be set per application class, in terms of the maximum sessions as well as maximum and minimum bandwidth per class, location, and time of day. Bandwidth policy is enforced without dropping any packets, by leveraging Zscaler's custom TCP/IP stack that enables modulating the throughput on each side of the ZEN proxy.

Comply

Zscaler is industry's first fully integrated cloud based web DLP solution. Employees can now send out data over webmail, social networks, blogs, or instant messaging. Zscaler's enforcement nodes scan every byte leaving an organization to look for sensitive data. It provides a variety of dictionaries and engines for enterprises to enforce compliance policies and protect its Intellectual Property (IP).

DLP Dictionaries

Zscaler provides three forms of dictionaries; special dictionaries that identify a specific type of number or content type, Artificial Intelligence (AI) engine based dictionaries that identify types of documents, and phrase based dictionaries

- **Special Dictionaries:** Special numbers such as credit cards (with full checksum validation), social security numbers (SSN), Singapore NRIC, Canadian Social Insurance Numbers, etc.
- **AI Dictionaries:** Types of documents such as financial statements (balance sheets, cash flow statements, income statements, etc), medical statement, source code (C, Java, etc), documents saved from Salesforce.
- **Phrase Based Dictionaries:** Custom dictionaries created by administrators containing company specific keywords (e.g., Zscaler Confidential). Zscaler's dictionaries use fuzzy matching techniques to ensure phrases match regardless of capitalization, spacing and noise words.

DLP Engines and Rules

A DLP engine combines one or more dictionaries to match specific compliance requirements. Zscaler provides sample templates for HIPPA, GLBA and PCI. Administrators can create their own engines by combining dictionaries. For example, a DLP Engine containing dictionary for Source Code with custom phrases for copyright notices can prevent Intellectual Property leakage.

All content leaving the organization is scanned. In addition to HTTP postings, Zscaler can decode all Microsoft documents, PDFs and content inside a zipped file. Granular rules can be specified for applying specific DLP Engines to a group of users, web application types or locations.

Analyze

Nanolog technology enables administrators to get complete transaction level detail in real-time for any user located anywhere and on any device. Zscaler's unique real-time reporting capability enables administrators to instantly drill into any sequence of events as they unfold. For forensic investigations, administrators can instantly access specific transactions that occurred at any time in the past, filtering by user, time, location, url categories, malware types, transaction size, and a number of other parameters. Zscaler's base package provides six months of detailed log storage and two years of summary data at individual user granularity and over 1500 different views.

Traditional web filtering solutions require administrators to install databases or third party correlation tools to manage their web logs. Zscaler's cloud service takes away the burden of log management by providing detailed log archiving capability that can be expanded to as long as 10 years.

Zscaler, Inc.
392 Potrero Avenue,
Sunnyvale, CA 94085
USA

+ 1 408.533.0288
+1 866.902.7811

sales@zscaler.com
www.zscaler.com

Zscaler®, and the Zscaler Logo are trademarks of Zscaler, Inc. in the United States. All other trademarks, trade names or service marks used or mentioned herein belong to their respective owners.