

Zscaler Cloud Security Offers Secure Web Gateway Features

By: Matthew Sarrel

The recent 3.0 release brings together just about every security feature that could be offered via proxy in a straightforward GUI, yet several flaws remain.

Zscaler's self-named cloud security service provides organizations with security covering integrated Web, instant messaging, peer to peer, Webmail and SMTP-based e-mail, and it does so without any on-premises hardware or software installation requirements. Rather, the Zscaler service spreads its proxy and relay load across the company's 40 data centers and presents administrators with a rich, flexible, Web-based management interface.

Based on my tests of Zscaler's security services, I'm impressed with its potential to protect users from HTTP- and SMTP-based threats, which are more prevalent than ever—particularly due to the increased use of blended malware threats. (For example, a phishing e-mail drives a user to a site that plants malware on his PC via a browser exploit.)

The service's policy configuration and reporting options stand out for their depth, and I expect that the Zscaler service should provide a strong complement to companies' existing end-point and perimeter security solutions. And, as a SAAS (software as a service)-based offering, Zscaler provides additional value by enabling administrators to inspect potentially malicious traffic on the Web, rather than allowing it to reach their network perimeter, or, perish the thought, their endpoints.

Zscaler's cloud security services are sold in a number of different editions, starting from a basic Web filtering edition that's priced at \$1.50 per user per month for 100 users, with

discounts available at higher volumes. For more information on the available editions and included functionality, check out the data sheets at tinyurl.com/23lwjtx.

Zscaler in the Lab

As a hosted service, Zscaler is very easy to install. All I really had to do was change the administrator password, agree to the terms of service and remember to save my policy changes. I noticed Zscaler guarantees only 99.99 percent uptime monthly, although a representative said "outside of the regularly scheduled maintenance windows affecting just the admin-console, Zscaler has delivered 100 percent availability since launching in August 2008."

Zscaler uses bandwidth from multiple carriers and maintains dedicated space in multiple data centers operated by different providers. There's a little more work required to integrate with a directory service or import users. Plus, browsers need to be configured to use the Zscaler service as a proxy; firewalls also need to be reconfigured to allow only Web traffic to and from the Zscaler proxy.

The admin interface launches with a clean, easy-to-read dashboard with prominent, context-sensitive options for Logout, Support, Getting Started, Help and Concept. Concept (and the little icons that the company tells me are light bulbs) is interesting and brings up a Flash demo that shares somewhat helpful information that was obviously developed by someone on the marketing, rather than

technical, staff.

Help opens in a new window and is pretty useful except that it lacks an index or search capacity. Clicking Support took me to a page where I could submit a trouble ticket. I clicked Getting Started, and a new window popped up that listed configuration steps and provided links to walk through them rapidly. I easily uploaded the eWEEK logo and customized end-user notification messages for when sites or files would be blocked.

The default security policy configuration is most likely acceptable for most organizations. I found it very easy to establish policies for inbound and outbound traffic inspection, scanning different file types and even whitelisting sites by URL where all content should be allowed.

Spyware also can be blocked by category. For example, I could allow password-stealers while blocking all others, although I'm not sure why someone would do that. Browser control is somewhat interesting. I could easily enforce policy to block older browsers with known vulnerabilities or simply block a browser entirely. In my testing, the service blocked about three quarters of the malware downloads that I attempted over http.

Under Advanced Threats, I found a lot of settings designed to address today's malware threats. There are settings for blocking botnet traffic to known command-and-control servers, ActiveX controls, known and suspected phishing sites, IRC tunneling, anonymizers, Cross-Site Scripting, and also

traffic destined to countries or regions. (It was preconfigured to block China, and I easily added Russia and Brazil.)

There also are extensive controls for allowing or blocking P2P file sharing such as BitTorrent and eDonkey, as well as P2P anonymizers such as Tor and P2P VOIP (voice over IP) such as Skype and Google Talk. For some strange reason, all this P2P stuff was configured to be allowed by default, but after a few clicks, I shut it all down. One thing I really liked is a reminder to save and activate changes: Far too many Web GUIs have allowed me to wander off a page without saving settings.

It was also very easy for me to set policies to block various content categories of Websites. It's possible to set different configurations for different locations, so I could block gambling sites from the office but allow them to be accessed from outside the office. I could also block or allow access to various Webmail sites. The same goes for streaming media sites and social networks and blogs.

Rules can be pretty complex. For example, instead of simply blocking Twitter, I could configure Zscaler to allow reading but not posting. However, content filtering worked about as well as it does with most of the other products in this category, meaning that the same weaknesses regarding identification of sites and correctly categorizing them by content and not by URL are present. For instance, blogs hosting image thumbnails of pornography are not correctly classified as pornography.

It started to get really interesting when I drilled down into bandwidth control features. The service comes preconfigured with seven types of application classes, including "general surfing" and "large files." It's also possible to add application classes.

Then, under bandwidth policy, I could allocate minimum and maximum bandwidth by application class. For example, I could allow 100 percent of bandwidth for Web conferencing but only 10 percent for streaming media. Although this was one of the most interesting features to play with, I was unable to assign bandwidth by user, which makes this feature moot because we all know that a security administrator could never subject the CEO to the same streaming media bandwidth rules as a regular employee. I worked around this by creating users and groups and then applying different bandwidth rules on a site-by-site basis.

Then I clicked on Administration and drilled down to Admin & User Accounts. Users can be forced to authenticate against a hosted user database, Microsoft Active

Directory or OpenLDAP. I could show an acceptable usage policy for every session, day or week—or never.

I liked the role-based system administrator options, which allowed me to limit access to the GUI and certain settings. I also liked the ability to define various real-time alerts, such as this: "If three virus download attempts are made within five minutes, then issue an alert via e-mail and/or RSS."

Under the Comply heading, I set DLP (data loss prevention) policies and enforcement options. Zscaler uses the term "dictionary" to describe a DLP rule. There are eight predefined dictionaries including "credit card leakage" and "social security leakage." I effortlessly created new dictionaries by clicking Edit, Add Dictionary and then entering strings to search for and their weightings.

Dictionaries are then grouped into DLP engines. For instance, the HIPAA (Health Insurance Portability and Accountability Act) engine contains the medical information leakage and Social Security leakage dictionaries.

It all comes together under Compliance Policy, where I enabled or disabled engines, set the order in which they should run, and assigned users and applications where I could apply the rules. My attempt to send a message on Facebook containing a Social Security number was logged and reported accurately.

Analyze gives you very close to real-time inspection of traffic, which can be sorted by user or transaction, and then filtered by department, location, URL classification, security threat and the amount of time you want to include in the report. I could easily see that my test user had spent the last half-hour browsing for barbecued brisket recipes. As a forensics tool, this is very helpful for answering the "what-in-blazes-just-happened" question.

Flexible reporting

Reporting is a strong suit for Zscaler. One of my favorite features is the ability to set any report as a Favorite, then organize Favorites and select them directly from the dashboard. Reporting is very flexible. I could slice and dice, subset and analyze, double-click on just about anything and get more detail. Any report can be generated as a PDF simply by clicking the little PDF icon next to the report title.

It's important to track Web activity on a per-user basis, which most companies will do through integration with LDAP or AD. I created user accounts and selected to force authentication, but this did not actually take place until I turned on authentication under

the gateway settings. (Incidentally, there's also an "enable bandwidth control" on the same screen.) In most cases the Zscaler administration GUI provided the ease and power I needed, but in cases like this, it left me stranded.

SMTP services, new in version 3.0, provide a similarly comprehensive, multitiered array of inspection and mail delivery services. As expected, anti-malware/spam/phishing services inspect mail before it reaches your e-mail servers. E-mail and Web security options are shown right next to each other; this integrated management ensures greater security-policy consistency than if policy were managed through multiple products.

Spam filtering worked fairly well: Settings are done via a slider (dial up the spam!), but I found them to be too general. Settings can be tweaked on a domain, user or group basis, but they can't be tweaked for content. So it's just dialing up spam versus dialing up subprime mortgages and dialing down Russian brides.

Over a 24-hour period, the default configuration delivered more spam than I would've liked, but there were no false positives. I could probably slide the dial around to find just the right spot, but users don't have access to their spam settings. More mature anti-spam products allow users to tweak their settings and access their quarantines.

It's possible to define "inflow" and "outflow" e-mail content policies. I easily configured all mail addressed to info@mattsarrel.com to be delivered directly to me. DLP works as well with e-mail as it does with Web traffic. I prevented myself from sending an e-mail full of Social Security numbers to an external account. The e-mail was blocked and I received a custom e-mail explaining why it was blocked.

Zscaler can also perform gateway-to-gateway SMTP encryption and delivery assurance. Reporting for SMTP security services is as helpful and easy to use as for Web security.



Zscaler, Inc.
392 Potrero Avenue,
Sunnyvale, CA 94085
USA
+ 1 408.533.0288
+ 1 866.902.7811
sales@zscaler.com
www.zscaler.com