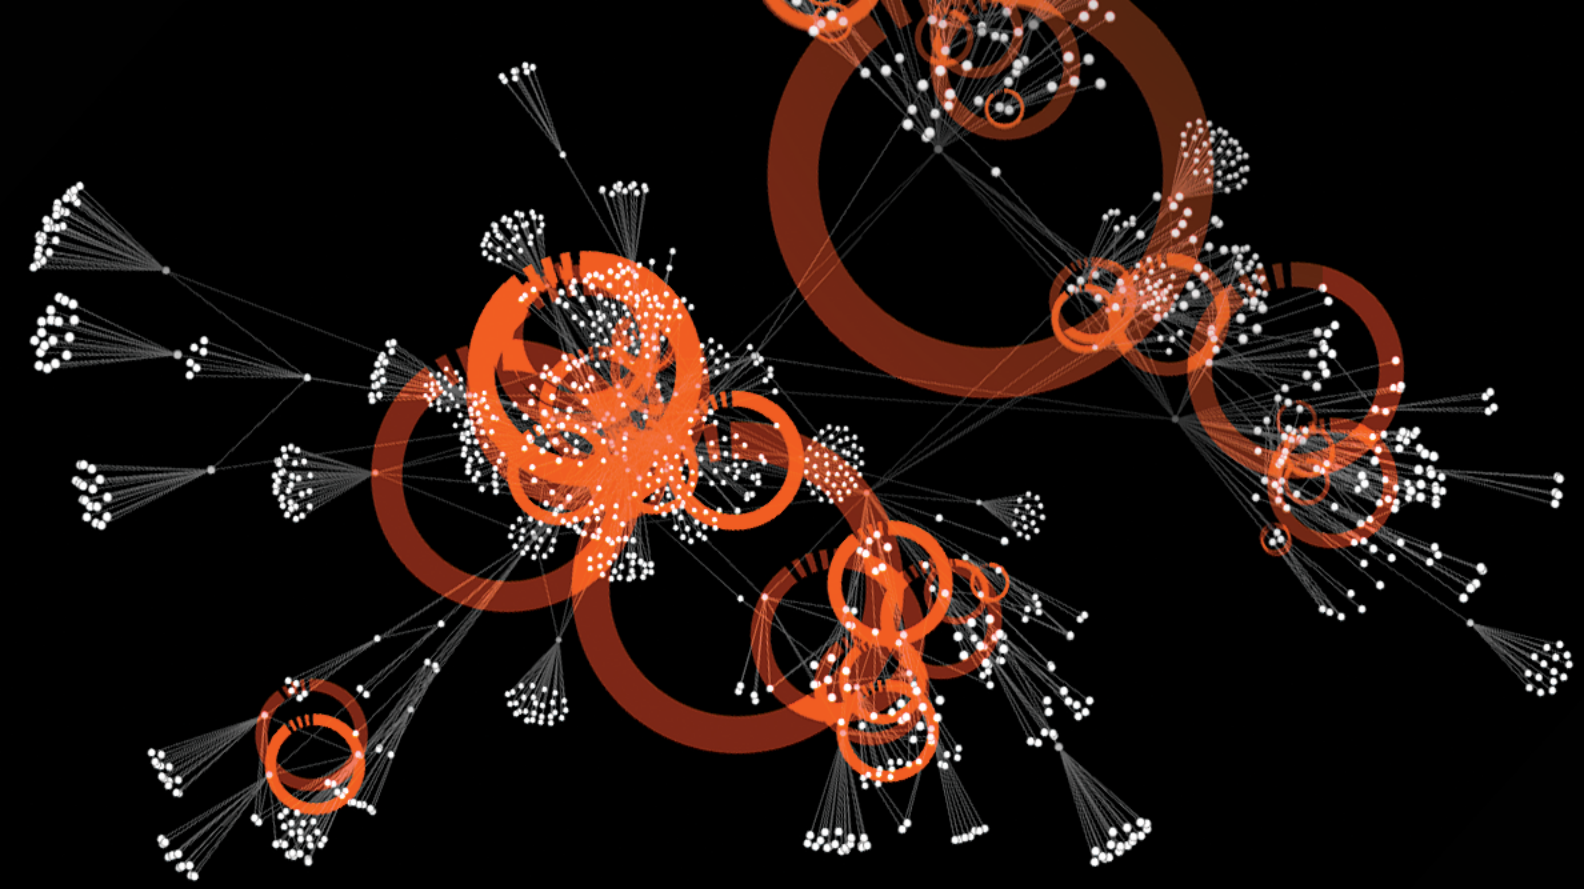

OpenDNS

Predictive Intelligence:

Automated Protection Against Advanced Attacks

Threat actors operate globally but most security tools act locally. OpenDNS uses its bird's eye view of Internet activity and data analysis systems to identify emergent threats before an adversary puts your organization in their crosshairs.

Hoping for the Best, Predicting the Worst

More than ever before, staying abreast of emerging threats to your IT environment demands that you look far beyond the boundaries of your network. While the perspective of your IT security team is necessarily local, your IT infrastructure is increasingly tied to mobile devices and cloud-based infrastructures that are not under your direct control. Moreover, threats to your organization are global. Well-financed and determined adversaries often back stealthy, targeted attacks aimed at individual users within your business and supply chain partners that you rely on.

Recognizing that existing security controls aren't reliable, industry experts counsel companies to rethink their security model while assuming the worst—malware and malicious actors will find a way through your defenses, if they have not already done so. Your IT security team needs to develop what the analyst firm Gartner calls an "Adaptive Security Architecture" that raises the bar for attackers who want to compromise your IT assets, while also nurturing "predictive" capabilities that can leverage large pools of information—"Big Data"—to identify evolving attacks and inform your company's response.

> *"Assessing current and potential future threats requires advanced threat intelligence that most organizations are not capable of developing cost-effectively in-house."*
> — Rob McMillan, Kelly Kavanagh
> "How to Select a Security Threat Intelligence Service" (G00206140)
> Gartner

To do that, however, you need reliable information about what is happening on your network and what is happening external to your network environment. Your IT security team needs to keep tabs on emerging threats to the Web-based and SaaS platforms your users rely on. They need to identify patterns of suspicious or malicious activity originating within and outside their network perimeter. Operationally, they need to be able to correlate global Internet activity to this local activity.

However, as the analyst firm Gartner has pointed out: few firms have the resources and know-how to conduct that kind of information gathering or even to extract useful information from the data they have.

A Science of Spotting Emergent Threats

OpenDNS combines two key elements for tracking emergent threats: exceptional visibility into global Internet activity and custom data analysis systems that enable us to extract vital information about emergent threats.

A recursive view of global Internet requests

The OpenDNS Global Network handles more than two percent of global Internet requests each day, including DNS requests that are a standard part of HTTP/S, FTP, P2P, and all other protocols.

We build our view of what's happening online via our network of 23 distributed data centers that resolve more than 50 billion DNS requests each day (as of May 2014). Each of those data centers maintains direct, BGP (Border Gateway Protocol) peering connections to the top-tier networks and content providers that make up the fabric of the Internet. We use these direct connections to ensure that DNS requests sent through our network are routed quickly, but they also give us a detailed picture of the interconnections between the various networks that make up the Internet.

On top of that, OpenDNS layers intelligence derived from the request patterns of 50 million users. From homes in more than 160 countries to enterprises in every industry vertical, nearly 2% of the world's daily-active Internet users send their Internet requests through our network. By aggregating data from billions of individual requests and using innovative data analysis systems to match DNS requests with intelligence gleaned from BGP routing tables, we can observe malicious or suspicious behavior across the Internet.

Finally, we layer data feeds from more than 200 partners on top of the data we collect directly. These third-party data sources include threat feeds from AV vendors and other security companies, as well as data from university, governments and independent researchers. All third-party feeds are run through our automated validation system and then added to our threat intelligence.

Breadcrumbs of data help identify attacks before they happen

By accumulating billions of pieces of discrete data from across the Internet, OpenDNS breaks the cycle of attack-detect-respond that has prevailed in the information security space for more than two decades.

We're able to spot activity consistent with malicious software and online scams soon after it occurs by observing significant sequences of events: web domains generated to host malicious infrastructure, spikes in traffic to those domains and patterns of traffic to and from other sites. Our technology can not predict the future exactly, but it can roll back the identification of new threats almost to the moment of their creation, while reliably predicticting how malicious infrastructure will behave in the near future. That ability gives our customers a leg up in defending their infrastructure from malware and new attacks.

How OpenDNS Classifies Threats

To create reliable security ratings for domains, OpenDNS tracks and measures a wide range of attributes of DNS infrastructures and IP networks that helps us to identify anomalies and calculate reliable risk scores.

What IP addresses and ASN are associated with the domain?

Our AnyCast routed network allows us to observe Internet activity at a global scale. Among other things, we track Autonomous System Numbers (ASNs)—groups of connected IP routing prefixes—and the IP addresses listed in DNS records for known-malicious domains, OpenDNS has identified a number of ASNs on the Internet that are almost entirely devoted to serving malicious traffic.

Is the domain the creation of a domain-generating algorithm (DGA)?

DGAs generate tens to thousands of pseudo-random domain name strings. Some of these become registered domains and actual websites that are used to host exploit kits, serve malware, or participate in botnet command and control (CnC) activities. These domains are active for only brief periods, and their DNS records often have a short TTL (time-to-live).

The randomness and large number of domains used and communicated make blacklist-based detection and blocking of domains ineffective. The use of DGAs is not new. The technique for botnet CnC has been noted since 2004. But the use of DGAs is new in connection with malicious software sites, and appears to many in the security community to be a growing trend.

Domains generated by DGAs are easy to spot with the human eye. However, using humans to cull out DGAs from millions of domains is naither practical nor scalable. OpenDNS Security Labs has pioneered methods for identifying the products of DGAs algorithmically using a wide range of strategies from lexical analysis to monitoring of domain WHOIS registration. High entropy and n-gram perplexity—measures of randomness in groups of Internet domains—are strong indicators that algorithms were used to create the domain name strings rather than human.

Is the domain routable?

As botnets and other kinds of malware increasingly rely on DGAs, we find more examples of requests for non-routable domains. That is, requests from endpoints for DNS records that do not exist—and may never exist. These kinds of requests are highly indicative of activity from an infected endpoint. We automatically flag the requested domain—even if it is not routable—while also adding the IP address of the requesting endpoint to our list of malicious requestors.

Unique Visibility of **Botnet Activity** at the DNS Layer

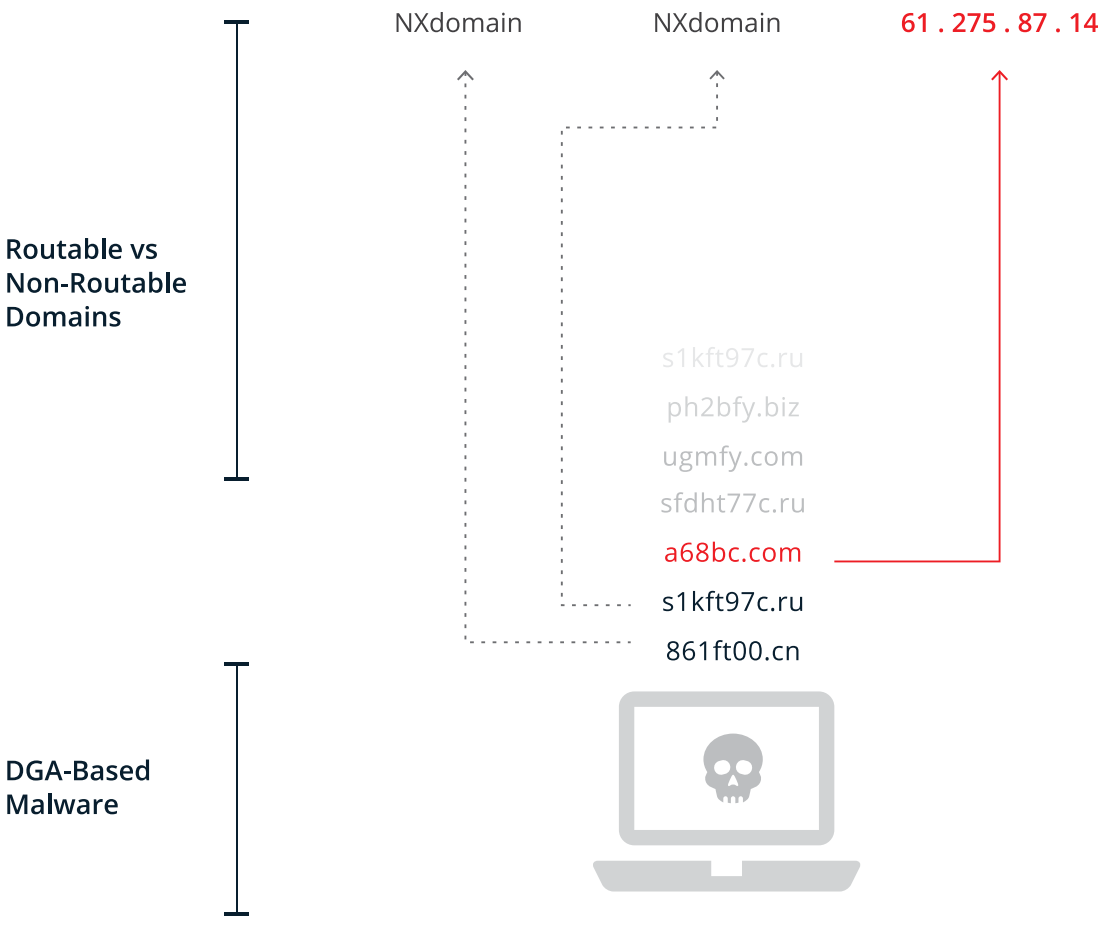

Using Data Analysis To Automate Protection

To make sense of all that data, OpenDNS researchers have created a range of data analysis systems to build a detailed map (or graph) of interactions between DNS requestors (individual endpoints) and requested hosts (Websites, FTP servers, P2P systems, and so on). These systems make it possible for OpenDNS to spot clusters of suspicious activity and track malicious incidents over time, flagging potentially malicious infrastructure even before it can be pressed into action by cyber criminals or other threat actors.

Automated Analysis Spots Malicious Infrastructure

OpenDNS's users provide us with a wealth of information about online behavior. To make sense of that data, we created OpenDNS Security Graph, an automated threat classification platform that comprises a number of algorithmic tools like machine learning, graph theory, anomaly detection, and temporal pattern mining. OpenDNS Security Graph automatically aggregates metadata associated with specific hosts and networks, scoring them as "good," "bad" or "indifferent." Rather than asking an army of threat researchers to manually analyze and rank sites, OpenDNS uses mathematical models, advanced classifiers, and artificial intelligence to make informed judgments about the trustworthiness of specific domains, networks, or IT assets.

OpenDNS is not the only company to use data analytics to help solve security problems. What makes us unique is our ability to leverage big-data analytics and a trove of information about users' online behaviors to make determinations about the reliability of IP networks and DNS infrastructure. Significantly, we can make determinations about the trustworthiness of IT assets and domains in the absence of malicious binaries, attack signatures, or even specific knowledge of malicious activity. That gives our customers an edge in blocking attacks against their networks, off-network devices, and users before they occur.

Exposing malicious domains

One way that OpenDNS is able to map malicious networks is by observing patterns of requests that happen immediately before and after requests for domains that we know are malicious. By observing and tracking this "co-occurrence" behavior, we identify which domains tend to get looked up together and assign them a higher co-occurrence (a.k.a. **C-Rank**) score. Similar to tracking co-occurrences, we also calculate a score based on how often domains are looked up in close succession. These links are calculated for each pair of domains. Mapping domain co-occurrence and links makes it easy to track the extent of a malicious network and expose other parts of a malicious command and control infrastructure that might not otherwise be apparent.

DNS address (A) records are also used to calculate reputation-based scores for the domains we resolve. The **RIP** (routing information protocol) score computes a reputation for domains based on the IP addresses linked to its A records. The **Prefix** score is computed using the entire subnet that the

IP address resides on, and the **ASN** score is computed based on the reputation for the autonomous system (AS) the subnet is associated with. Using A records, we can see if the IP address, subnet, or even the entire AS is suspicious—and to what degree.

Geographic diversity (or disparity) is also scored. The **Requester Geo-Diversity** score is calculated based on the number of queries from clients visiting the domain, by country. The **Sum Geo Distance** and **Mean Geo Distance** are used to score the total distance between the actual country location for the IP of the domain and the average distance between the IP addresses for a domain, respectively. This geographic scoring helps us track and influence the overall scoring of the domain or IP address.

SecureRank, an automated tool that ranks an Internet domain's trustworthiness or its "guilt by association," is also used to assign a score to each domain. At a high level, SecureRank notices if a particular domain is looked up by IPs that never look up malicious sites and gives those domains a higher (more trustworthy) SecureRank. If domain names are looked up by IPs that we've observed looking up malicious sites, their SecureRank goes down, indicating that they are less trustworthy. Under the covers, SecureRank comprises a number of link analysis methods, which are used to create a global map of transactions between hundreds of millions of endpoints and tens of millions of domains. On top of that map, we layer reputation data derived from the OpenDNS Security Labs as well as information on malicious activity culled from third party providers and our own user base. Patterns emerge. For example, by combining our traffic data with historical data on known malicious endpoints and malicious domains ("bad neighborhoods") and those with no history of malicious activity, we find that domains that are visited by known infected endpoints, but never by endpoints that we know to be "clean" and free of infection are more likely to be malicious domains.

SecureRank Uses **Bipartite** Graph Theory

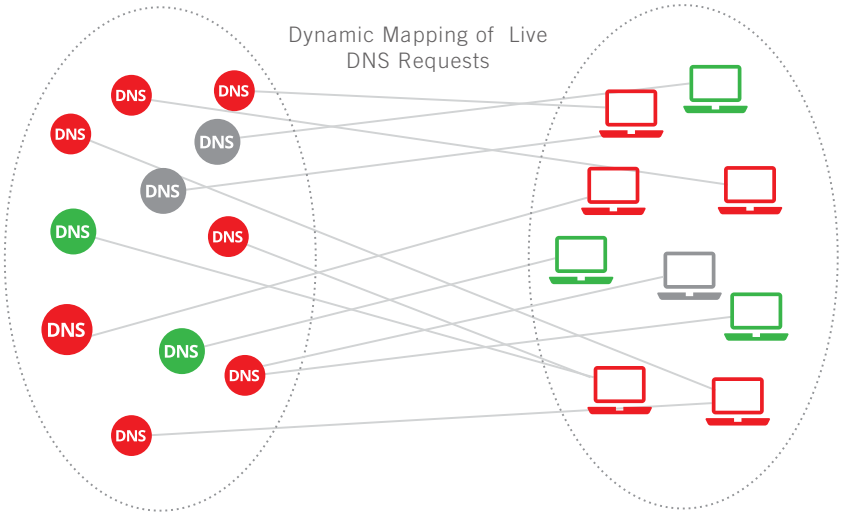

Using SecureRank in our back-end classification engine, we identify more than 200,000 domains each day that are malicious or likely to become malicious. Those domains include sites used as part of botnet CnC, drop sites used by data-stealing malware, and domains used in targeted "spear phishing" and "watering hole" attacks.

Using OpenDNS to Improve Incident Response

It is not enough to have great technology. In today's threat environment, any security solution also needs to be able to work in concert with a wide range of detection, alerting, analysis and incident response tools.

> *"By 2020, 60% of enterprise information security budgets will be allocated to rapid detection and response approaches—up from less than 10% in 2014."*
> — Peter Firstbrook, Neil MacDonald
> "Malware is Already Inside Your Organization, Deal With It" (G00259857)
> Gartner

As reports of fines linked to data breaches become a matter of public record, what is clear is that the harshest penalties are meted out not to organizations that experienced a breach—but to those that failed to adequately respond to the security incident once it was identified. In an age of long-lived, stealthy and targeted threats, incident response has become just as important as defense and attack prevention.

In just one measure of this trend, Gartner's "Adaptive Security Architecture" weighs incident response and recovery as much as traditional activities like prevention and detection. Given that organizations cannot stop all malware or thwart every attack, Gartner says investments should be balanced across four stages of security incident response: prevention, detection, response and prediction.

Using the investigative features provided by OpenDNS, customers get access to the fruits of our automated security analysis via a simple search interface. Just enter a suspicious domain name, an IP address, or ASN into the Web-based Investigate interface. Instantly view interconnections between suspicious or malicious infrastructures and patterns of Internet activity that correlate with other known threats—malware, bot networks, and phishing websites.

When analyzing suspicious hosts, algorithms like SecureRank make it simple to see—at a glance—whether a particular domain name or IP address is associated with safe or malicious activity. You can also study the list of co-occurring domains to see if the asset you are investigating is part of a malicious botnet CnC infrastructure or other online scam.

Investigating Security **Incidents** with OpenDNS

Beyond that, the Investigate interface makes it easy to explore complex IP networks of DNS infrastructures. By simply clicking a link, you can perform reverse lookups, which inform you what other domains are associated with the same IP address and any malicious or suspicious activity related to those domains.

Using OpenDNS in concert with other security tools

OpenDNS's security services are a potent aid when used as part of incident response. They become even more powerful when they are coupled with complementary technologies like real-time detection and security incident and event management (SIEM) platforms.

Umbrella is often the first product OpenDNS customers reach for when responding to security incidents. It gives them quick access to threat rankings and historical information so they can corroborate observations from their own environment with other incidents that have occurred or are occurring around the globe.

Industry-standard RESTful Web APIs allow developers to build queries against the Investigate interface to be conducted programmatically by other security tools like SIEMs and incident response tools. Furthermore, OpenDNS data can be pulled back into those platforms, speeding investigations and alerting with data that can help with attribution of threats and attacks on your network.

Bird's Eye To Bullseye: Targeting Threats to Your Environment

Your organization is being challenged as never before. Your workers are more mobile than ever. You manage an increasingly blended computing environment comprised of cloud-based and on-premises applications. Your job demands that you manage the security of your network, while also keeping tabs on fast-moving and elusive threats that target your users and data from outside your network perimeter. How do you know if you have become ensnared in a wide-ranging attack, or whether you are the sole victim? Which suspicious activity is worth investigating, and which is a distraction? Are the malicious actors targeting your organization also targeting other firms?

OpenDNS is a cloud-delivered network security service that provides easy access to predictive threat intelligence. We use data from the Internet's "phone book"—the Domain Name System—and the backbone of the "information highway"—the Border Gateway Protocol. By combining OpenDNS's view of Internet activity, derived from 50 billion DNS requests we receive each day from more than 50 million users around the world, OpenDNS allows you to investigate security incidents in a targeted fashion. Correlate incidents with high-risk threats in the public domain, and provide your staff with the ability to spot emergent threats to your IT infrastructure before they happen.

Visit us online to learn more about how OpenDNS technology can help your organization stay on top of emerging threats and attacks on your networks and employees.

Stopping Hacktivists in Their Tracks

In August 2013, readers who pointed their Web browser to the website of the New York Times (nytimes.com) found that, instead of the day's articles and columns, they were shown an image celebrating the Syrian Electronic Army (SEA)—a hacktivist group sympathetic to the regime of Syrian dictator Bashar al Assad.

What happened? Behind the scenes, the SEA used a targeted phishing attack on an Australian domain name registrar to gain control of the Times' DNS records. They redirected nytimes.com to the SEA web server's IP address, and changed the WHOIS information for nytimes.com to list the Syrian Electronic Army as the rightful owner of the domain.

The SEA wanted to make a political statement. But the attack could have been much worse. Even with temporary control over a high-traffic site like nytimes.com, the SEA could have configured their website to launch surreptitious Web-based attacks that plant data-stealing malware on visitors' computers.

As it happened, OpenDNS was able to protect our customers and users even before the attack occurred. How?

Mapping the Internet's dark alleys

Even before the attacks against The New York Times, the IP addresses and domain names used in the redirection attack triggered algorithmic classifiers within OpenDNS Security Graph, our automated threat classification engine. Specifically, it knew that the name server used in the attacks, 141.105.64.37, also hosted domains including malware and phishing websites, so we blocked access to that IP address as a matter of course.

A sudden change

Our global network and our direct links to core ISPs and content providers gives us access to DNS data from across the Internet. With that, we noted the sudden change in the IP address that the Times website resolved to. While it is not uncommon for websites to change their IP address, established web properties like nytimes.com rarely do it. The change was enough to generate an alert.

Good site, bad neighborhood

The combination of a sudden change in the IP address that nytimes.com resolved to, coupled with what we knew about the address—that it was hosted on an obscure Indian ISP and was associated with malware and phishing attacks was proof that something had gone badly wrong.

Resolution

We worked with the New York Times and other organizations affected by the attack (including Huffingtonpost.com and Twitter) to recover and get their domains resolving properly.

OpenDNS

OpenDNS provides a cloud-delivered network security service that delivers automated protection against advanced attacks for any device, anywhere.

Visit www.opendns.com/enterprise-security
to Instantly Start a Free 2-Week Trial

OpenDNS, Inc.
www.opendns.com
1.877.811.2367