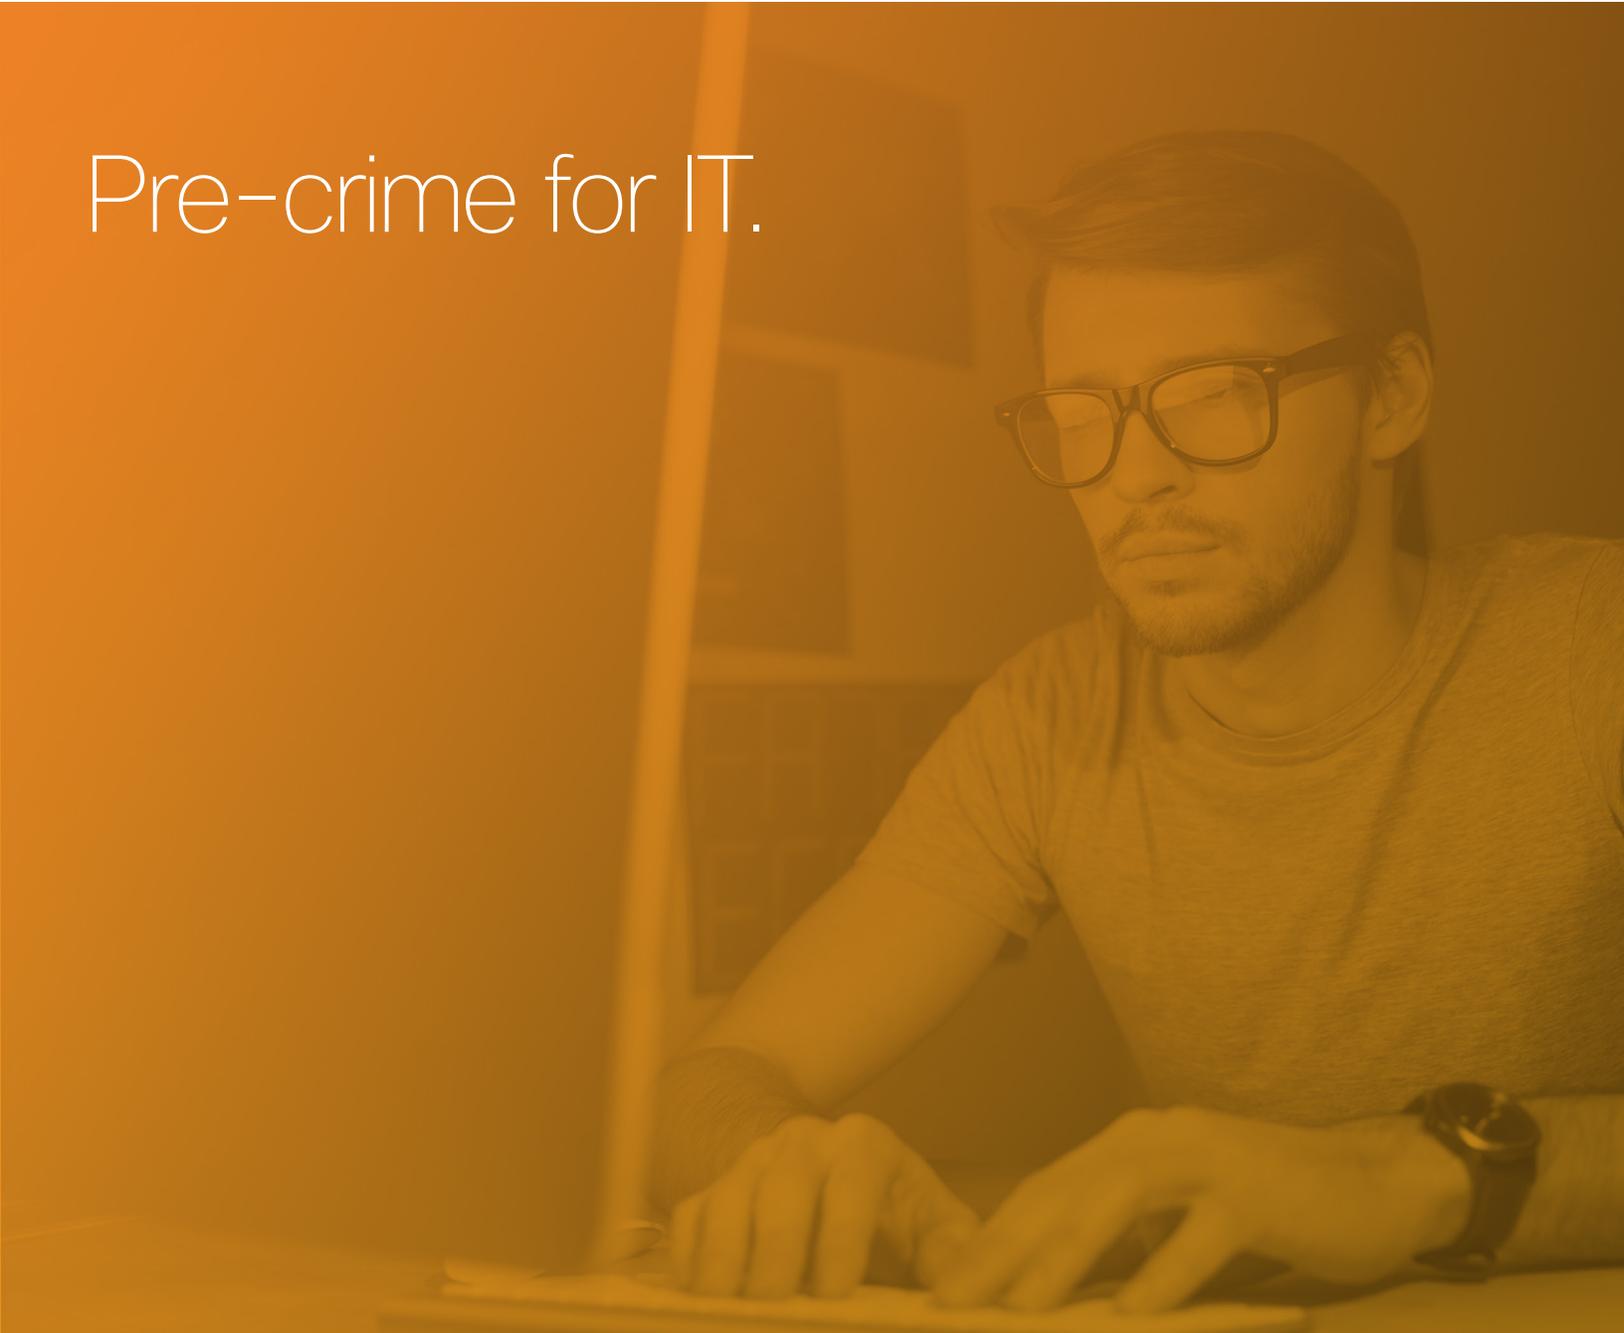
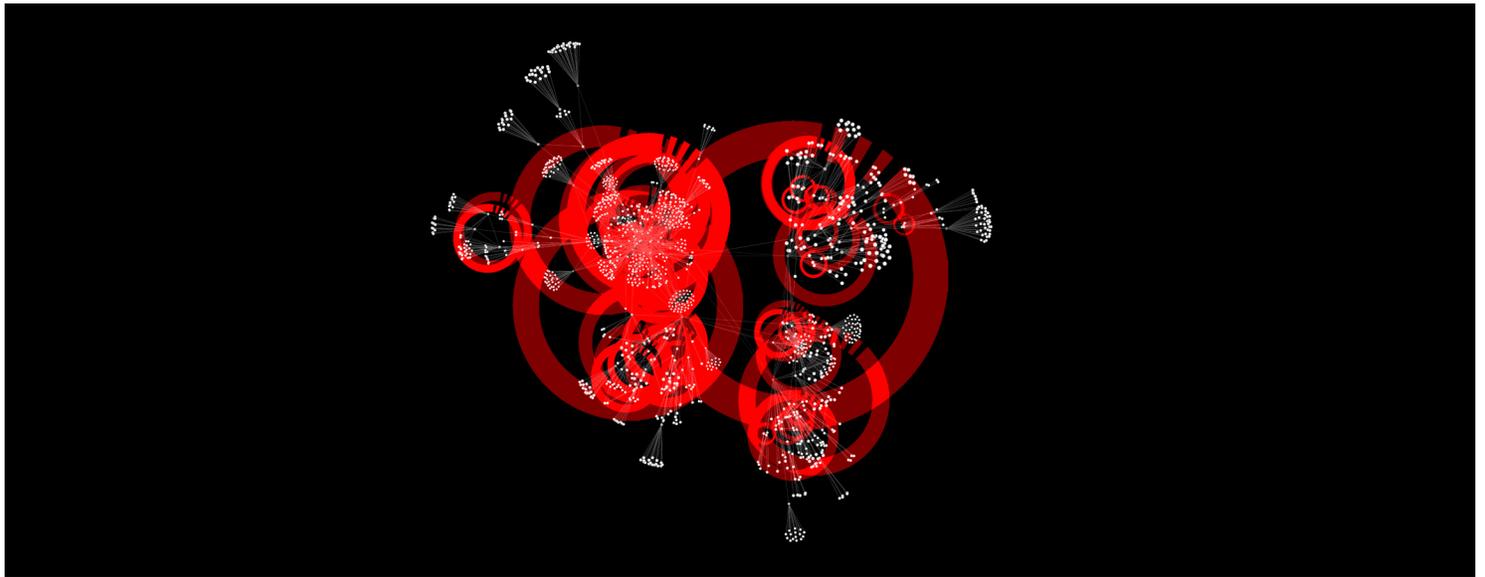


Pre-crime for IT.



Leveraging internet insight to prevent attacks — instead of just responding to them

In the movie “Minority Report,” police use predictive insight to stop crimes before they happen. Effective IT security now demands the same. Today’s sophisticated attacks routinely evade conventional after-the-fact technologies such as firewalls and signature-based malware detection. Therefore, it’s essential to adopt new measures that predictively neutralize these new threats.



Fortunately, it’s possible to predict and prevent attacks before they’re fully launched. It’s also possible to stop command-and-control exfiltrations before they do real harm. This proactive protection requires:

- A large, statistically significant volume of real-time internet infrastructure intelligence
- Statistical models that leverage this intelligence to immediately pinpoint and neutralize attack infrastructure without waiting for detection of attack artifacts

By implementing a cloud-delivered security service possessing these attributes, you can block phishing attempts, bespoke malware, and other evolving threats from the moment attackers first start spinning up their attack infrastructure. You can also quickly identify and neutralize potential destinations for exfiltration. Just as important, you can gain this improvement in protection immediately, without disrupting your existing environment.

Identifying attack infrastructure with statistical models

IT security has historically focused on identifying attack artifacts such as malicious payloads after an attack is fully launched – and then attempting to defend against those specifically identified attacks. The first step in any type of attack, however, is to create attack infrastructure from which to launch that attack.

Predictive IT security therefore focuses on identifying this infrastructure – and pre-emptively neutralizing it. This way, attacks

can be stopped even before the specific nature of the attack is fully identified.

As noted above, this pre-emptive infrastructure intelligence requires a substantial volume and diversity of internet infrastructure data. But the right statistical models are also extremely important for identifying attack infrastructure in real time as it is being created.

This automated identification is especially important given the huge volume of attack domains

constantly being generated worldwide – especially using domain generation algorithms (DGAs) – and the immediacy with which IT needs to pinpoint any and all malicious infrastructure “needles” in the real-time DNS (Domain Name System) “haystack.”

Predictive identification of attack infrastructure depends on three general types of statistical models.

The infrastructure imperative: How “attackops” enables proactive defense

Attacks don’t just suddenly happen. There is a development life cycle to creating new threats similar to the legitimate creation of new business applications. We build something, test it a few times, and then launch the new service. To launch new threats, attackers work through an AttackOps process that includes:

- Coding the malicious payload
- Staging server infrastructure on the internet
- Registering domain names
- Testing the payload on random targets

If the attack works, the attacker launches it. If not, the attacker goes back to the drawing board and tries again.

Furthermore, when their payloads and/or infrastructure are

discovered days, weeks, months, or even years later, attackers make changes to evade security technologies such as signature-based malware detection, reputation systems, and blacklists.

This AttackOps process enables a DNS-based defense to detect and neutralize attacks far earlier and more effectively than other security technologies for three main reasons:

- **Attack infrastructure precedes full-scale attack activity.**

Attackers must stage servers and register domains before sending payloads and managing botnets. By detecting that infrastructure, a DNS-based defense can protect the enterprise proactively – regardless of how stealthy an attacker’s payload or social engineering tactics may be.

- **Attackers use common methods to provision attack infrastructure.**

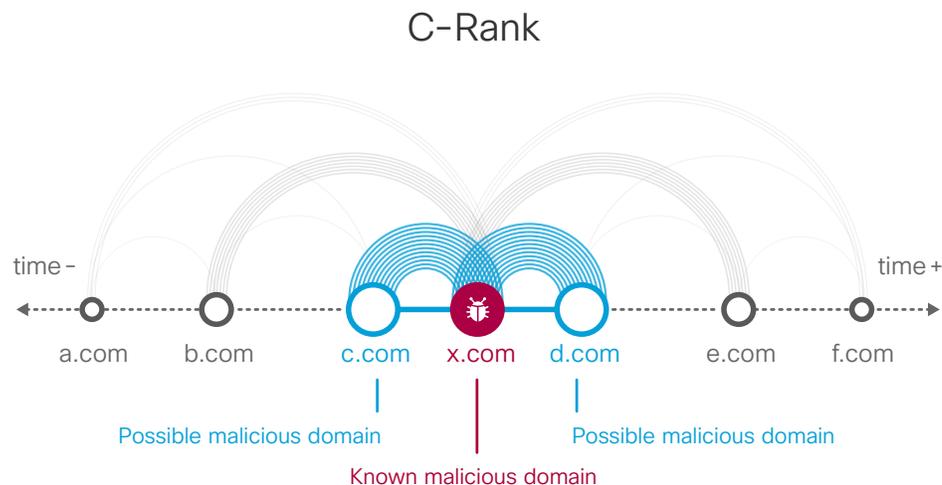
Malicious payloads and social engineering schemes vary considerably; the process of provisioning attack infrastructure does not. Therefore, with the right detection algorithms, that infrastructure can be detected quickly and reliably.

- **Attackers reuse attack infrastructure.**

Enterprise system administrators know how time-consuming and expensive it is to completely rebuild server infrastructure. It’s no different for attackers. Attacker’s often reuse much of what they’ve already built, making them highly susceptible to DNS-based detection.

Guilt by inference

One particular difficulty with present-day attacks is that there's often nothing noticeably anomalous about the actual payloads moving between the attacker and the target. So attack infrastructure can't always be quickly or accurately identified simply by tracing the origin of malicious payloads. There are, however, several ways to reliably identify attack infrastructure by inference.



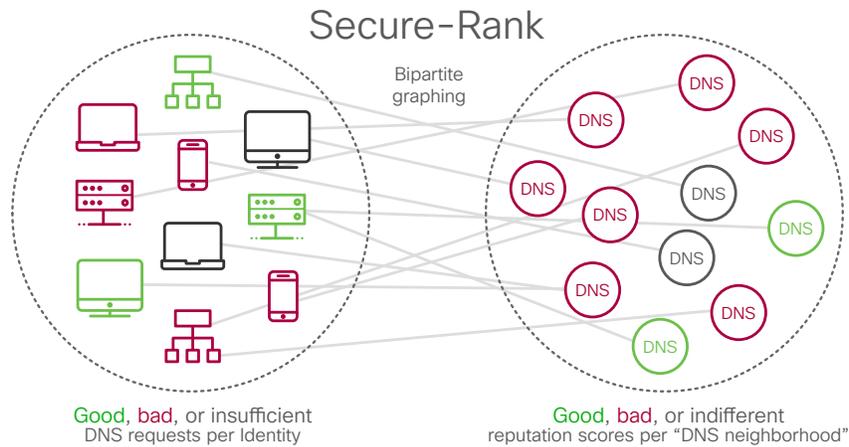
- **C-Rank.** C-Rank, as in co-occurrences rank, is one method of tracking requests that occur both immediately before and immediately after requests for domains that are known to be malicious or suspicious. Like a boxer who studies an upcoming opponent's films to look for actions that habitually precede a certain kind of punch, analysis of thousands of unique DNS requestors can reveal the requests that habitually precede or follow problematic ones. By discovering patterns in these "co-occurrences," domains that might otherwise appear innocuous can be correctly identified as malicious with statistical significance.

- **Geographic models.** Attack infrastructure can also be inferred through geographic analysis of internet activity. That's because even surreptitious command-and-control attacks can reveal themselves through the anomalous distance between the attack infrastructure and its attack targets.

One geographic model creates a behavioral baseline by capturing the normal geographic diversity of identities making DNS requests to top-level domains (.com, .edu, .ru, .cn, etc.) for any specified domain. Potentially malicious attack infrastructure can then be identified by DNS requests that diverge from this

baseline (e.g., identities located only in Germany that request domain.us are suspicious).

Another geographic model uses as its baseline the typical total distance between all IP hosts in most domains. Domains that significantly diverge from this baseline can then be identified as suspicious. This model is especially useful for discovering "fast flux" networks – where the IP addresses are rapidly swapped among a botnet of compromised proxy hosts – often associated with phishing attacks.



- **Secure-Rank.** A third method of predictively identifying attack infrastructure by inference applies a statistical model known as “bipartite graphing” to DNS requests.

Command-and-control botnets generate a huge volume of DNS requests from infected endpoints calling back to the attack infrastructure. So, over time, there are identities that look up known malicious destinations with some frequency – while other identities look up legitimate destinations and steer clear of attack infrastructure.

Bipartite graphing combines both affirmative evidence and negation to rank unknown destinations as malicious or safe. And unlike traditional reputation, the rank is always changing based on live internet activity. To use an analogy: bad people tend to hang out in bad neighborhoods, whereas good

people tend to hang out in good neighborhoods.

- **DNS neighborhood scores.** The likelihood that a domain is being used for malicious purposes can be further inferred based on its “neighborhood” – i.e., the IP addresses linked to its DNS records. Are there known malicious hosts on the same subnets as any IP addresses associated with the domain? Are there known malicious hosts on all subnets within any autonomous systems that include IP addresses associated with the domain?

These inferences alone are insufficient to identify an unknown domain mapped to a given IP as malicious, but they are very useful when combined with bipartite graphing – and they are certainly more effective than a reputation score based on individual destination IP addresses alone.

Cumulatively, these inference-based methods identify attack infrastructure much earlier and more reliably than conventional reputation scoring that checks only to see how recently a domain was registered and/or whether it has already been involved in a recognized attack.

Guilt by association

The potential “guilt” of attack infrastructure can be deduced by association as well. Two ways of doing this are:

- **Passive DNS and WHOIS**

correlation. Data from authoritative DNS and DNS registrars can be used to build rich passive DNS and WHOIS databases. Using these databases, it is possible to correlate domains based on everything from name servers to registrant email addresses. This enables rapid, early identification of domains related to other domains used by any criminal – regardless of whether that domain has yet been associated with any particular spearphishing attack or malicious payload.

- **Predictive IP space modeling.**

Starting with compromised domains identified through SP-Rank (Spike Rank) as initial clues, this model scores each step in the process that attackers engage in as they set up their attack infrastructure – from choosing a hosting provider to deploying server images – to determine which domains are likely to be the source of future malicious activity. By focusing on the unchangeable characteristics of infrastructure provisioning, this model can identify more than 300 new potentially malicious domains every hour

and can block them before they are used in an attack campaign, thereby overcoming the evasion techniques that criminals typically employ.

Similarly, Border Gateway Protocol (BGP) routing data can be used to detect shared suspicious behaviors of a domain that may be hosted on completely different autonomous systems on different dates – or to correlate all the IP blocks and domains being used as part of a given attacker’s total infrastructure portfolio.

Of course, attackers can theoretically avoid such telltale associations by, say, using a unique email address for each registration. But in the real world, no crime is committed perfectly. By using a comprehensive set of associations, an attacker’s “fingerprints” can be detected even if the attacker does only one thing wrong. This approach turns the tables on the conventional security dynamic that allows attackers to be successful even if they do many things wrong – while requiring enterprises to do everything right if they hope to remain safe.

One particular difficulty with present-day attacks is that there’s often nothing noticeably anomalous about the actual payloads moving between the attacker and the target.

Patterns of guilt

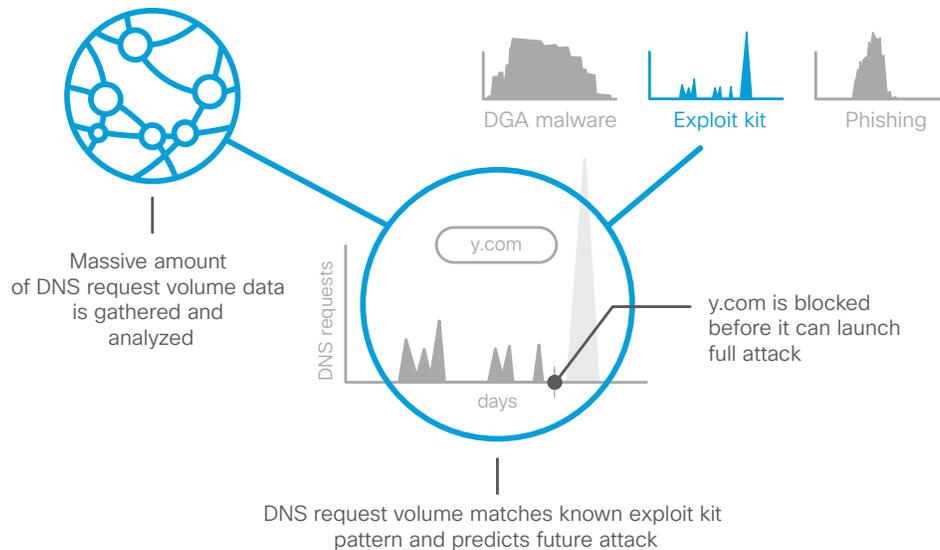
A third critical set of statistical models for detecting attack infrastructure revolves around discernible patterns of malicious behavior. This pattern-based approach echoes the self-learning techniques being used for everything from predicting customer buying preferences to anticipating power demands across the electrical grid. In this case, though, the patterns being detected are those of attackers as they create and manage their attack infrastructure.

Here are two prime examples:

- **NLP-Rank** (natural language processing-rank) was developed by Cisco in 2015 after a security researcher noticed that certain spearphishing attacks increased their odds of fooling users by employing fraudulent domain names that combined well-recognized brands (such as PayPal, Gmail, and Adobe) with particular English words (such as install, update, and download). This observation was tested – and led to the discovery that NLP heuristics could be used to further aid in the pre-emptive discovery of attack infrastructure.

Those heuristics have since been complemented with others that also “read” Autonomous System Number (ASN) mapping, WHOIS data patterns, and HTML tags to assign infrastructure with an NLP-Rank score that provides still another effective means of identifying potential attack vectors while avoiding false positives. NLP-Rank has also been enhanced with alignment techniques from computational biology to grade permutations of domain names that fit a fraudulent pattern (such as “install-ad0be” instead of “install-adobe”) to assess the likelihood they, too, will be used in a spearphishing attack.

SP-Rank



SP-Rank was developed by a Ph.D. data scientist with specific expertise in distributed sensor networks. The underlying principle is similar to that used in technologies such as Pandora's Music Genome Project. The levels of DNS activity associated with any given domain over time are akin to a sound's waveform. A domain such as Gmail or Facebook has a very consistent pattern, since its massive total volume tends to minimize transient peaks and valleys. Other domains, such as those of universities or travel sites, will peak at certain predictable seasonal intervals.

Attack infrastructure also exhibits its own distinctive wave signature. This signature is the unavoidable consequence of the techniques attackers use to spin up algorithm-generated domain infrastructure across dispersed autonomous systems – an activity that, by definition,

is unlike anything else that transpires on the internet.

The special value of this technique is that it can detect the existence of new attack infrastructure even when that infrastructure includes captive infected botnet hosts operating within fully legitimate domains, as has been the case with several especially successful and damaging exploits.

Just as important, as noted above, Predictive IP Space Modeling can immediately be applied to the clues provided by SP-Rank to yield a 340x increase in the discovery of malicious domains.

While each of these techniques is individually very useful for discovering attack infrastructure, their full value is collective. The combination of inference, association, and pattern recognition models – all fed

with high volumes of internet data and continuously refined by leading security experts – make it virtually impossible for attackers to launch any type of attack without first revealing themselves. And the moment such potentially malicious infrastructure is discovered, it can be instantly neutralized by simply enforcing it globally within minutes.



Predictive IP Space Modeling can immediately be applied to the clues provided by SP-Rank to yield a 340x increase in the discovery of malicious domains.

Why change now?

Enterprises have made significant investments in conventional security technologies.

Unfortunately, these technologies are proving to be inadequate for defense against an intensifying barrage of increasingly sophisticated threats.

By adopting the “pre-crime” statistical models described above, however, enterprise IT security leaders can dramatically improve their defenses. More specifically, they can gain three extremely valuable benefits:

- **Pre-emptive protection from even the most advanced zero-day threats.** No matter how clever or surreptitious attackers may be, they first have to spin up attack infrastructure. By identifying and neutralizing that infrastructure, a cloud service enforcing statistical model-based intelligence

provides pre-emptive protection unlike any other security technology.

- **Early interdiction of command-and-control traffic.** The worst high-profile security breaches often occur because an enterprise failed to detect exfiltration for several months. Statistical model-based intelligence nips these damaging long-term exfiltrations in the bud, because they can quickly identify and block the external destination of command-and-control traffic – long before enterprise security teams typically discover and pinpoint an internal (and potentially well-camouflaged) compromise.
- **A 50% to 90% reduction in security alert volume.** Enterprise security isn’t just about effective defense. It’s

about effective defense within finite resource constraints. By blocking attack infrastructure before most attacks can even touch the enterprise, DNS-layer security dramatically reduces security events and alerts – freeing staff resources for other tasks.

Umbrella’s statistical models have been proven in the field to be nearly 100% successful at discovering new malicious destinations before next-generation firewalls, secure web gateways, or sandboxes find them. Given this efficacy – and given the magnitude of the risk posed by next-gen attacks – these “pre-crime” techniques have become indispensable to IT security.

Feeding statistical models with statistically significant internet data

The internet is a big, busy place. Attackers rely on its scale and complexity to hide their malicious activity. That’s why, to uncover that activity in every dark corner of the internet, it’s essential to feed statistical models with large samplings of geographically diverse internet data, including DNS, WHOIS, BGP routing, IP geolocation, malware file connection, and SSL certificate information.

Cisco Umbrella processes 80 billion DNS requests daily from 65

million users across 160 countries – or about 2% of the internet’s total activity. This in-house data is already statistically significant. Umbrella further complements that data with WHOIS records and BGP routing data from 500+ peering partners, as well as malware alerts from Cisco and other global security leaders. As a result, Umbrella currently protects 7 million total malicious destinations at any give time – and identifies 60 thousand new malicious destinations daily. The existence of 7 million malicious

domains underscores just how intense attack infrastructure activity has become – especially as attackers increasingly use domain generation algorithms (DGAs) to spin up thousands of domains at a time. It also underscores how important it is to choose a security provider capable of detecting and blocking that infrastructure before an attack compromises an organization’s digital integrity.

30 minutes to a safer enterprise

It's not enough to understand how statistical models work. If you're an enterprise IT security decision maker, you also have to determine exactly how – and when – to put those models to work.

For the “how,” consider these solution attributes:

- **Data volume, variety, and quality.** Accurate analytic results are entirely contingent on statistically significant data inputs. If you don't have enough data – or if that data does not meet requirements for accuracy, diversity, immediacy, etc. – your outputs will suffer. That's why IT security leaders need to carefully evaluate the data collection resources of any solution provider on their short list. (See callout on page 9)
- **Sophistication and variety of statistical models.** As attackers become smarter about how to attack, defenders must get smarter about how to defend. That's why it's essential to bring the most advanced statistical models available to bear on the enterprise security challenge. It's also important to work with a service provider whose industry leadership and track record of model innovation

strongly suggests it will continue keeping pace with future attack innovation.

- **The ability to quickly and reliably translate predictive identification of attack infrastructure into predictive blocking of attack traffic – anywhere users/laptops go.**

There are a variety of ways to enforce DNS-layer security. Ideally, however, it should be executed in a high-performance cloud environment that can protect any device on your network without introducing latency or outage risk into legitimate unblocked traffic flows.

The “when” is even simpler to answer. Every hour that passes without the protection of predictive threat intelligence puts the business at risk, so sooner is better than later.

And with Umbrella, that “pre-crime” protection can be put in place within 30 minutes by simply pointing enterprise DNS requests to Cisco Umbrella

Sixty-five million daily-active users and 12,000 businesses are already doing so – and operating much more safely as a result.

The Cisco Umbrella advantage

#1 fastest & most reliable DNS with/65M+ daily active users

80B+ daily internet requests or connections

3M+ daily new domain names discovered

60K+ daily malicious destinations identified

7M+ total malicious destinations enforced at any given time

80M+ daily malicious requests blocked

To instantly launch a free 14-day trial of predictive DNS-layer security, simply click signup.opendns.com. There's no cost, no obligation, and no phone calls. Just immediate, substantial improvement of your organization's security – along with a noticeable and immediate reduction in your security alert volume.