PipelinePub.com
September, 2006 | Volume 3, Issue 4

Skype:
The Future of Traffic Detection and Classification

By Antonio Nucci, Narus CTO

With the increased deployment of high-speed ("broadband") Internet connectivity, a growing number of businesses and individuals are using the Internet for voice telephony. The proprietary VoIP system that is having the most dramatic impact on carriers' revenue streams and network security is Skype. It uses a unique peer-to-peer technology, making it especially challenging for carriers to identify, classify and manage associated traffic.

Skype: New Communication, New Concerns

Many enterprise IT managers fear the introduction of Skype onto their networks. Both Skype's vulnerability to hackers, either for the purposes of eavesdropping or virus launching and the amount of bandwidth it could potentially consume are issues that could have a far-reaching impact on not just the enterprise network, but the telecom carrier as well.

Security

In line with the claims of its creators, Skype appears to encrypt or otherwise scramble information that is transmitted over the Internet. Although it is generally accepted that Skype is secure against casual snooping, it is not clear how it would fare against sophisticated attackers.

The security of any data sent over an encrypted connection depends upon many factors, including the specific encryption algorithms used and how encryption keys are chosen or exchanged (known as key management). Also of critical importance is the protocol that employs the algorithms, and how well both the algorithms and protocols are implemented. An analysis of the packets sent between Skype clients indicates that a combination of protocols appears to be used for actions such registering oneself on the network, searching for other participants, or making a voice telephone call.

Skype claims that its system employs RSA's encryption for key exchange and 256-bit AES as its bulk encryption algorithm. However, Skype does not publish its key exchange algorithm or its over-the-wire protocol. Despite repeated requests, Skype refuses to explain the underlying design of its certificates, authentication system, or encryption implementation. It is therefore impossible to validate the company's claims regarding encryption. A poor implementation of the RSA algorithm could provide encryption, but no actual security.

In order to avoid detection, many peer-to-peer applications, including Skype, change the port that they use each time they start. Consequently, there is no standard "Skype port" like there is a "SIP port" or "SMTP port." In addition, Skype is particularly adept at port-hopping with the aim of traversing enterprise firewalls. Entering via UDP, TCP, or even TCP on port 80, Skype is usually very successful at passing typical firewalls. Once inside, it then intentionally connects to other Skype clients and remains connected, maintaining a "virtual circuit." If one of those clients happens to be infected, then the machines that connect to it can be infected with no protection from the firewall. Moreover, because Skype has the

ability to port-hop, it is much harder to detect anomalous behavior or configure network security devices to block the spread of the infection.

Supernodes

Like its file sharing predecessor Kazaa, Skype employs an overlay peer-to-peer network. There are two types of nodes in this overlay network, ordinary hosts and super nodes. An ordinary host is a Skype application that can be used to place voice calls, send text messages, etc. A super node is an ordinary host's endpoint on the Skype network, meaning that any ordinary host must first connect to a super node and authenticate itself with the Skype login server. Any node with a public IP address having sufficient CPU, memory, and network bandwidth is a candidate to become a super node - including machines that reside on enterprise networks. Because Skype super nodes are created dynamically, and could conceivably consume as much bandwidth as is available to them, enterprise IT managers consider these super nodes a significant risk to the health of their network.

Privacy and Authenticity

When you initiate a Skype conversation, how sure are you that you are actually reaching the user that you specified? Every Skype user has a username and a password. It appears that the network is used by Skype to perform username/password verification, but it isn't clear how this is done. For example, hosts on the Skype network could relay the encrypted username/password combination back to Skype's servers for approval. Alternatively, they could relay an unencrypted username/password combination. If the Skype network is indeed involved in the communications, several types of attacks may be possible:

- A malicious Skype client may learn the username/password combination of registered Skype users;
- If a Skype user accesses the Skype network through a malicious Internet Service Provider, the ISP may direct that user's Skype communications to the malicious Skype node. Thus, it may be possible for a malicious ISP to learn any of their user's Skype passwords;
- A malicious node may fake a valid authentication, allowing a client to log in with a particular Skype username even though the password for that username is not known.

When using Skype as a voice communications system, its users can often rely on identifying a person by the sound of their voice. This layer is absent, however, if Skype is used only for text messaging and exchanging files. These challenges are forcing carriers to look for accurate ways to detect Skype (and

other P2P protocols). In some cases the telecom Marketing departments are highly interested in what percentage of their customers are using Skype so that they can decide whether or not to launch their own commercial VoIP service. In other cases, unpredictable bandwidth consumption and security issues are concerning enterprise IT managers- the customers of the telecom carrier. Many of these enterprise IT managers are responding by requiring that the carrier actually block Skype traffic *before* it hits their private networks.

Challenge: Detection of Skype Traffic

In general, effective Internet traffic detection and classification requires three key elements:

1. Accuracy: the technique should have low false positive (identifying other protocols as targeted protocol X);
2. Scalability: the technique must be able to process large traffic volumes in the order of several hundred thousands to several million connections at a time, with good accuracy, and yet not be computationally expensive;
3. Robustness: traffic measurement in the middle of the network has to deal with the effects of asymmetric routing (two directions of a connection follow different paths), packet losses and reordering.

There are usually tradeoffs in terms of the level of accuracy, scalability and robustness that can be achieved relative to the detection of any given protocol or service.

INDUSRTY PARTNER

One current classification practice consists of TCP/UDP port number application identification using known TCP/UDP port numbers to identify traffic flows. This method is highly scalable since only the TCP/UDP port numbers must be recorded to identify a particular application. It is also highly robust since a single packet is sufficient to make a successful identification. Unfortunately port number-based identification is increasingly inaccurate primarily due to the fact that P2P networks tend to intentionally disguise their generated traffic in order to circumvent filtering firewalls (as well as legal issues associated with organizations like the Recording Industry Association of America). Most P2P networks now operate on top of custom-designed proprietary protocols and their clients can easily operate on any port number - even HTTP's port 80, making port-based detection schemes incapable of accurate and robust classification of Internet protocols.

To overcome the issues with port-based detection, a new technique has emerged based on payload-signature methods, in which packet payloads are processed for patterns or signatures that univocally identify any given protocol. One challenge facing payload-signature techniques on telecom networks is the high speed at which such pattern matching algorithms must be executed, e.g. 2.5Gbps (OC48) and above. It is therefore critical to design algorithms that can efficiently perform pattern matching while simultaneously dealing with memory and CPU limitations. Another key challenge is the lack of openly available, reliable protocol specifications. This is partially due to developmental history and partially a result of the proprietary nature of many protocols. For example, most P2P protocols are both proprietary and constantly evolving. Some of these (Gnutella for instance) provide some documentation, but it is often incomplete, or not up-to-date. To make matters worse, there are various implementations of Gnutella clients, some of which do not comply with the specifications in the documentation (raising potential interoperability issues). For application detection and classification to be accurate, it is important

to identify signatures that span all the variants (or at least the dominantly used ones). However, it is increasingly common to see new applications (such as Skype or GCN) employing 128-bit or 256-bit encryption techniques to defend the privacy of the information exchanged between their users. As a consequence, the payload-signature method fails when traffic is encrypted, because the signatures in the packet payload are scrambled by the encryption.

Skype offers a combination of challenges that make it notoriously difficult to detect with scalable, accurate algorithms:

- The Skype agent does not run on any standard source port. Skype randomly selects a source port for the agent to run on, then communicates via either TCP or UDP, or both. The choice of the protocol that Skype uses depends on whether the agent is behind a proxy/NAT or has a public IP address. The destination IP addresses are not the same every time Skype runs, and the destination port numbers are also not standard.

- All communication via Skype is encrypted. This also means that phone numbers called (SkypeOut) or other data are also encrypted. In many cases, there is no direct communication between end users in Skype. All communication passes through intermediate nodes, and these nodes may be different for every call.

- Skype is a peer-to-peer protocol, which means that the peers (IP addresses) to which a Skype agent connects are many and the network is very dynamic, so these peers (and thus their IP addresses) keep changing.

- Skype provides voice, chat, file transfer and video services. It appears that all of these services are passed together, making it difficult to separate out voice, from chat, from video, etc.

To accurately detect and classify these unfriendly applications, it is necessary to provide a systematic methodology that overcomes the lack of well-known port numbers or user payload signatures. Instead, any new methodology should analyze flow connections at the transport layer (Layer 4) to extract and profile key features from the packet streams processed. Such a method could be referred to as "classification in the dark."

Solution: Traffic Classification in the Dark

"Traffic classification in the dark" is a particularly effective protocol detection technique that involves the "pipelining" of two different detection applications: the first based on a *payload-signature model* and the second based on a *behavioral-signature model.* In this method, all TCP and UDP streams are processed first by the payload-signature application. If no match with current known signatures is found, the stream is then forwarded to the behavioral-signature application that analyzes the characteristics of the packet streams and very accurately detects even the most complex Internet Applications.

Payload-signature model: TCP and UDP streams of packets are processed first by the payload-signature application. The payload of each incoming packet is matched against a large set of constantly signatures. A match is achieved using proprietary algorithms that guarantee excellent performance at very high-speed (up to OC48). The majority of standard protocols (and their associated applications) are promptly classified by this application.

Behavioral-signature model: Any TCP and UDP streams not classified by the Payload-signature application are forwarded to the Behavioral-signature application. Streams of packets with encrypted payloads,

emerging P2P protocols for which a signature is not available, or multimedia applications using proprietary technologies (such as VoIP, Video, Gaming, File Transfer, Chat, etc) fall into this family.

The Behavioral-signature application profiles the behavior of hosts at different levels by exploring its social level (hosts that it communicates with), its functional level (servers vs. clients vs. peer-nodes), its application level (transport layer interactions between particular hosts on specific ports) and specific dynamics, with the intent to identify the application of origin.

As P2P services such as Skype continue to evolve, and find new ways to avoid detection, new entropy-based classification methods such as “traffic classification in the dark” offer great promise to network managers who wish to manage these services to ensure the health and profitability of their networks.