Technologies and Foundations for Robust and Secure Networked Systems

WInter 2012 Newsletter

Center for Networked Systems

Cloak and Dagger: Striking at the Heart of Scammer Web Sites

Have you ever entered a term into a search engine, clicked on one of the resulting links, and ended up on a page that was utterly unlike the one represented to you in your search list? Whether you knew it or not, you most likely discovered a 'cloaked' Web site, one of many strategies employed by Internet scammers to drive traffic to their grey- and black-market services. Cloaking is a bait-and-switch technique that hides the true nature of a site by delivering significantly different content to different user segments.

David Wang *(pictured above)* is a Computer Science and Engineering Ph.D. student at UC San Diego who has been studying this cat-and-mouse game played by Web scammers with search engine operators and Web users. The ploy stems from our reliance on the use of search engines. Anyone can post anything on the Internet, but without search engines, the chances that those postings could be found are minute. Search engines are the primary method for connecting information seekers to content providers, and customers to vendors. Moreover, search engines assign rankings based on a number of factors (e.g., popularity, relevancy, etc.) that give the site better placement in the list of search results.

This dependency on the use of search engines means that a Web site's search-engine ranking has a meaningful impact on the level of traffic directed to the site. Because businesses recognize their interest in maintaining the health of their ranking, search engine optimization has become an important strategy in Web site design and maintenance. While many of these methods are considered legitimate, many others – such as blog spamming or using link farms – are not. The effectiveness and broad use of illegitimate methods by those who run scam Web sites have made it a priority for those who operate search engines to discover and block such troublesome sites. In response, web scammers have employed a host of strategies to evade detection, and so-called 'cloaking' is one of the most potent weapons in their arsenal.

Continues on page 3

Protecting Anonymity for Users of Location-Based Services

While much private information has been publicly available for decades, the cost of physically accessing, recording and tabulating the data has been a de facto barrier to its use and dissemination. But with the rise of the Internet and increasingly comprehensive database technologies, untold volumes of information about individuals have been collected and made accessible either publicly or commercially. We now live in a world where public records can be searched and downloaded with the touch of a keystroke, and data is collected about everything from grocery-buying habits to patterns of political donations. This easy accessibility and the development of powerful tools for the management and analysis of this data has been a boon for a number of entities, from advertisers to government institutions – and CNS faculty member Alin Deutsch *(pictured at left)* is exploring better ways to permit use of such data while maintaining the anonymity and privacy of users.

Continues on page 7

CNS Members

cisco ERICSSON Google ORACLE

In This Issue

New CNS Sponsored Projects Announced

Ericsson inaugurated its second year of CNS membership by providing support for two new CNS research projects to study switches in data center architecture. The first, "Optical Switches in the Datacenter", will be conducted by Tajana Rosing. The second study will be led by Alex Snoeren and CNS Research Scientist Kenneth Yocum, and they will focus on "Topology Switching for Data Center Networks".

Oracle (following the merger of Sun Microsystems into Oracle), has been a member of CNS since its inception in 2004. Oracle has chosen to fund Tajana Rosing on her project "Novel Methods for Maximizing Performance as a Function of Power."

Cisco Systems has also been a member and supporter of CNS research since 2004. This year, Cisco is supporting George Porter in his work on "Incorporating Networked Storage with Highly Efficient Data-Intensive Computing".

George Porter leads the new CNS project funded by Cisco Systems.

ACM Recognizes CNS Faculty Members

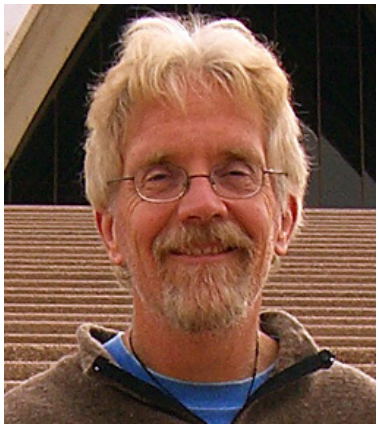

On December 8, 2011, ACM named 46 new inductees as ACM Fellows "for their achievements in computer science and information technology and for their significant contributions to the ACM." Being named an ACM Fellow is one of the highest forms of recognition bestowed upon a computer scientist. This year, former CNS Director Amin Vahdat and CNS faculty member Keith Marzullo *(at right)* were both so recognized. Vahdat was cited for his "contributions to data center scalability and management," while Marzullo was honored for his "contributions to distributed systems and service to the computing community." The new inductees join the ranks of other CNS faculty who have been recognized by ACM as Fellows: Stefan Savage and George Varghese.

Separately, CNS faculty member Yuanyuan (YY) Zhou is one of only 49 members worldwide to be named an ACM Distinguished Scientist in 2011. The Distinguished Member Grade recognizes ACM members with at least 15 years of professional experience and five years of continuous professional membership who have achieved notable accomplishments or have made a significant impact on the computing field.

CNS's Stefan Savage Delivers NSF WATCH Talk; George Porter at University of Washington

CNS Director Stefan Savage addressed NSF's "Washington Area Trustworthy Computing Hour" in November 2011. The WATCH Series aims to provide thought-provoking talks by innovative thinkers with ideas that illuminate challenges and provide signposts toward solutions. The talk focused on "Why the Hard Problem of Computer Security Needs the Soft Sciences." *http://www.nsf.gov/events/event_summ.jsp?cntn_id=122083&WT.mc_id=USNSF_13*

CNS Assistant Research Scientist George Porter spoke Nov. 3, 2011, as part of the University of Washington's Computer Science and Engineering Department Colloquia series. A podcast of Porter's talk, "Towards Balanced, Data-Intensive Scalable Computing," is available at: *http://norfolk.cs.washington.edu/htbin-post/unrestricted/colloq/details.cgi?id=1096*

Graduating Students

In February 2012, **Mohammad Al-Fares**, a CSE Ph.D. student, presented his dissertation, "A Scalable, Adaptive, and Extensible Data Center Network Architecture." His advisors were Amin Vahdat and George Varghese. Dr. Al-fares has accepted a position at Google as a Software Engineer.

In September 2011, CSE Ph.D. student **Avinash Vyas** defended his dissertation, "Policy-aware Sender Anonymity in Location-based Services". His advisor was Alin Deutsch. Dr. Vyas has become a Technical Staff member in the Computing and Software Principles Research Department at Alcatel-Lucent's Bell Labs.

Record Attendance at February 2012 CNS Research Review

On February 8-9, CNS welcomed a record crowd of 138 industry representatives, faculty, researchers and graduate students to its 16th half-yearly Research Review held at the Price Center on the UC San Diego campus. Following the new center funding model (see below), this was the first winter Research Review to feature new project proposals. Topics ranged from: managing energy in distributed, battery-operated systems; new methods for debugging noSQL analytics; a study of Internet abuse-driven affiliate programs; ways to improve energy efficiency in wireless local-area networks; and how to protect sender anonymity in location-based services.

Winter Research Review attracted 138 attendees

New Membership Structure Announced

In a response to the changing needs of our industry members and affiliates, CNS announced last August a revision of its membership structure and grant funding model. Instead of the old model of collecting membership fees to fund a number of proposals that members chose jointly, members have been invited to participate either as affiliates or as direct sponsors of research. Affiliate members join by paying annual dues that give them access to CNS events, brainstorming meetings with faculty and research scientists, and recruitment opportunities with our graduate students. Companies that wish to support specific research projects will have a seat on the CNS Advisory Board, will receive an annual visit to the company from our Director, and gain further access to networking opportunities with our faculty and graduate students.

At the Winter Research Review, CNS Director Stefan Savage announced that the new structure has so far met with approval from participating companies. He also explained that future CNS Research Reviews will be closed to all but CNS sponsors or affiliates. Previously unengaged companies who wish to explore the possibilities of membership will be given a grace period when first invited to attend CNS events. For more information about the center's new membership structure and benefits, see: *http://cns.ucsd.edu/memberbenefits.shtml*. If you have any questions, please email Director Stefan Savage at *cns@ucsd.edu*.

Cloak and Dagger

Continues from page 1...

When a search engine queries the content of a cloaked page through a crawler, it will only detect benign content, while normal visitors referred to the page through a specific search request will view the true, scam-related content. For example, the Web user may query the name of a popular singer and a list of search engine results will come up. The overview of one page looks to be of interest and the Web user clicks on the link. However, once the user has reached the page, it turns out that it has nothing to do with the popular singer and everything to do with spurious pharmaceutical sales. This is achieved when web scammers structure the benign version of their page by employing a number of strategies (e.g., by lacing their cloaked pages with popular search terms, a practice known as keyword stuffing) so that they can drive a maximum number of unwitting visitors to their sites.

Wang, working with CSE professors Stefan Savage and Geoffrey M. Voelker, analyzed the dynamics of the cloaking phenomenon. They were specifically interested in trying to measure how prevalent cloaking behavior was on different popular search engines, such as Yahoo! and Google, how cloaking behavior changes for targeted versus untargeted advertising, and how search engine providers respond to the practice. Over a five-month study, they employed a custom Web crawler that they named Dagger to track popular search terms and targeted keywords. In near-real time, Dagger was able to tell the researchers when it received distinct results for crawlers and browsers, and with that information, it could measure how many cloaked sites existed associated with the studied search terms and the average lifetime of cloaked sites and the sites they mask.

What the researchers found was that the majority of cloaked search results remain high in rankings for 12 hours. While this may not sound like a big window of economic opportunity, so long as the overhead for site placement is less than the revenue obtained from the resulting 12 hours of traffic, it remains a feasible strategy. Additionally, the economic utility of the pages is extended because they often persist for much longer before the search engine service provider discovers and blocks them. According to Ph.D. student Wang, this information could be used by search engine operators to help diminish the process. If search engine providers can "further reduce the lifetime of cloaked results they will effectively demonetize the underlying scam activity," said Wang, and they could thereby remove the incentive for cloaking altogether.

Making the Cloud More Transparent

More and more consumers are catching on to the advantages of storing their data in the cloud. Cloud storage providers such as Amazon, Microsoft, Google, and HP are able to pool their resources of hardware and technical expertise to deliver on-demand and easy-to-retrieve access to data resources for less money and perhaps with greater security and reliability than could be created by the consumer. Clients contract with providers to store their data securely, meaning that the data is replicated in multiple data centers spread across several geographic regions so that they are protected from local outages or catastrophic events like natural disasters. The result is that the owner of the data pays a nominal fee to ensure that his or her data is easily accessible, secure, and continuously available.

But how do consumers know that they are actually receiving the services that they pay for? After all, outsourcing data stewardship only makes sense if the outside provider can deliver superior service. According to CNS faculty member Hovav Shacham, consumers have three concerns in this area:

- Is the data actually being stored as claimed?
- Is the data backed up through replication? And,
- Is the replicated data stored in a variety of geographic locations to provide maximum assurance of its replicability?

These concerns were once merely theoretical, but recent, well-publicized cloud service failures have resulted in the permanent loss of data or email accounts that has injected them with a new relevancy. These incidents highlighted that the relationship between the client and the provider is one of trust and dependency, but it lacks transparency. This opacity creates discomfort in the client and opens up the possibility for careless or even dishonest business practices on the part of the provider. "It is conceivable that to cut costs, a cloud provider may claim one level of service but actually provide an inferior level," warns Shacham, who co-authored a paper that addressed this issue, "Do You Know Where Your Cloud Files Are?"

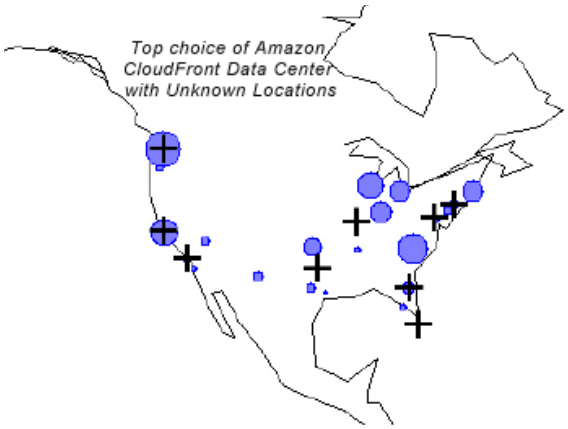

While it is true that customers sign a service agreement with storage providers, the contract is only useful as a method of enforcement or redress if the client has a way to prove that the stipulated level of service has not been rendered. To tender consumers with proof that they are receiving the services for which they pay, companies would need to devise some other method. Previously, some work has been done in the academic community to develop ways to verify that a client's data is actually being stored as claimed by the cloud storage service provider. Some CNS researchers have even proven a method by which they can verify whether a client's virtual machines are co-hosted on a physical machine with another user's virtual machine. But proving that your files are, in fact, being stored as promised is only the beginning of service verification. "Knowing that one's files are being stored is not always enough," says CNS Ph.D. student Karyn Benson, a co-author on the paper. "What if they are stored on a single hard disk and that hard disk fails? Or in a single data center that later experiences outage?" Clearly, clients cannot have peace of mind without the assurance that their data is being replicated across disks and across geolocations.

In their paper, the UC San Diego researchers outline their methods for detecting where data was being stored in Amazon's many datacenters spread throughout the United States. "We made an account on Amazon's CloudFront and uploaded a file," explains Shacham. "We then made an HTTP request for our file and the image was fetched three times from each node, and almost every time the image had a different IP address." Although the researchers' method cannot pinpoint the exact location for client data, Ph.D. student Rafael Dowsley says precise locations are not so important. "When storing data in the cloud," says the study's co-author, "customers are more concerned about the diversity of the locations than the exact locations."

The researchers caution that their work is a theoretical framework and not a full-fledged system that could be used to verify the location of a client's data within the cloud. Future work could usefully be extended to test their method against other cloud providers than Amazon, and the creation of a complete system that would account for several modifications that were not addressed in the preliminary study. Regardless, the CNS team believes that their research offers a potential solution to a pressing problem facing users and providers in the economic development of cloud services.

To read the full paper, visit *http://cseweb.ucsd.edu/~hovav/dist/cloudloc.pdf*.

Greening the Cloud

New methodology for scheduling green energy in data centers offers cloud-service providers three times better energy usage and almost eight times fewer terminated tasks

Online file sharing and storage, outsourced IT security, Internet searches and social networking Web sites are all cloud computing services that have become popular and economic choices for both individuals and enterprise-level consumers. Cloud computing is pragmatic because it leverages economies of scale so that the expensive computing hardware and technical expertise required to perform these services is moved out of the hands of the consumer and into the data centers of service providers.

Because economics is one of the main attractions for using cloud-based services, service providers face constant pressure to minimize their costs. And since one of the greatest expenses incurred by operators of data centers is energy, a cloud computing provider would ideally use the cheapest energy source possible. Yet the popularity and growth of cloud-based services have pushed up the carbon footprint of data centers in the past decade. Data centers have become so ubiquitous and are such huge consumers of energy that it has been estimated they now are responsible for 2% of global greenhouse-gas emissions. For businesses that are doubly concerned about the impact of the rising cost of so-called 'brown' – non-renewable – energy sources on their bottom line and on the impact of their operations on the environment, turning to green energy sources such as solar or wind might seem to be a viable alternative for cloud services that seek to reduce their reliance on non-renewable energy sources.

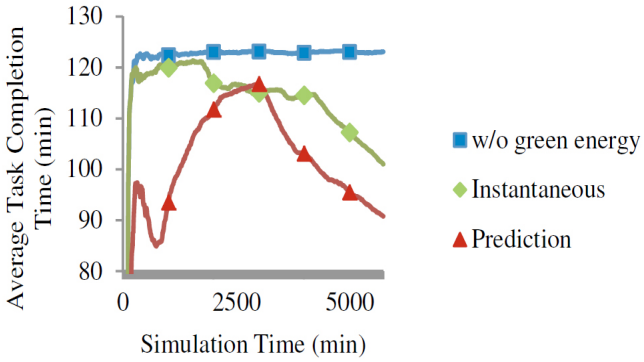

Average completion time of MapReduce tasks

How average completion time of MapReduce jobs changes over time as a function of how green energy is used

Unfortunately, green forms of energy generation come with a notable flaw: their output is unpredictable. Changing weather patterns can becalm wind turbines and cloudy skies can slow solar power production to a trickle. To complicate matters further, cloud services are susceptible to wild variations in the number of jobs being processed, i.e., on demand. Moreover, the time-sensitive nature of many services hosted by cloud-computing providers makes it difficult for them to tolerate unpredictability in power sources. This makes for a potentially disastrous scenario, where high demand occurs during a period of low power production, resulting in slower responses and cancelled tasks.

CNS graduate students Baris Aksanli, Jaggannathan Venkatesh, and Liuyi Zhang, working with CNS faculty member Tajana Rosing, decided to address this issue in a recent project funded by sponsors including CNS, National Science Foundation, and CNS member companies Google and Oracle. The team is focusing on the use of an adaptive job-scheduling methodology for data centers that more efficiently leverages green energy.

The team applied time-series prediction algorithms to estimate near-future availability of energy and then scaled the data center workload so that it does not exceed the expected level of supply. Even if the levels of green energy drop below the necessary level to support the execution of the tasks or jobs currently demanded of the system, explains Ph.D. student Aksanli, "the system offsets the remainder of the immediate need with brown energy with the assurance that over the prediction interval the average green energy will ultimately be available."

When compared to the instantaneous use of green energy, the group's tests of the scheduler on a simulation platform showed that its predictive policy leads to three times better energy usage. Moreover, the policy makes data center operation far more efficient, with almost eight times fewer terminated tasks. The scheduling methodology, says CSE professor Rosing, "allows a more efficient use of the available energy [and reduces] the amount of wasted green energy and the number of tasks/jobs that must be re-executed." It also works to increase the total throughput and productivity of the data center. The methodology results in an overall win for both the environment and for the bottom line of cloud service providers and users.

To read the paper, "Utilizing Green Energy Prediction to Schedule Mixed Batch and Service Jobs in Data Centers", go to: *http://sigops.org/sosp/sosp11/workshops/hotpower/05-aksanli.pdf*

CNS faculty member Tajana Rosing

New Approach to Log Analytics

Modern data centers not only must store and process primary data and applications but also gather and analyze information about the data they process. Evaluating this collected metadata is called 'log analytics', which has become a critical component for managing the operation of Web sites and Web-based applications. Example: application and click logs built into commercial Web sites can now capture an unprecedented wealth of information about customer behaviors and preferences, and these kinds of logs are critical to debugging, optimizing and monitoring the security and service quality of the infrastructure that supports cloud-based services and applications.

The current method for data analysis is for businesses to pull their data from tens of thousands of machines and multiple data centers into a centralized location for analysis. Most sites utilize frameworks such as MapReduce (MR) run on a Hadoop cluster to process the information deluge. But there are drawbacks. "First, it fundamentally limits its scale and timeliness," says Kenneth Yocum, a CNS Research Scientist who led a team of graduate students Dionysius Logothetis and Kevin Webb, and Salesforce.com software engineer Chris Trezzo, studying the issue. "Second, the approach must sacrifice availability or blindly return incomplete results in the presence of heavy server loads or failures."

So why delay analysis until all logs are completely delivered? "It is often possible to accurately summarize or extract useful information from a subset of log data," notes grad student Logothetis, so all that is needed is a "systematic method for characterizing data fidelity." The team's solution was to design an in situ MapReduce (iMR) architecture that moved analytics onto the log servers themselves. This reduced both the volume of data crossing the network and the time required to transform and load the data into stable distributed storage. The iMR approach enhances log analytics by allowing continuous MapReduce jobs – permitting incremental updates – and making results available when sourcing logs from thousands of servers by supporting lossy MapReduce processing. The team also developed efficient strategies for internally grouping key-value in the network and explored the impact of failures on result fidelity and latency.

In situ MapReduce is designed to ensure that log processing is scalable, responsive, available, efficient and compatible. By moving initial log analysis onto the data sources themselves and not into dedicated clusters, iMR efficiently extracts and transforms data and thereby reduces analysis time while improving system scalability. See: *http://cseweb.ucsd.edu/~kyocum/pubs/USENIX-2011-CR.pdf*

CNS In the News

Planning Cybersecurity to Defeat the Threats of the Future

In the December 5, 2011, issue of the ***New York Times***, CNS Director Stefan Savage published an essay, "In Planning Digital Defenses, the Biggest Obstacle Is Human Ingenuity". In it, he reframed the question of how security researchers must think to anticipate threats to online security that are likely to be faced in the next decade. Some of the more interesting and thought-provoking threats that he suggests we may face in the near future: the use of 'social bots' that use social networking to manipulate human behavior; the increasing use of cyber warfare; and the vulnerability of the networked devices that have become commonplace to our lives. To read the article: *http://www.nytimes.com/2011/12/06/science/stefan-savage-girding-for-digital-threats-we-havent-imagined-yet.html?_r=2&ref=science*

Interest in the Spam Economy

On October 15, 2011, ***The Economist*** published an article titled "Measuring the Black Web". It questioned whether cybercrime is as serious a threat as many claim it to be. The article featured work led by CNS Ph.D. student Chris Kanich. The referenced project was a four-year exercise during which CNS researchers tracked 20 shadow businesses that use spam to advertise the goods and services of illegal online pharmacies. First, the researchers secretly monitored the spammers' payment systems, and they later obtained logs from one of the servers that power the illegal pharmaceutical sites. This provided them with detailed information about the probable number of actors fronting the many Web sites, their sales volume, and the profitability of the spam campaigns. To complete their map of the spam economy chain from beginning to end, the researchers even ordered (and, perhaps surprisingly, received) some of the non-prescription drugs on sale. The study findings suggest that only two of the 20-plus operators net a gross income in excess of $1 million per month. To read this article: *http://www.economist.com/node/21532263*

A related groundbreaking study by the same CNS researchers on the underground spam economy was profiled in the December 12, 2011 issue of ***Bloomberg Businessweek***. The article provides a high-level overview of spammers' business models along with a number of interesting statistics about the spam economy, such as the percentage of Americans who have purchased goods and services through spam advertisements (12%) and the probable number of businesses (45) behind 69,002 Web sites. To read this article: *http://www.businessweek.com/magazine/spam-works-12082011-gfx.html*

CNS Car Security Work Grabs Attention

On October 13, 2011, ***Bloomberg Businessweek*** published an article on "Making Cars More Hacker-proof", which featured work by CNS researchers who showed how to infiltrate automobiles that are equipped with networked devices. The project proved that hackers could remotely access and control such key vehicle operational systems as brakes and steering. Though there are no known incidents of this actually occurring, the findings of this research put the automobile industry on notice about potential threats to the security of their products. Industry members have subsequently collaborated with CNS researchers to address these issues in the next-generation versions of their products. The car security project was a joint venture between UC San Diego and University of Washington led by CNS Director Stefan Savage and University of Washington Associate Professor (and UCSD CSE alumnus) Tadayoshi Kohno. The CNS team was comprised of Professor Hovav Shacham, graduate students Stephen Checkoway, Damon McCoy, and Daniel Anderson, and staff programmer Brian Kantor. To read this article: *http://www.businessweek.com/magazine/making-cars-more-hackerproof-10132011.html*

Protecting Anonymity for Users of Location-Based Services

Continues from bottom page 1...

However, the many kinds of value that can be extracted from the analysis of this new wealth of data must be balanced against the demands of individuals that their information be stewarded in such a way that their privacy is protected. To achieve this, data holders typically strip their released data of 'identifying information', such as a person's name, street address, or social security number. Yet CNS's Deutsch and other researchers in the field believe that it is not unduly difficult to cross-reference various pieces of unscrubbed information with publicly available records to successfully reduce the uncertainty about the identity of the individuals described. When this was demonstrated, it was suggested that one way to prevent this from happening would be either (a) to suppress more pieces of information, or (b) to swap it between individuals within the data set. However, these methods diluted the accuracy and therefore the usefulness of the released data. In order to protect privacy while also protecting the truthfulness of the data, software researchers developed the concept of 'k-anonymity'.

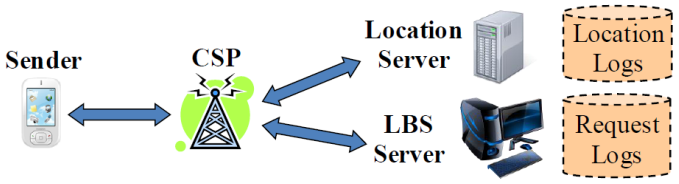

Each logged request is associated with the identifier of the device that sent the request (e.g. IP-address or MAC-address). This allows the LBS provider to identify the requests sent by the same user in the LBS request log and assemble a history of LBS requests for each user.

K-anonymity dictates that released sets of data be such that any combination of values of information that could be used to identify people are indistinctly matched to at least 'k' respondents. This is accomplished through the generalization and suppression of certain pieces of data (e.g., a ZIP code can be generalized by eliminating the least relevant digit).

Location-Based Services: New Threat to Privacy

One of the most exciting developments in mobile computing is also proving to be a new threat to the privacy of individual data. This development: the proliferation of so-called 'location-based services' (LBS). These services pinpoint an individual mobile device within the mobility network so that users can access useful information relevant to their current location, such as finding highly-rated restaurants nearby, or showtimes at the neighborhood cinema. The privacy concern arises because LBS providers keep logs of these queries that record a treasure trove of information not only about people's interests and behaviors, but also about their patterns of interests and behaviors. This information can, in turn, be analyzed and used in a variety of ways, such as in marketing products or services to that person or to other people who fit a similar behavioral profile.

Sender k-anonymity was developed to hide the identity of mobile device users, even from malicious attackers, by employing an algorithm that cloaks the specificity of information from the request log and precise location from which the request was generated. The result: someone viewing the information could not distinguish the individual from k-1 other possible requesters.

Policy-Aware Sender Anonymity

However, the UCSD research team headed by professor Alin Deutsch has shown that current methods for providing anonymity to LBS users are insufficient. In a recently produced paper on the subject, Deutsch – who has openly his concerns over the privacy of personal information stored in databases – states that sender k-anonymity defends "only against naïve attackers who have no knowledge of the anonymization policy that is in use." That is, when attackers gain access to a few key pieces of information (these being the request from the LBS log, the location of the mobile device at the time of the request, and the design of the system used to provide protection), the identity of the individual can be deduced.

While it might seem unlikely that the anonymization policy could be known, a Ph.D student who worked on the project, Kevin Keliang Zhao, explains that this is "a realistic threat since an attacker with subpoena powers (e.g., a federal agency) or a disgruntled ex-employee can obtain the 'design' of the system" and also because these system designs are based on well-accepted principles that are "not secret."

According to Deutsch, the group's work is "to keep the requestor's interests private even from attackers who (via hacking or subpoenas) gain access to the request and to the locations of the mobile user and other nearby users at the time of the request."

One of the high-performance priorities for LBS is speed, so any noticeable effect on responsiveness would meet with disapproval by the vast majority of users. Because of the seriousness of the concern, Deutsch's team studied performance effects, and also concluded that the proposed algorithm results in only a minimum reduction in utility to users. The CNS researchers show that only 16 servers are sufficient to "provide anonymization time of about half a second for one million users" in the San Francisco Bay area – "only a 1% divergence of the cost from optimum." Deutsch concludes that their work "strikes a paradigmatic balance in the trade-off between strength of the privacy guarantee, utility, and running time for enforcement."

CNS Well Represented at 2011 USENIX Security Conference

The USENIX Security Symposium is one of the premier systems and networking academic conferences, where the best researchers in the field present their latest findings. CNS was privileged to be well-represented at the latest USENIX Security conference, held last August 8-12 in San Francisco. Two of the 14 papers presented at the USENIX Workshop on Offensive Technologies (WOOT) were authored by CNS research teams, as were two of the 12 papers presented at the Workshop on Cyber Security Experimentation and Test (CSET). Additionally, CNS faculty member Hovav Shacham served as a member of the USENIX Conference Program Committee, while CSE Ph.D. student Stephen Checkoway sat on the WOOT '11 Program Committee. (Paper co-authors below are from the University of California, San Diego, except where noted otherwise.)

WOOT '11 Papers:

Heat of the Moment: Characterizing the Efficacy of Thermal Camera-Based Attacks
Keaton Mowery, Sarah Meiklejohn, and Stefan Savage

Putting Out a HIT: Crowdsourcing Malware Installs
Chris Kanich, Stephen Checkoway, and Keaton Mowery

Abstracts, .pdf copies of the WOOT '11 papers, and videos of the presentations can all be found on this page: *http://www.usenix.org/events/woot11/tech/*.

CSET '11 Papers:

No Plan Survives Contact: Experience with Cybercrime Measurement
Chris Kanich, Neha Chachra, Damon McCoy, David Wang, Marti Motoyama, Kirill Levchenko, Stefan Savage, Geoffrey M. Voelker (UCSD), and Chris Grier (UC Berkeley)

ExperimenTor: A Testbed for Safe and Realistic Tor Experimentation
Kevin Bauer (University of Waterloo), Micah Sherr (Georgetown), Damon McCoy (UCSD); and Dirk Grunwald (University of Colorado)

Abstracts, .pdf copies of the CSET '11 papers, and slides from the presentations can all be found here: *http://www.usenix.org/events/cset11/tech/*

Get Connected

Stay up-to-date about upcoming CNS events, including lectures and Research Reviews, by signing up for the CNS Events RSS feed. To do so, visit CNS online at:

http://cns.ucsd.edu

and click on the link "CNS Events RSS Feed."

Upcoming Events

May 4, 2012
CNS Security Day
Price Center Forum, UCSD, La Jolla, Calif.

CNS research is structured around the three major research themes: data centers and cloud computing; access and mobility; and security. In order to give our member companies with special interests in these areas more in-depth exposure to the work being conducted by CNS researchers and their colleagues, the center will begin to host day-long symposia focused around these themes. The first event, *CNS Security Day*, will feature speakers from CNS and beyond. They will present their newest and most innovative approaches to solving some of the most pressing problems facing security researchers. There will also be numerous opportunities for informal interactions and follow-up meetings between our industry partners, researchers and graduate students.

For information about the event: http://cns.ucsd.edu/upcoming.shtml Email Kathryn Krane at kkrane@ucsd.edu re: event registration and attendance information.

Mission and Objectives of CNS

The mission of CNS is to develop key technologies and frameworks for networked systems. By combining our research talents and strengths in partnership with industrial leaders, CNS achieves critical mass and relevant focus, accelerating research progress and creating key technologies, frameworks and systems understanding for robust, secure networked systems and innovative new applications. CNS also works to educate the next generation of top students with a perspective on industry-relevant research and to train students on how to continue their leadership throughout their careers. This is accomplished by bringing together leading faculty, students, and companies to investigate the most challenging, interesting and important problems in computer networks.

If you are interested in joining the Center, please contact Director Amin Vahdat at vahdat@cs.ucsd.edu.