Towards Trusted Adaptation Dynamics in Computing Systems and Networks

Azer Bestavros

http://www.cs.bu.edu/~best

Computer Science Department
Boston University

The Devil You Know

Largely, exploits of computing systems and networks – and consequently research in system and network security for countermeasures to such exploits – have targeted *static* properties – namely characteristics or features of a system that are fairly independent of the system's workload. Examples of widely-known (and highly-researched) exploits of such static properties include "Denial of Service" (DoS) attacks and virus/worm attacks.

Security breaches due to computer worms and viruses exploit known bugs or bad software practices in the implementation of services and protocols. Examples include buffer overflows and related code-injection attacks, which were estimated to be the culprit for roughly 50% of the major security flaws over the last two decades. Clearly, the characteristic property of the underlying system that permits such an exploit is *"static'"* in the sense that memory safety (or lack thereof) is an inherent feature of the programming language used to implement such systems. Indeed, systems developed using memory safe languages such as CCured and Cyclone, or those using other defensive technologies such as StackGuard and ProPolice are not prone to such exploits.

Another example of exploits that target static properties of a computing system or network are DoS attacks. DoS attacks target the *fixed* capacity of a system component. An adversary bent on limiting access to a network resource could simply marshal enough client machines to bring down an Internet service by subjecting it to sustained levels of demand that far exceed its capacity, making that service incapable of adequately responding to legitimate requests.

While such exploits may be viewed as mere nuisances – mounted for curiosity's sake, benign acts of free speech, or even commercial advantage – their impact on critical resources and services may cripple our increasingly cyber-dependent economy. Luckily, exploits of system static properties are not easy to mount because they do require control of a fairly large base, *e.g.*, 100K-200K zombie clients in the case of MyDoom. More importantly, by their very nature, such attacks are easily anticipated, allowing countermeasures to be taken, including the collection of information that could be used to prosecute the attack perpetrators. Indeed, the ability to anticipate a DoS attack and/or to trace-back perpetrators thereof act as powerful deterrents.

The Devil You Don't Know

But, what if victims of an attack cannot anticipate or even detect that they are under an attack? What if the attack's purpose is not to necessarily cripple a service, but rather to inflict significant degradation in some aspect of the service – *e.g.*, resource utilization, system stability, or service quality – or to gain an unfair advantage over competing parties using a shared infrastructure? And what if what seem to be the perpetrators of the attack are nothing more than the regular "innocent" users of the system or network?

In recent and on-going work of ours, we exposed a new breed of exploits that target the *dynamics* of a system's operation – *i.e.*, the characteristics of a system's transient behavior as opposed to its limited steady-state capacity or some other known static feature – to achieve the above adversarial goals. In particular, we have shown that a determined adversary could, for example, bleed a system's capacity or significantly reduce service quality by subjecting the system to a fairly low-intensity (but well orchestrated and timed) request stream that causes the system to become very inefficient, or unstable. We instantiated such attacks – which we term Reduction of Quality (RoQ; as in "rock") attacks – on a number of common adaptive components in modern computing and networking systems, including routers, admission controllers, load balancers, and virtual hosting environments. RoQ attacks stand in sharp contrast to traditional brute-force, sustained high-rate DoS attacks, as well as recently proposed low-rate "shrew" attacks that exploit specific protocol settings. Indeed, as our recently-published results (see references) show, RoQ attacks are potentially more *potent* than both DoS-like and shrew-like attacks.

It is important to note that RoQ exploits do not target the data plane of a system or network, but rather its *control plane*. Indeed, the attacks we have mounted exploited not the resources used in the data planes (such as buffers or link capacities), but rather the adaptation/algorithms used in the control plane. In particular, a RoQ attack on a network router exploits the characteristics of the AIMD and RED controllers that regulate the traffic injected at the end points by *normal* connections. Similarily, a RoQ attack on an admission controller exploits the PI or PID controller that regulates the percentage of *normal* requests admitted to a server or service, for example. We emphasize *normal* to underscore the fact that *unlike data-plane attacks, RoQ attacks exploit the control plane in such a way that normal traffic or workloads ends up becoming adversarial!*

The Challenge

Computing systems and networks may exhibit elaborate dynamic behaviors due to resource management strategies in general (as in scheduling, load balancing, caching, *etc.*) and system adaptation strategies in particular (as in admission control, congestion control, *etc.*) These dynamics are quite hard to capture analytically or even empirically. As a result, our models of computing system components tend to abstract away such dynamics and focus instead on static properties obtained through aggregations over time scales that are long enough to hide the transients of adaptation; metrics used to monitor and evaluate a system's performance (such as utilization, delay, jitter, and admission rates) are typically expressed as shapeless mean values, which do not give us insight into the inefficiencies caused by transients over time scales shorter than those used in measuring such metrics.

System dynamics could be *"safely ignored"* if one can ensure that such dynamics will not interfere, or that they will have negligible impact on the overall performance of the system. Such assurances are warranted for closed systems with predictable, non-adversarial workloads. However, in open environments, system dynamics cannot be safely ignored as they could be exploited by adversaries. Notice that while system dynamics could be shown not to interfere with, or significantly impact the fidelity of an end-system or network service under non-adversarial (even if bursty or erratic) workloads, the same could not be said for adversarially-engineered workloads.

The relatively little attention by computing/networking system designers and practitioners to system dynamics is appalling and stands in sharp contrast to how other engineered systems, such as electric and mechanical systems are evaluated. For such systems, the characterization of system dynamics is front and center to protect against oscillatory behaviors and instabilities. Paradoxically, this state of affairs is despite the fact that computing system dynamics *dwarf* in their complexities the dynamics of such traditional systems, not to mention the intentional and possibly adversarial (as opposed to the random "acts of God") nature of harmful stimuli from which the system must be protected.

In our exposition so far, we may have given the reader the impression that exploits of system dynamics are mounted for adversarial purposes, whereby the goal is simply to "hurt" an infrastructure or service. While this is certainly a concern and a focus of ours, it is not the only one. Exploits of system dynamics may be for non-adversarial purposes as well. As we have demonstrated in another recent result of ours, RoQ-like attacks could be mounted for the sole purpose of giving one class of requests an unfair advantage over other classes, without necessarily reducing the efficiency of the underlying infrastructure. Independent of what motivates such exploits, the fundamental questions that need to be answered are similar, and should indeed be central to our quest for trustworthy cyberinfrastructures.

A Road Map

We believe that there is critical need for systems and networking researchers to systematically examine the dynamics of existing as well as proposed control planes of computing systems and network designs: (1) uncovering susceptibilities to malicious behaviors that could compromise the availability and/or efficiency of such systems, (2) quantitatively assessing the extent of such exploits, enabling a solid basis for comparing the trustworthiness of competing designs, (3) developing a general understanding of design principles that could be adopted to protect against such exploits, (4) applying such principles to the design and implementation of common adaptive resource management components that would be resilient to such exploits, (5) the creation of a repository that would act as a knowledge base of known vulnerabilities and defenses in various adaptive systems. The mere availability of agents that exploit specific vulnerabilities (whether for adversarial purposes or not) would be quite instrumental in hardening newly developed systems by enabling practitioners access to what amounts to benchmarks for testing their designs.

Tackling the above goals will also result in important intangible broader impacts, including heightening the research community's appreciation of the importance of system dynamics, which will undoubtedly lead to concrete advancement in other dimensions of basic research.

References [1]

- Mina Guirguis, Azer Bestavros, Ibrahim Matta, and Yuting Zhang. Reduction of Quality (RoQ) Attacks on Internet End-Systems. In *Proceedings of Infocom'05: The IEEE International Conference on Computer Communication*, Miami, Florida, March 2005.
- Yuting Zhang, Azer Bestavros, Mina Guirguis, Ibrahim Matta, and Richard West. Friendly Virtual Machines: Leveraging a Feedback-Control Model for Application Adaptation. In *Proceedings of the 2005 ACM/USENIX Conference on Virtual Execution Environments*, Chicago, Illinois, June 2005.
- Mina Guirguis, Azer Bestavros, and Ibrahim Matta. Bandwidth Stealing via Link-targeted RoQ Attacks. In *Proceedings of CCN'04: IASTED International Conference on Communication and Computer Networks*, MIT, Cambridge, MA, November 2004.
- Mina Guirguis, Azer Bestavros, and Ibrahim Matta. Exploiting the Transients of Adaptation for RoQ Attacks on Internet Resources. In *Proceedings of ICNP'04: The 12th IEEE International Conference on Network Protocols*, Berlin, Germany, October 2004.

[1] Please contact the author (best@cs.bu.edu) for other publications, presentations, and for unpublished results and on-going work.