Information Instability and Disturbance: Goal or Menace

Richard Ford, Mark Bush, Alex Boulatov,
Florida Institute of Technology

Conventional wisdom in the anti-virus arena is that viruses and worms are a unique menace and incapable of serving any useful purpose. Indeed, this position is so strongly entrenched that there is an adage within the anti-virus community which states that there is "no such thing as a good virus". This paper critically re-examines this argument, not by following the well-trodden approach of finding useful application for viruses, but by drawing a close analogy with the effect of disturbances in biological systems. Furthermore, the paper goes one step further, and argues for the deliberate design and release of security threats when environmental conditions merit it.

When considering the effect of computer viruses and worms, one normally considers the damage that these self-replicating programs inflict on individual computers and the network as a whole. Traditional wisdom is that the property of self-replication is inherently dangerous, and the system as a whole would be more stable if such malcode could be eradicated.

Attempts to legitimize malware have mainly centered upon malcode which patches other machines or eradicates other existing worms – essentially, the focus has been on creating self-replicating code that directly serves a "good" purpose. Such work has met with significant controversy within the anti-virus industry; indeed even teaching graduate students how to *write* viruses (let alone releasing them) caused a letter of objection to be sent to the "offending" faculty member's university. The idea that the release of computer viruses in general can be good for the overall health of the global infrastructure is heresy.

This paper examines the problem differently. By drawing an analogy between biological systems and the global computer "ecology" we explore the role of disturbance in driving diversity and robustness in large systems. Based upon this analysis a somewhat startling conclusion is reached: the disturbance caused by viruses and worms is highly beneficial in preventing a catastrophic Internet failure. Empirically, the wave of patching that accompanies large virus outbreaks points towards the validity of this conclusion.

The idea of deliberately perturbing a system before a critical threshold is reached is commonplace in other disciplines. Controlled burns, for example, prevent catastrophic wildfires that cause long-term ecosystem damage. Similarly, we argue that when the density of a particular vulnerability becomes too high it is the *duty* of those tasked with protecting the Internet to perturb the system by releasing an exploit that forces systemic change.

This view will be highly unpopular in traditional anti-virus circles and has little chance of being accepted for publication elsewhere – not because of its lack of scientific validity, but because it is so far against conventions in this area. We would like the opportunity to present it at NSPW because self-replicating code represents a major threat to the Internet, and it would be highly valuable to get feedback from open-minded peers as to the value of our approach.