Involuntary Computing: Hacking the Cloud

Sebastien Goasguen, *Senior Member, IEEE*, Lance Stout, and Michael A. Murphy, *Member, IEEE*

***Abstract*—In this short wild and crazy paper we argue that cloud computing has a lot to learn from current cybercriminals who plague the internet with botnets, malware, and viruses, with the intent of financial gain through extortion or exploitation. We present basic distributed system designs from the perspective of a hacker and argue that the same design features can be and are indeed used in cloud computing: decentralized scheduling, network overlays, heterogneous and intermittent resources, redundancy of task execution, and asynchronous messaging. In botnet architectures, these features are implemented via secure channels, typically using the ubiquitous HTTP protocol. We propose the replication of these nefarious distributed architectures for legitimate applications by means of Kestrel, a job scheduling framework built around XMPP messaging and customized for cloud resources. Our results show that Kestrel scales well over tens of thousands of nodes in heterogeneous environments containing NAT and firewall boundaries, and we envision a one million agent deployment in the near future. Since a botnet is an inexpensive cloud – at a reported cost of $200 a week for 6,000 cores – we hope to capitalize upon this architecture to lower the cost of cloud resources for scientific and educational applications.**

***Index Terms*—cloud computing, virtualization, botnet, XMPP, wild and crazy idea**

I. Introduction

Volunteer computing was coined in the mid nineties and is best represented by the BOINC middleware [1]. In parallel with volunteer computing, grid computing was developed to support the scientific community in search for increasing computational power. By the end of the nineties, another kind of distributed system started to spread: botnets. According to a recent report by Microsoft [2], botnets are currently being rented for an inexpensive cost. The report points out that while early botnets were based on IRC command and control mechanisms, newer botnets are making use of standard http(s) communication mechanisms (~29%), while a small number (~2.3%) use P2P techniques such as Distirbuted Hash Table to build reliable command and control infrastructures. IRC botnets still dominate the market (~38%), but the P2P implementations are becoming more common. P2P techniques have been widely covered in theoretical research but adoption in national grid infrastructure has been extremely slow, while most of the internet traffic is actually caused by P2P-based file sharing [3].

When provisioning cloud resources, which in most cases are virtual machines started on remote sites, we argue that the most difficult challenge is to configure the networking of these intermittent resources in a world with a shortage of IP adresses. Increasingly, cloud resources will be started behind firewall and NAT boundaries and will only feature outbound network connectivity. However, cybercriminals have long worked in such a hostile environment. In fact, they have worked in the toughest of environments: one where the administrator, owner, and users of the resources disagree with the intent of the cloud user. Therefore, we believe that they have developed useful technologies and practices to develop middleware that can be of good use to a non-malicious cloud community.

In particular, we should learn from the cybercriminals, how to deal with adverse network environment, how to scale to a large number of heterogneous workers (millions) running on supercomputers with hundreds of thousands of cores, and how to build a reliable, fault-tolerant system (e.g. P2P command and control layer, redundant execution of tasks). Finally, the many infection vectors used (e.g. drive by download, USB stick, viruses, worms) present an unorthodox view of cloud "API", if anything a non-standard one that has proven to be extremely efficient in coalescing various resources. Compared to clouds, botnets are also much cheaper: a reported $200 a week for 6,000 cores [2]. Unfortunately, cybercriminals do not benchmark their resources with linpack, nor do they care about latency. In general, bandwidth is a much more important metric for their business (e.g. DDoS).

In this paper presentation we will expand on this wild and crazy view that does not adhere to the regular scientific thinking. First, we will introduce a new paradigm: involuntary computing, which deals with resource owners that provide cloud resources without being aware of it. Second, we will introduce a new cloud-specific job scheduling system: Kestrel [4], which leverages cybercriminals' concepts using an XMPP messaging system as a base communication protocol and a P2P layer of reliable job managers.

II. Kestrel

Kestrel has been presented in [4] and [5]. It takes its origin from the observation that instant messaging systems have scaled to hundreds of thousands of clients and can provide interactive communications between agents as well as work seamlessly in complex network environment. If a messaging system can scale and work well behind NAT then it can certainly be a solid communication protocol for many nodes in a grid/cloud environment. The second argument for Kestrel is that virtual machines started on a cluster may not be allocated public IP addresses both due to the shortage of IPv4 addresses and because of the intermittent nature of virtual machines. Moreover, the sheer increase in scale that virtual machines can create via overprovisioning of nodes via consolidation (i.e. more virtual machines than physical cores can be started

S. Goasguen and L. Stout are affiliated with the Clemson University School of Computing, Clemson SC, USA (e-mail: <sebgoa,lstout>@clemson.edu), M. Murphy is affiliated with Coastal Carolina University, Department of Computer Science and Information Systems, Conway, SC, USA (e-mail: mmurphy2@coastal.edu)

on individual nodes) would complicate any attempt to provide inbound traffic routing to the individual VMs, even with IPv6 addressing.

With these motivating factors, Kestrel has been built on XMPP, formerly known as Jabber, the protocol behind Gtalk and other instant messaging services. We have demonstrated efficient job dispatching as well as a scale of over ten thousand agents. Current experiments at Clemson University are reaching 40,000 virtual machines running on a Top 500 production Linux cluster.

In addition the federation capabilities of XMPP servers provide a strong substrate to create inter-cloud mechanisms. Indeed, agents reporting to various XMPP servers can use either a single, shared Kestrel manager or a different manager linked via database sharing mechanisms. This latest development should enable us to use Kestrel in networks of well over a hundred thousand agents and perhaps millions. Figure 1 shows an architecture diagram of Kestrel. Instant Messaging (IM) clients are used to submit and manage jobs, while XMPP servers provide the key communication substrate over which Kestrel managers are deployed to dispatch jobs. Workers started on corporate clouds, grid sites, or local resources can all join the same Kesrel infrastructure. SSL encrypts the communication between workers and servers. Multiple XMPP servers provide a natural scaling mechanism.

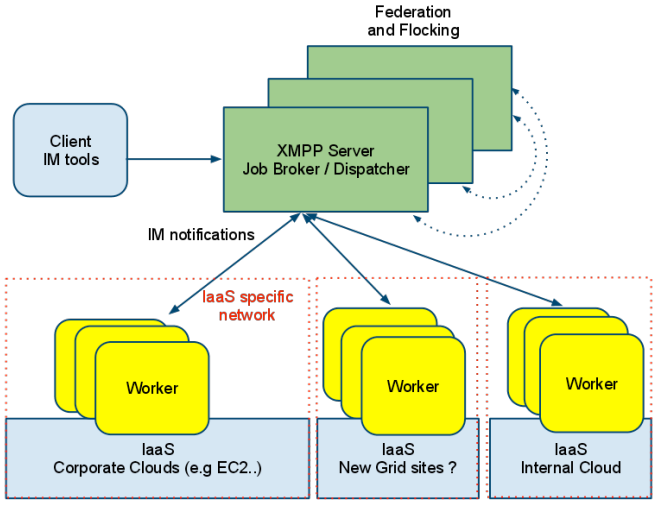

Figure 1. Kestrel [4] Architecture resembles botnet architecture and leverages the XMPP protocol, which has been used by attackers in lieu of IRC channels. Workers started in virtual machines have a Jabber ID and communicate via instant messaging with the managers responsible for job/task scheduling.

Figure 2 shows a snapshot of a standard instant messaging client showing how it can be used to interact with jobs in the Kestrel system. Interaction with jobs is very similar to any standard batch queuing system. A new command line client is also being developed to avoid using a GUI based XMPP client. The STAR collaboration [6] has recently used Kestrel to do one of their largest monte carlo production runs. They used over 400,000 walltime hours on the Clemson Palmetto cluster – a 10,000 core linux cluster – and ran their jobs for one month straight using Kestrel.

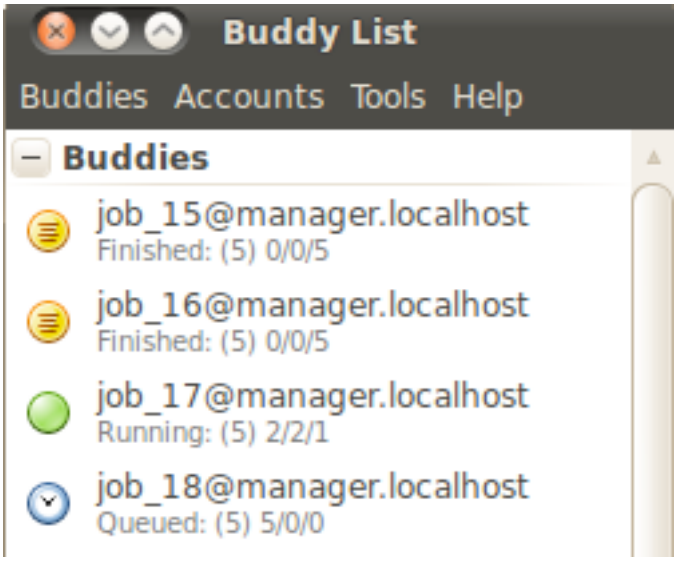

Figure 2. Snapshot of a standard IM client used to control jobs submitted to Kestrel. Each worker has its own Jabber ID (jid). IM commands are used to submit and delete jobs, as well as to check the status of the queues and workers. When a worker goes offline, Kestrel receives a *worker offline* message.

III. Conclusions

While clouds have seen significant growth in the last few years, cybercriminals have long borrowed resources over the Internet and are now offering on-demand scalable networks of resources. At $100 per day for 100,000 bots, the involuntary clouds appear much cheaper than current corporate cloud providers. In addition, involuntary cloud APIs are diverse. Rather than converging toward a standard, these involuntary systems respond quickly to innovations that enable some of the largest clouds presently in use. We argue that by learning from the cybercriminals who design these systems, we can design better cloud infrastructures and find new innovative ways to leverage computational resources across the Internet.

References

[1] D. Anderson, "BOINC: A system for public-resource computing and storage," in *proceedings of the 5th IEEE/ACM International Workshop on Grid Computing*. IEEE Computer Society, 2004, pp. 4–10.

[2] C. T. Anselmi D., Boscovich R. and C. N., "Microsoft Security Intelligence Report January-June 2010," 2010.

[3] J. Grizzard, V. Sharma, C. Nunnery, B. Kang, and D. Dagon, "Peer-to-peer botnets: Overview and case study," in *Proceedings of the first conference on First Workshop on Hot Topics in Understanding Botnets*. USENIX Association, 2007, p. 1.

[4] L. Stout, M. Murphy, and S. Goasguen, "Kestrel: an XMPP-based framework for many task computing applications," in *Proceedings of the 2nd Workshop on Many-Task Computing on Grids and Supercomputers*. ACM, 2009, pp. 1–6.

[5] L. Stout, M. Fenn, M. Murphy, and S. Goasguen, "Scaling virtual organization clusters over a wide area network using the Kestrel workload management system," in *Proceedings of the 19th ACM International Symposium on High Performance Distributed Computing*. ACM, 2010, pp. 692–698.

[6] J. Harris *et al.*, "The STAR experiment at the relativistic heavy ion collider," *Nuclear Physics-Section A*, vol. 566, pp. 277–286, 1994.