The Science of Cloud Computing:
Scientific Foundation for Scalable Trust in Cloud Computing

Hakim Weatherspoon

Cornell University, Department of Computer Science, Ithaca, NY 14853, hweather@cs.cornell.edu

Topics: (3) Data Portability, Consistency, and Management, and (2) Network Support for the Cloud

Current Research Activities Related to Cloud Computing

We are, at last, on the verge of realizing the computer utility vision [CV65, F65, G68]. Its name today is cloud computing. It promises to catalyze the technology economy, revolutionize health care, financial systems, scientific research, and and of course society. Yet many of these uses demand properties that today's cloud platforms either struggle to provide efficiently or lack altogether: energy efficiency, robustness and availability despite failure or attack, security of data and integrity of computation (especially across administrative domains).

- Cloud platforms today are inefficient. Studies of the data centers constructed by Yahoo, IBM, Google and Microsoft reveal that a 10-fold reduction in energy consumption may be possible, with no loss of performance [GHM$^+$09].
- Cloud platforms today are poorly backed up. A US Treasury Dept study [PPG05] found that banks considered "too big to fail" nevertheless fail to maintain synchronous remote backups, so they would be unable to weather certain system failures.
- Cloud platforms today entail significant risk. Software developers and users no longer have physical control over the security and integrity of their computation or data.

My research assumes that such issues reflect technical problems: network layers that are poorly optimized for the patterns of communication that occur in large-scale wide-area cloud settings, operating systems that are hard to customize with the needed communication software, file systems that operate in energy-inefficient ways and that are mismatched to the forms of network communication most appropriate for wide-area backups, and lack of data structures that can be used to ensure integrity of data and computation. My work applies a principled approach to understanding these problems, then validates the understanding gained through analysis by design, implementation, measurement and evaluation of a system. Such an approach offers unique insight and assists into revealing unexpected issues, validating hypotheses, developing remedies.

For example, in recent work (BiFocals: IMC 2010 and DSN 2010) my group constructed a novel network testbed consisting of a series of 10Gbit network loops that run over various lengths from 500 to 8000 miles, but with all the endpoints at Cornell. We used this testbed to study performance properties of non-congested optical wide-area networks (i.e., "private lambdas" that interconnect today's data centers used for cloud computing) and discovered, surprisingly, that packet loss is a serious issue. Common wisdom held that fiber-optic networks did not lose packets. To explain the loss, we built our own optical network interface card using regular physics lab equipment: a pulse pattern generator, laser, beam splitter, an optical to electrical converter, and a high-speed oscilloscope. This setup enabled us to measure packet spacings and delays to a heretofore inaccessible level. And we can now explain the wide-area packet loss, as well as having a basis for improving network interface design (Maelstrom: NSDI 2008).

In a second effort, my group has built a new enterprise file system that leverages a simple observation about log-structured file systems. Since all writes occur at the head of the log or the tail of such file systems, we can reduce power by spinning down lightly loaded disks managing the middle of the log. The power savings can be dramatic (HotOS 2007). Moreover, by combining this file system with my wide-area communication layer, we showed how a heavily loaded data center can maintain a continuous remotely-located backup, without performance loss (indeed, totally transparently) (SMFS: FAST 2009). And in a third effort, we showed that all of these mechanisms can be secured against attack or media failure and integrity of data can be maintained despite lack of physical control over storage resources (Antiquity: Eurosys 2007 and RACS: SOCC 2010).

[CV65] F. J. Corbato and V. A. Vyssotsky. Introduction and Overview of the Multics System. in Proc. *Fall Joint Computer Conference*, 1965, v27.

[F65] R. M. Fano. The MAC System: The Computer Utility Approach. *IEEE Spectrum*, vol. 2, pp. 56-64, January 1965.

[GHM$^+$09]A. Greenberg, J. Hamilton, D. A. Maltz, and P. Patel. The cost of a cloud: Research problems in data center networks. *SIGCOMM Computer Communication Review*, 39(1):68-73, January 2009.

[G68] F. Gruenberger. Computers and Communications; Toward a Computer Utility. Prentice-Hall, 1968.

[PPG05] D. S. Parsons, B. Peretti, and J. M. Grochow. Closing the gap: A research and development agenda to improve the resiliency of the banking and finance sector. In *U.S. Department of the Treasury Study*, March 2005.

Future Research Problems

The trend is a growing spectrum of societally critical, highly sensitive applications are shifting towards cloud computing to benefit from lower costs. These include medical computing applications in which patient care might be directly linked to the quality of the solution (issues arise not just of privacy for sensitive records but also high availability and consistency, for example if a community of home-care patients have their insulin pumps controlled in a closed loop manner, with blood sugars relayed to the cloud system and adjusted insulin doses relayed back). Medical HPC applications are being explored, for example real-time image processing tools to support surgical procedures. These, obviously, would need guaranteed response time. There are other examples: cloud-hosted financial systems, technology to control a smart electric power grid, disaster response, and information systems used by law enforcement or the military. Each brings its own requirements.

The picture, then, is of a class of important applications that will need some form of strong property, from a list that includes preservation of privacy, other kinds of secure access control (for example HIPPA auditability), fault-tolerance, high availability, real-time responsiveness, etc.

Some issues can be solved with simple mechanisms. Some cannot be resolved no matter what we do. But if we focus on properties related to replication, we encounter an intriguing issue: today's cloud is often "inconsistent by design". Many cloud platform developers believe that this design element is central to cloud scalability and cost reduction. With this in mind, we plan to look closely at the consistency issue as it arises in large-scale systems that replicate information or execution, or that indirectly depend on replication, for example as part of a time-critical response loop, or in support of a cloud-scale security or privacy guarantee. We believe this problem can be solved, and that doing so will enable progress on the full spectrum of properties reducible to replication. Moreover, because the central question here is to study scalability in previously overlooked dimensions (for example, by looking at efficient sharing when protocols contend for scarce resources, and at robustness relative to metrics such as freedom of oscillatory behaviors, convoy phenomena, or bottlenecks), our work will also illuminate questions that might usefully be studied in other domains. Success will contribute towards a scientific foundation for scalable trust in cloud computing.

There is a lot of excitment about opportunities related to the cloud, however today's situation is unsatisfactory. Will future banking systems be safe in the cloud? Will cloud-based healthcare be trustworthy? Can we count on the future "smart grid" to robustly deliver electric power if it were in the cloud? A rigorous science of trusted cloud computing will offer convincing responses to such questions.