The Platform Laboratory at Stanford University

John Ousterhout, Sachin Katti, Christos Kozyrakis,
Phil Levis, Nick McKeown, Guru Parulkar, and Mendel Rosenblum

March 3, 2015

Department of Computer Science
Stanford University

Over the last 15 years there has been a revolutionary explosion of new applications, operating at a scale never before imagined and providing features that are changing not just the ways we use information but also the ways that people interact with each other. These applications would not have been possible without the development of an equally revolutionary collection of new *platforms*: general-purpose substrates that simplify the creation of higher-level platforms and applications. Examples of new platforms over the last 15 years include large-scale storage systems such GFS, Bigtable, Hadoop, Cassandra, and memcached, new computational models such as MapReduce, Pregel, and Spark, new approaches to networking such as software-defined networks, new Web development frameworks such as Django, Ruby on Rails, and node.js, and virtual machines, not to mention the Web itself. These new platforms represent the "seed corn" for Silicon Valley, enabling the next generation of new applications.

The revolution is not over. Over the next 10 years there will be new generations of platforms, and these will enable the development of new generations of applications.

We are assembling a diverse group of computer systems faculty at Stanford University into a new Platform Laboratory, whose mission will be to research new platforms that enable future generations of planetary-scale applications. The Platform Laboratory will explore a variety of platforms at many different levels covering the three fundamental legs of all applications: computation, communication, and storage. Here are a few examples of domains ripe for the creation of new platforms:

- The low-latency datacenter. Most of the datacenter platform work to date has focused on the *scale* aspect of datacenters: enabling thousands of machines to work together on a single application. However, little has been done to take advantage of the sub-microsecond speed-of-light distance between machines in a datacenter. The next big opportunity in datacenters is around *latency*: building new hardware and software platforms that operate at microsecond-scale latencies. Low-latency platforms will enable new applications that use information much more intensively than has been possible before. The RAMCloud project at Stanford has begun to create a low-latency storage platform, but there are many more opportunities, such as raising the level of the platform with features such as secondary indexes and multi-object transactions. There is also an opportunity to create a new hardware-software networking platform for communication within low-latency datacenters.
- Software defined networks. We have not yet achieved the SDN goal of turning networks into programmable platforms for innovation. One opportunity is to create a "top down" programming model for networks analogous to compilation in software, where those who own and control networks (such as a datacenter network or a WAN that delivers services from datacenters) can describe how the network is to behave. Given the description, network "compilers" implement the behavior automatically by configuring the data plane. These compilers must target a platform, a "network operating system." This platform, in turn, must meet the requirements of SDN applications and compilers.
- Smart objects. Digital technology is becoming an integral component of physical objects all around us, ranging from cars to thermostats to wearable devices. However, the tools, software, and algorithms for smart objects are still primitive. There is an exciting opportunity to create a software platform for coordinating swarms of smart objects and connecting them with Internet services.
- Mobile interconnect. Mobile networks have evolved from bit/voice-pipes to become the compute interconnect for personal computing: users expect to be able to access all their data and stitch together their applications anytime, anywhere from the cloud over mobile networks. However, the user experience is currently quite poor, primarily due to the difficulty of providing high bandwidth, low latency wide-area mobile access over scarce spectrum, and the rigidity of current mobile networks in tailoring themselves dynamically

to provide the personalized connectivity and context (e.g. location awareness, physical social connectivity) that users and their applications desire. The opportunity here is to evolve mobile networks into a programmable hardware/software substrate for our applications; they should be able to use context from the networks and program them dynamically to realize the data plane connectivity they desire, as well as gather and mine contextual information such as location, proximity etc.

Most universities no longer attempt to develop significant new platforms, out of fear that a meaningful scale of effort for this can only be achieved in an industrial setting. At Stanford we take a different view. We believe that universities can still drive the development of major new platforms, and recent projects such as OpenFlow, RAMCloud, and TinyOS serve as examples. Furthermore, university projects have several advantages over those based in industry. Industrial platforms tend to be developed in support of specific applications. As a result, they are often quirky and narrow in their feature sets. In a university laboratory we can explore general-purpose approaches that might not make sense for a "gotta get this app running yesterday" industrial project. We can search not just for something sufficient for a particular application, but for the best approach, and we can take time to evaluate a platform thoroughly and compare it to alternatives, which helps to develop the science behind the platforms. In addition, we can be more forward-thinking in a university project. For example, we can create new platforms before there is a market for them, and thereby enable the development of new classes of applications that create new markets.

Platform research is not new to Stanford: we have been carrying it out in individual groups for many years. The Platform Lab will bring all of these diverse groups together into one organization that will encourage interactions and synergies between the platforms. For example, there is a natural synergy between the RAMCloud storage system and the underlying networking platform: RAMCloud provides an application to drive and measure networking improvements, and RAMCloud depends on improved networks in order to achieve maximum performance. Platform Lab events such as retreats and seminars will increase interactions between the groups, resulting in a greater flow of ideas. Finally, the Platform Lab will make it easier to share infrastructure and other resources between projects. For example, it may be advantageous to for us to hire staff programmers who can help with implementation issues that are less researchy in nature.

We expect 6-10 faculty members to participate in the Platform Laboratory along with 20-40 graduate students. We expect to add a few more faculty members in addition to those listed on this proposal. John Ousterhout will serve as Faculty Director of the laboratory, and Guru Parulkar will be Executive Director. We are planning a major renovation of the Gates 3A wing to create a new working space that brings together the laboratory members in an environment that stimulates collaboration.

As of March, 2014, the Platform Lab is still in an early formative stage. We expect both the research agenda and our mechanisms for interaction and collaboration to develop over the next 6-12 months.