Research Statement – Eric Keller

I design and build secure and reliable networked systems using a cross-layer approach that draws from networking, operating systems, distributed systems, and computer architecture. My focus is on virtualization and the movement toward cloud-based services. Rather than simply patching existing systems, I question the underlying assumptions to create new designs that are fundamentally more secure and reliable.

1 Research Area and Approach

The idea of a networked society which tears down borders and eliminates barriers to innovation, through a globally connected and ubiquitous network, is inching closer to reality. This type of technology will continue to change the way we live through advancements in all aspects of life – from health care to entertainment to the environment. As we increase our dependence on the network and networked services, the security and reliability of the underlying infrastructure becomes more important. My overall research aim is to create a secure and reliable end-to-end infrastructure for dependable networked services.

Towards this goal, my particular interest is virtualization and its role in improving the underlying infrastructure. While virtualization is currently most associated with server virtualization technology, virtualization is a broad concept captured by the idea of a layer which presents the appearance of something real to a layer above – e.g., virtual memory presents the appearance of a large amount of contiguous memory to applications. This enables (i) efficient resource usage through sharing, (ii) secure execution through isolation, and (iii) an evolvable system where the physical resource and what is using the virtual resource can evolve independently.

Because of these benefits, virtualization is already a core technology in cloud computing – getting us closer to the promise of a ubiquitous, always-on "computer utility." I believe virtualization can be applied to the end-to-end infrastructure of networked services – including the wireless infrastructure people use to connect to the network, the core Internet which enables global communication, and the data centers hosting the services. While virtualization as a concept has great security properties, realizing the ideal in practice is difficult – evident by the many vulnerabilities in server virtualization technology. In my graduate studies I explored both the opportunities and threats in virtualization and will continue this path in ongoing and future work.

My work is systems oriented in nature. Systems-oriented work doesn't naturally fall into hard-and-fast divisions between computer architecture, operating systems, networks, and applications. This is especially true in today's networked services where (i) end-user devices range from embedded sensors to smart phones to computers, (ii) cloud services are dependent on multiple, very distinct networks – e.g., the Internet and data center networks, and (iii) servers heavily utilize advanced computing architectures and virtualization technology. Any given solution might straddle several of these areas. My background mix of hardware, software, and networking is particularly important in this area.

Within my area of systems research, security is of utmost concern and becomes a motivating factor in many aspects of my research. Here, my combination of skills and perspective is valuable for security research as many security problems are "systems security" problems – in some cases, the right solution is to add some cryptographic mechanism, and in others it is to change the design of the system with security as a first-order goal. I view security as an integral aspect of my research.

Finally, great systems research requires great engineering alongside great academic rigor. Engineering skills are important both in being able to master a complex domain (such as routing protocols or virtualization) as well as being able to take ideas to reality. Coupling this with the academic skills of good taste in problem selection, the ability to convey motivation to others, and abstracting the ideas underlying a solution provide a well balanced research approach. I have shown great promise in each regard through building complex systems and taking the approach of challenging existing assumptions.

2 Prior Research Work

In my graduate work I focused on securing the virtualization technology used in cloud computing and applying virtualization concepts to the core Internet.

2.1 Virtualized Cloud Infrastructure without the Virtualization

The key underlying technology in cloud computing infrastructures is virtualization – so much so that many consider virtualization to be one of the key features rather than simply an implementation detail. Unfortunately, the use of virtualization is the source of a significant security concern. The virtualization layer is quite complex and forms a very large trusted computing base that is practically impossible to ship without bugs. A malicious virtual machine can exploit these bugs to attack the virtualization software. Exploiting such an attack vector would give the attacker the ability to obstruct or access other virtual machines and therefore breach confidentiality, integrity, and availability of the other virtual machines' code or data. For cloud computing to gain widespread acceptance, this security concern needs to be addressed.

We did just that with our NoHype architecture where we eliminated the attack surface by going to the extreme of removing the virtualization layer altogether, without sacrificing the key features enabled by virtualization as used in cloud computing infrastructures [1]. While our NoHype architecture is named to indicate the removal of the hypervisor, it has an intended double meaning that it is "no hype" and that we designed, implemented, and evaluated the NoHype architecture on today's hardware [2]. Our prototype utilizes Xen 4.0 to prepare the environment for guest VMs, runs on Intel Nehalem processors, and supports a slightly modified version of Linux 2.6 for the guest OS. Our evaluation with both SPEC and Apache benchmarks shows a roughly 1% performance gain while our security analysis shows that, while there are some minor limitations with current commodity hardware, NoHype is a significant advance in the security of cloud computing.

2.2 Virtualizing the Internet Core to Minimize Disruption

As a first step in exploring the opportunities enabled by virtualization, we virtualized the core Internet to make it more accommodating of planned change as well as more resilient to attacks.

Accommodating Change with Migration (Moving Routers and Links Around)
Network operators often need to repair faulty equipment, deploy new services, perform traffic engineering, and install new equipment to deal with the ever increasing traffic. Unfortunately, this change causes disruption, affecting the availability of networked services. As more devices come online and users consume more traffic, the operational practices of today will not scale. Because of this, researchers have proposed extensions to the routing protocols to make change less disruptive (e.g., by pre-calculating alternate paths). Of course, this suffers from deployment problems as it requires an Internet-wide upgrade.

Instead, we address the root of the problem – the monolithic view of a router where the hardware, software, and links are all considered one entity. We capitalized on the decoupling of the logical and physical that is inherent in virtualization to better accommodate change. We first applied virtualization at the router level, enabling (virtual) routers to freely move from one physical router router to another with VROOM [3]. To eliminate the disruption that would appear when migrating the data plane, we observed that we can have two physical routers with the same data-plane state that are both capable of forwarding traffic. This double data plane allows transitioning traffic to start using the new physical router without disruption. We then applied virtualization at a finer granularity to the individual link and routing session level, enabling links to move around to change the topology with Router Grafting [4]. The mechanisms to virtualize the links largely existed through research in session migration. The key was determining that the mechanisms are the right approach, leveraging exiting building blocks, showing correctness of the routing protocol behavior, and introducing optimizations specific to routing. Our prototype is based on modifications to OpenVZ, Linux 2.6, and the Quagga open source routing software, and supports both an FPGA and software based data plane. In addition to demonstrating migration without disruption, we demonstrated through empirical and analytical evaluation that dynamically re-homing edge links with router grafting can improve the utilization of the available network bandwidth [5].

Becoming More Resilient to Attacks by Virtually Eliminating Router Bugs
The main protocols in the Internet are implemented in software. However, software systems often suffer from bugs and have limited processing capability – both of which are starting to affect the reliability of the underlying network and leave the network vulnerable to attack. Software bugs in routers lead to network outages, security

vulnerabilities, and other unexpected behavior. Rather than simply crashing the router, bugs can violate protocol semantics, rendering traditional failure detection and recovery techniques ineffective. The impact of bugs has become especially evident in the past few years with several high profile outages and periods of global instability. Had these been coordinated attacks, the Internet could possibly have been brought to its knees.

Rather than attempt to make the software bug-free, we instead built a system to operate correctly in the presence of bugs. Building on the years of byzantine fault tolerance research, we tailored software and data diversity (SDD) to the unique properties of routing protocols [6] – using virtualization to make this diversity transparent from the rest of the network. Here, we note two main differences from other settings. First, in routing it is legal for there to be temporary disagreement between the diverse instances during the convergence process. In our system, we tolerate transient inconsistencies during convergence by still doing voting but not over-reacting to temporary inconsistencies. Second, in routing, there is limited dependency on past behavior. We capitalize on the lack of past dependencies to simplify the bootstrapping of new diverse instances and to correct any mistakes. Our Linux-based router hypervisor prototype supports running the three main open-source routing software packages (Quagga, XORP, and BIRD). Experiments with real BGP message traces and traces targeting known bugs demonstrated that our solution scales to large networks and masks buggy behavior.

3 Ongoing and Future Research Directions

I intend to continue pursuing a cross-layer approach toward the goal of dependable networked services with a focus on providing a secure and reliable end-to-end infrastructure. This begins with extending the work I began in my graduate studies on securing the virtualized computing infrastructure as well as applying virtualization to new areas such as data center networks and the wireless access infrastructure.

In support of this, I have capitalized on my post-doctoral appointment to learn more about the funding process and how large, government-funded projects are organized and run. I have written a proposal for an NSF-Medium grant for which I am a principal investigator (PI), I am in the process of submitting another NSF proposal for which I am a PI, and I have started to become involved in the DARPA Mission-oriented Resilient Cloud program which awarded a grant to a team of several UPenn professors and BAE Systems researchers.

3.1 Threats in Virtualization Technology and the Cloud Model

As evident with our NoHype work, there are security threats in the cloud computing model and the technologies used. I intend to continue researching the threats in virtualization technology and the cloud computing model.

As an example, with NoHype we addressed the threat of a malicious party leasing a virtual machine and attacking the virtualization layer from a guest virtual machine. However, we made the assumption that the cloud provider itself is not malicious. While the provider as a business entity (e.g., Amazon) might not be malicious, its employees might be. This insider threat is amplified in the the cloud computing model where the cloud infrastructure provider will employ a large number of people in support of its many customers. This increases the likelihood of a malicious employee slipping through the cracks. We intend to investigate this threat of insider attacks and technological solutions for mitigating the risk. This is no simple undertaking as insiders not only have access to perform the attack, but they also have knowledge of the system. As such, researching this will require looking at the problem from a number of angles. We will explore accountability techniques which hold the cloud provider accountable for any breaches, software analysis techniques to help prevent a breach from occurring in the first place, and hardware solutions for providing a root of trust and enforcing security.

Providing a secure cloud will not come from a single technological innovation. This is why I'm excited to be getting involved with the project at UPenn which was awarded a grant under the DARPA Mission-oriented Resilient Cloud program. This project brings together a number of researchers in a variety of areas and operates under a single, larger objective. This type of collaboration will be important for my research going forward.

3.2 Virtualization for a Secure and Elastic Wireless Access Infrastructure

While there are vulnerabilities with the current virtualization technology that need to be addressed, as a concept, virtualization provides a very nice security property – isolation. It is used in cloud computing to enable companies to run software on a hosted infrastructure. This enables a company that uses the cloud to (i) deploy a new service without having to purchase equipment, (ii) rapidly expand and contract resources to match demand, and (iii)

allocate resources in diverse geographic regions to improve service for its customers. We see great opportunities in applying this model to the domain of wireless infrastructure.

For example, the flexibility enabled by cognitive radio provides a great opportunity for innovation in wireless communication (e.g., the recent work in dynamic spectrum access). However, these innovations are often limited by deployment issues. Borrowing from the cloud computing model of enabling customers to lease server resources, we envision a cognitive access point which allows multiple applications to access a virtualized radio. As an example, consider a femto cell provided by cellular phone carriers to give customers better service in homes where signal strength is low. Rather than creating a dedicated box, the carrier could dynamically expand its footprint by downloading a virtual femto cell to a nearby cognitive access point (which might also happen to function as someone's Wi-Fi router and be used by another provider to deploy a novel cognitive radio algorithm). In addition to exploring the application possibilities that are opened up with this programmability, we intend to research the technology to virtualize the radio resources and the security implications of such a model.

We are in the process of submitting an NSF proposal, for which I am a PI, to explore the potential and limitations of the cognitive access point and a cloud based wireless processing model.

3.3 Virtualized Datacenter Networks for Scalability, Manageability, and Security

Today's simple abstraction of leasing a virtual machine on a hosted infrastructure makes applications simple to create, rapid to expand, and economical to run. However, companies must sacrifice the control they normally have with their own private infrastructure—e.g., the control of security measures within the network as well as the isolation between networks. As a result, many companies continue to run applications on their own private infrastructure, forgoing the many benefits of the public cloud.

Rather than leasing individual virtual machines, we argue that cloud providers should offer customers an entire *cloud resident data center*— effectively providing isolation of the shared network that is comparable to physical isolation while giving each customer the illusion of full control over the network. With the cloud resident data center, the provider presents each customer with a virtualized infrastructure, including a virtualized network topology that the customer can configure. This allows the customer to forward packets over paths with different performance properties, control the sharing of bandwidth across different services, or direct traffic through its own firewalls, load balancers, and intrusion detection systems.

While the security properties are universally desirable, the abstraction of full control may not be appropriate for all companies – the users of infrastructures such as Amazon EC2 are currently dominated by small, new companies who want the simplicity of today's abstraction. However, by going to the extreme of providing the abstraction of controlling an entire data center, we can support all customer requirements and better understand the security threats in public clouds. A set of tools can provide the simpler abstraction to those that do not need the full control. Further, supporting this extreme will help develop technology that is useful independent of the abstraction. Two initial examples include technology to enable a more seamless transition to software defined networks and technology to enable live migration of a collection of virtual machines and associated network without disruption.

Given its scope and broad impact, we have submitted a proposal for an NSF-Medium grant for which I am a PI, to design and build a cloud which provides the cloud resident data center abstraction.

References

[1] E. Keller, J. Szefer, J. Rexford, and R. B. Lee. NoHype: Virtualized cloud infrastructure without the virtualization. In *Proc. International Symposium on Computer Architecture (ISCA)*, 2010.

[2] J. Szefer, E. Keller, R. B. Lee, and J. Rexford. Eliminating the Hypervisor Attack Surface for a More Secure Cloud. In *Proc. ACM Conference on Computer and Communications Security (CCS)*, 2011.

[3] Y. Wang, E. Keller, B. Biskeborn, J. van der Merwe, and J. Rexford. Virtual routers on the move: live router migration as a network-management primitive. In *Proc. ACM SIGCOMM*, 2008.

[4] E. Keller, J. Rexford, and J. van der Merwe. Seamless BGP Migration with Router Grafting. In *Proc. Networked Systems Design and Implementation (NSDI)*, 2010.

[5] E. Keller, M. Schapira, and J. Rexford. Rehoming Edge Links for Better Traffic Engineering. Technical Report TR-917-11, Princeton University Computer Science Department, 2011.

[6] E. Keller, M. Yu, M. Caesar, and J. Rexford. Virtually Eliminating Router Bugs. In *Proc. International Conference on emerging Networking EXperiments and Technologies (CoNEXT)*, 2009.