

Shaya Potter

<http://yucs.org/~spotter/>

Contact Information	Shaya Potter 496 West Charleston Road. Apt 102 Palo Alto, CA 94306	e-mail: spotter@gmail.com cell: (646) 408-8460
Objective	To utilize and improve my skills as a systems engineer, software designer, and researcher	
Education	Columbia University - Doctor of Philosophy: Computer Science Advisor: Dr. Jason Nieh Dissertation: <i>Virtualization Mechanisms for Mobility, Security and System Administration</i> Emphasis on containerization, including security, management and storage. Introduced the concept of building containers out of composable layers, later popularized by Docker.	Mary, 2003 – Dec, 2009
	Columbia University - Master of Science: Computer Science Yeshiva University - Bachelor of Arts: Computer Science	Sept, 2001 – May, 2003 Sept, 1998 – May, 2001
Work Experience	AppOrbit <i>Research Scientist / Advance R&D</i>	San Jose, CA July, 2016 – Dec, 2018
	In charge of building new technologies with a primary focus on enabling AppOrbit's clients to move their legacy infrastructure to a container and Kubernetes environment	
	<ul style="list-style-type: none">• SS-GD - Reimagining Docker images with shared storage 2018 Changes the way Docker images are stored, pushed and pulled. Docker is inefficient when used with large images that are measured in multiple GBs or more. Instead of pushing and pulling the images, SS-GD reimagines what a Docker images is, enabling it to leverage network shared block devices (ex: Ceph) to enable instant pulling of images within a cluster.• AppPorter - Automated modernization of legacy applications into Docker images 2017-2018 Modernizes legacy machines by converting them into proper Docker container images. As opposed to coarse lift and shift methods, AppPorter deconstructs the targeted system into set of layers that correspond to individual software components enabling it to reconstruct multiple container images corresponding to the multiple applications installed on the targeted system.• Infranetes - Cloud VMs as first class citizens within a Kubernetes cluster 2016-2017 Enables Kubernetes to control the lifecycle of a cloud VM (AWS, GCP, Azure, VSphere...) in the same way it controls the life cycle of a Docker container. Additionally, enables a higher level of container security by enabling Kubernetes to run individual pods within individual cloud VMs	
	CoreOS <i>Software Engineer</i>	San Francisco, CA Nov, 2015 – June, 2016
	Member of the Rktnetes team integrating CoreOS's Rkt runtime with the Kubernetes container orchestration system. Work focused on Open Source development	
	<ul style="list-style-type: none">• Extend Cadvisor (a container monitoring tool) to support Rkt by making its interfaces generic, instead of Docker specific, and implementing Rkt versions of these interfaces• General work on implementing and debugging Rkt integration with Kubernetes	
	Bloomberg LP - Bloomberg Law <i>Senior Software Engineer</i>	New York, NY Nov, 2014 – Nov, 2015
	Member of the Integrated Solutions team with cross cutting responsibilities in support of the entire application stack (back-end services in C++ and Java, front-end in Ruby and support services, including Apache, Redis, Memcached...)	
	<ul style="list-style-type: none">• Deep dive into Bloomberg Law's stack in order to identify performance bottlenecks and improve performance. Improved average page load time by 33%• Prototyping/defining usage of AWS services within our enterprise• DevOps improvements to application deployment and management.	

Thomson Reuters
Senior Software Engineer

New York, NY
Nov, 2012 – Nov, 2014

Ingestion, Searching, Screening and Analytics of and on Thomson Reuters' large collection of financial information as part of the Analytics Data Cloud (ADC)

- Ingestion – Improved performance 50-75%
Rewrote ETL tools to handle weekly bulk data loads (200GB compressed, ~1 billion records) and real time loading (~750K messages a day).
- Search - Rewrote data fetching to enable queries to work across absolute and relative time periods
Data is stored with absolute dates and periods, however analysts desire to perform searches on relative dates and periods. Designed and led small team that implemented a generic infrastructure that enables any content set to be stored in an absolute date/period format but efficiently retrieved based on relative dates.
- Screening and Analytics – Improved SQL query design to enable efficient execution of complex analytics
ADC enables screening of securities on both the simple/database view of the data and complex analytics that are calculated from the underlying data. Implemented a number of complex analytics as well improved underlying screening performance by pushing down more complicated screens, such as currency conversion, into the SQL queries.

IBM T.J. Watson Research Center
Post-Doctoral Researcher

Hawthorne, NY
Dec, 2009 – Nov, 2012

Services research with a focus on discovery and migration for cloud computing.

- Workload Transformation Analysis (WTA) - Created consulting methodology and tools to help determine what machines can be migrated to a public and private clouds.
WTA defines what a consultant needs to gather from a client and provides tools that determine which of the client's machines are candidates for a migration, what the difficulty of the migration will be, and a cost comparison between their current data center costs and the clouds they are targeting.
- Galapagos discovery - Extended the Galapagos discovery toolkit to be able to run on dormant (not actively running) virtual machines.

Network Computing Lab, Columbia University
Lab member - Graduate Research Assistant

New York, NY
May, 2001 – Dec, 2009

Engaged in a variety of research, primarily involving operating systems, process migration, containerization, security, and system administration.

- Strata – A novel file-system framework that changes how one manages large numbers of containers and machines. Introduces the Virtual Layered File System (VLFS) that enables defining a file-system as a composable set of shared layers. This approach to container file-systems was later popularized by Docker.
- Apiary – Leverages Strata to create a secure desktop composed of individual applications running within isolated containers. Introduces ephemeral containers to provide a safe environment for application execution with untrusted and possibly malicious data. This impedes exploits from accessing other parts of the system and from having a persistent impact.
- Helped develop ZAP, a system for migrating computing environments intact from one computer to another. Involving ensuring that processes, file systems, devices all migrated together. ZAP inspired by CRIU the modern mechanism for checkpoint/restart of Linux applications.
- AutoPod – Extended ZAP to support heterogenous process migration to support migrating processes between different kernel versions. This enables processes to continue running across security necessitated kernel upgrades.

- PeaPod – An extension of the ZAP architecture to support secure isolation and migration of cooperating processes. PeaPod enables the easy creation of least privileged environments for each component of an application. PeaPod enabled Linux to have a namespaced root user who would act as root within the container for file-system and network access, but had no other root privileges on the host machine.
- *Pod – Leveraging the ZAP architecture to support persistent and portable application specific environments (desktop, web browser, multimedia content) that are stored on removable storage, such as a thumb drive or iPod.
- DejaView – Brings DVR semantics to the computer desktop. DejaView records a visual and textual record of a user’s desktop. Users can search through both records of their computing history. DejaView’s desktop is able to “time-travel” by letting the user recreate the any previous state and continue computing from the state as it existed at that point in time.
- ISE-T – Prevents malicious and accidental system administration faults from entering a system by applying the two person control model to system administration.

IBM T.J. Watson Research Center

Hawthorne, NY

Summer Intern

Summer 2006

Worked on productization of SoulPad, IBM’s research project that leverages virtual machines to enable a user to carry one’s desktop computing environment between physical machines

Microsoft

Redmond, WA

Summer Software Development Engineer Intern

Summer 2004

Worked on Windows Terminal Services Client, creating a cross platform core. This enables a single highly tested code base to be used across all of Microsoft’s supported platforms.

IBM T.J. Watson Research Center

Hawthorne, NY

Summer Intern

Summer 2003

Investigated algorithms related to peering arrangement and search propagation in P2P networks.

IBM T.J. Watson Research Center

Yorktown Heights, NY

Summer Intern

Summer 2001

Worked on IBM’s Linux WatchPad prototype

- Implemented tools to analyze the WatchPad’s power usage.
- Created a reduced memory footprint version of the WatchPad software.
- Created a native development environment for the WatchPad, in order to allow compilation of many programs that will not cross-compile.

IBM Extreme Blue

Cambridge, MA

Summer Intern

Summer 2000

Built trace and benchmarking tools to predict application performance across multiple platforms.

Student Organization of Yeshiva Seforim (“Book”) Sale

New York, NY

Software Designer and Developer

1999 – 2001

Designer and developer of inventory and point of sale software to support this student run sale.

- Written on Linux and MySQL in TCL/Tk, Perl and C. Complete retail solution including bar code labeling, credit card processing and receipt printing.
- Software handled over \$1 million in sales that occurred in the sales 2 week run each year.
- In charge of ensuring all IT services ran smoothly during the sale.

Naval Research Laboratory, Information Technology Division
Engineering Aid

Washington, DC
1995 – 2000

Intern for Fleet Network Security group of the Center for High Assurance Computing.

Major Projects

- Designed a large district wide firewall and load balancing infrastructure that US Navy deployed around all their computing resources in San Diego. 1999 – 2000
- Investigated many Network Intrusion Detection systems and other network security tools in order to determine if and how the US Navy could use them. 1998
- Participated in software evaluation and development for the US Navy's IT-21 Initiative for the 21'st Century. 1997 – 1998
- Continued work started as a SEAP intern and helped build the Secure Tactical Access Terminal (STAT) which was deployed on the aircraft carrier USS Theodore Roosevelt to rave reviews. 1996
- NRL & George Washington University SEAP Intern 1995
 - Worked on porting Unix applications to MultiLevel Secure Workstations running HP-UX and SCO Unix

Honors and Awards

- LISA 2005 Best Student Paper Award 2005
- Mobicom 2004 Best Student Paper Award 2004
- Golding Distinguished Scholar (full 4 year undergraduate scholarship) 1998 – 2001
- ACM Programming Competition 5th place (NY Region) 1999, 2000
- SEAP's Most Outstanding Intern 1995

Selected Publications

- Shaya Potter and Jason Nieh, "Improving Virtual Appliance Management through Virtual Layered File Systems", *Proceedings of the 25th Large Installation System Administration Conference (LISA 2011)*, Boston, MA, December 4-9, 2011.
- Shaya Potter and Jason Nieh, "Apiary: Easy-to-Use Desktop Application Fault Containment on Commodity Operating Systems", *Proceedings of the 2010 USENIX Annual Technical Conference (USENIX 2010)*, Boston, MA, June 22-25, 2010.
- Shaya Potter, Steven M. Bellovin, and Jason Nieh, "Two-Person Control Administration: Preventing Administration Faults Through Duplication", *Proceedings of the 23rd Large Installation System Administration Conference (LISA 2009)*, Baltimore, MD, November 1-6, 2009.
- Shaya Potter, Jason Nieh and Matthew Selsky, "Secure Isolation of Untrusted Legacy Applications", *Proceedings of the Twenty-first Large Installation System Administration Conference (LISA 2007)*, Dallas, TX, November 11-16, 2007.
- Oren Laadan, Ricardo Baratto, Dan Phung, Shaya Potter, and Jason Nieh, "DejaView: A Personal Virtual Computer Recorder", *Proceedings of the 21st ACM Symposium on Operating Systems Principles (SOSP 2007)*, Stevenson, WA, October 14-17, 2007.
- Shaya Potter and Jason Nieh, "Breaking the Ties that Bind: Process Isolation and Migration", *login Vol 30(6)*, pp 14-17, Dec. 2005
- Shaya Potter and Jason Nieh, "Reducing Downtime Due to System Maintenance and Upgrades", *Proceedings of the 19th Large Installation System Administration Conference (LISA '05)*, San Diego, CA. Dec. 4-9 2005. (Best Student Paper Award)
- Shaya Potter and Jason Nieh, "WebPod: Persistent Web Browsing Sessions with Pocketable Storage Devices", *Proceedings of the 14th International World Wide Web Conference (WWW 2005)*, Chiba Japan. May 10-14 2005.

Ricardo Baratto, Shaya Potter, Gong Su and Jason Nieh, "Mobidesk: Mobile Virtual Desktop Computing", *Proceedings of the Tenth Annual ACM International Conference on Mobile Computing and Networking (Mobicom 2004)*, Philadelphia, PA. Sept 26-Oct.1 2004 (Best Student Paper)

Skills

Operating systems/Platforms: Linux, Windows

Languages: C, C++, Perl, Java, C#, Go, Unix Shell Scripting, SQL

Selected Patents Systems, methods, means, and media for recording, searching, and outputting display information:

US Patent 8,214,367 , Filed: Feb 27, 2008, Issued: Jul 3, 2012

Detailed Inventory Discovery on Dormant Systems: US Patent 8,583,709, Filed: Mar 18, 2010, Issued: Nov 12, 2013

Methods, Systems, and Media for Application Fault Containment: US Patent 8,589,947 , Filed: May 10, 2011, Issued: Nov 19, 2013

Workload-to-Cloud migration analysis based on cloud aspects: US Patent 9,495,649 , Filed: May 24, 2011, Issued Novemeber 15, 2016