

CURRICULUM VITAE

## CONTACT INFORMATION

**Address:** 27 Yakinton St., Haifa, 34792, Israel  
**Telephone:** +972-54-6976101 (cell), +972-4-8257062 (home)  
**E-mail:** mulix@mulix.org, muli@cs.technion.ac.il  
**Web:** <http://www.mulix.org>

## EDUCATION

**2011–current** Graduate studies in Computer Science, Technion—Israel Institute of Technology  
**2010** B.A. in Computer Science, *cum laude*, Open University, Israel

## PROFESSIONAL HISTORY

**2009–2010** *Manager, Virtualization and System Architecture—IBM Haifa Research Lab.*

I managed a group of ten to fifteen researchers working in the areas of machine virtualization, operating systems, and high-speed interconnects. In addition to making technical contributions to most research projects in the group, I was responsible for all research activities carried out by individual researchers, set technical direction to the group, secured internal (from IBM) and external (from, e.g., EU) grant funding, and established collaborations with industry and academia.

**2003–2012** *Researcher—IBM Haifa Research Lab.*

I conceived, researched, designed, and implemented systems.

- **I founded the Turtles nested virtualization research project.** I lead the project from its inception in 2008, and made key contributions to the design and implementation. Nested virtualization is the ability to run multiple hypervisors simultaneously and efficiently on architectures (such as x86) that support only a single hypervisor in hardware. We achieved this ability by employing *multi-dimensional page tables* for MMU virtualization and *multi-level device assignment* for I/O virtualization. The resulting performance was within 6-8% of single-level (non-nested) virtualization. A paper describing the project won the *Jay Lepreau Best Paper award* at the prestigious OSDI 2010 conference [19].
- **I conducted pioneering research into high performance I/O virtualization** starting in 2006. I researched, designed, and implemented operating system and hypervisor support for the Calgary, CalIOC2, and VT-d IOMMUs on the x86-64 platform, exploring for the first time the benefits and costs of *direct device assignment* for unmodified virtual machines [2, 4, 16, 26]. This work won an IBM Research accomplishment, and I personally received an Outstanding Technical Achievement Award for it. I also explored other software and hardware methods for efficient I/O virtualization [10, 17, 21, 22, 23].
- **I made substantial contributions to other research projects:** TCP acceleration via flexible hardware/software split [1] and via dedicated CPU cores [18, 23], using machine learning methods for monitoring the well-being and performance of virtual machines [7, 11, 12], combining memory and storage checkpointing of VMs to enable “virtual machine time travel” [6], adding operating system support for programmable devices [8], and out-of-band storage intrusion detection [14].
- **I designed and implemented system software** such as firmware for the “IP Only Server” [3], an Ethernet driver for an experimental adapter, and embedded firmware for the DS6000 storage controller.
- **I gave invited talks and taught courses** at the IDF, the Technion, Tel-Aviv University, Ben-Gurion University, Intel, and IBM, including a 16-hour course on machine virtualization and 8-hour courses on operating systems and on the Linux kernel.

**2002** *Linux kernel Engineer—Qlusters, Inc., Tel Aviv, Israel.*

I worked on the core of Qlusters’ Linux-based cluster OS, including on its distributed shared memory subsystem.

**1997–2000** *Programmer and Programming Team Leader—IDF, Intelligence Corps.*

I worked on parsing, understanding, and representing large amounts of data.

## PUBLICATIONS

- 27 A. Gordon, N. Amit, N. Har'El, M. Ben-Yehuda, A. Landau, A. Schuster, D. Tsafirir "*ELI: Bare-Metal Performance for I/O Virtualization*". In *17th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '12)*
- 26 M. Ben-Yehuda, E. Borovik, M. Factor, E. Rom, A. Traeger, B.-A. Yassour "*Adding Advanced Storage Controller Functionality via Low-Overhead Virtualization*". In *10th USENIX Conference on File and Storage Technologies (FAST '12)*
- 25 O. Agmon Ben-Yehuda, M. Ben-Yehuda, A. Schuster, D. Tsafirir "*Deconstructing Amazon EC2 Spot Instance Pricing*". In *3rd IEEE International Conference on Cloud Computing Technology and Science (CloudCom '11)*
- 24 M. Hines, A. Gordon, M. Silva, D. Da Silva, K. D. Ryu, M. Ben-Yehuda "*Applications Know Best: Performance-Driven Memory Overcommit With Ginkgo*". In *3rd IEEE International Conference on Cloud Computing Technology and Science (CloudCom '11)*
- 23 N. Amit, M. Ben-Yehuda, D. Tsafirir, A. Schuster "*vIOMMU: Efficient IOMMU Emulation*". In *The 2011 USENIX Annual Technical Conference (USENIX ATC '11)*
- 22 A. Landau, M. Ben-Yehuda, A. Gordon "*SplitX: Split Guest/Hypervisor Execution on Multi-Core*". In *The 3rd Workshop on I/O Virtualization (WIOV '11)*
- 21 A. Gordon, M. Ben-Yehuda, D. Filimonov, M. Dahan "*VAMOS: Virtualization Aware Middleware*". In *The 3rd Workshop on I/O Virtualization (WIOV '11)*
- 20 A. Gordon, M. Hines, D. Da Silva, M. Ben-Yehuda, M. Silva, G. Lizarraga "*Ginkgo: Automated, Application-Driven Memory Overcommitment for Cloud Computing*". In *Runtime Environments/Systems, Layering, and Virtualized Environments (RESoLVE) Workshop (RESoLVE '11)*
- 19 M. Ben-Yehuda, M. D. Day, Z. Dubitzky, M. Factor, N. Har'El, A. Gordon, A. Liguori, O. Wasserman, B.-A. Yassour. "*The Turtles Project: Design and Implementation of Nested Virtualization*". In *9th USENIX Symposium on Operating Systems Design and Implementation (OSDI '10)*  
**Received the Jay Lepreau Best Paper award**  
**Received the IBM Research 2010 Pat Goldberg Memorial Best Paper Award**
- 18 L. Shalev, J. Satran, E. Borovik, M. Ben-Yehuda "*IsoStack—Highly Efficient Network Processing on Dedicated Cores*". In *The 2010 USENIX Annual Technical Conference (USENIX ATC '10)*
- 17 N. Amit, M. Ben-Yehuda, B.-A. Yassour "*IOMMU: Strategies for Mitigating the IOTLB Bottleneck*". In *The Sixth Annual Workshop on the Interaction between Operating Systems and Computer Architecture (WIOSCA '10)*
- 16 B.-A. Yassour, M. Ben-Yehuda, O. Wasserman "*On the DMA Mapping Problem in Direct Device Assignment*". In *The 3rd Annual Haifa Experimental Systems Conference (SYSTOR '10)*
- 15 A. Landau, D. Hadas, M. Ben-Yehuda "*Plugging the Hypervisor Abstraction Leaks Caused by Virtual Networking*". In *The 3rd Annual Haifa Experimental Systems Conference (SYSTOR '10)*
- 14 M. Allalouf, M. Ben-Yehuda, J. Satran, I. Segall "*Block Storage Listener for Detecting File-Level Intrusions*". In *The 26th IEEE Symposium on Massive Storage Systems and Technologies: Research Track (MSST '10)*
- 13 D. Cohen, F. Petrini, M. D. Day, M. Ben-Yehuda, S. Hunter, U. Cummings "*Applying Amdahl's Other Law to the Data Center*". In *IBM Journal of Research and Development special issue on Hybrid Computing Systems, 53(5), 2009 (IBM JR&D)*
- 12 M. Ben-Yehuda, D. Breitgand, M. Factor, E. Kolodner, V. Kravtsov, D. Pelleg "*NAP: A Building Block for Remediating Performance Bottlenecks via Black Box Network Analysis*". In *The 6th IEEE International Conference on Autonomic Computing (ICAC '09)*

- 11 N. Parush, D. Pelleg, M. Ben-Yehuda, P. Ta-Shma “*Out-of-band detection of boot-sequence termination events*”. In *The 6th IEEE International Conference on Autonomic Computing (ICAC '09)*
- 10 B. Rochwerger, D. Breitgand, E. Levy, A. Galis, K. Nagin, I. M. Llorente, R. Montero, Y. Wolfsthal, E. Elmroth, J. Caceres, M. Ben-Yehuda, W. Emmerich, F. Galan “*The RESERVOIR Model and Architecture for Open Federated Cloud Computing*”. In *IBM Journal of Research and Development special issue on Internet and Enterprise-Scale Data Centers, 53(4), 2009 (IBM JR&D)*
- 9 J. Satran, L. Shalev, M. Ben-Yehuda, Z. Machulsky “*Scalable I/O—A Well-Architected Way to Do Scalable, Secure and Virtualized I/O*”. In *The First Workshop on I/O Virtualization (WIOV '08)*
- 8 Y. Weinsberg, D. Dolev, T. Anker, M. Ben-Yehuda, P. Wyckoff “*Tapping into the Fountain of CPUs—On Operating Systems Support for Programmable Devices*”. In *Thirteenth international conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '08)*
- 7 D. Pelleg, M. Ben-Yehuda, R. Harper, L. Spainhower, T. Adeshiyan “*Vigilant—Out-of-band Detection of Failures in Virtual Machines*”. In *ACM SIGOPS Operating Systems Review, 42(1) (OSR)*
- 6 P. Ta-Shma, G. Laden, M. Ben-Yehuda, M. Factor “*Virtual Machine Time Travel Using Continuous Data Protection and Checkpointing*”. In *ACM SIGOPS Operating Systems Review, 42(1) (OSR)*
- 5 M. Ben-Yehuda, E. Van Hensbergen “*Open Source as a Foundation for Systems Research*”. In *ACM SIGOPS Operating Systems Review, 42(1) (OSR)*
- 4 M. Ben-Yehuda, J. Xenidis, M. Ostrowski, K. Rister, A. Bruemmer, L. Van Doorn “*The Price of Safety: Evaluating IOMMU Performance*”. In *The 2007 Ottawa Linux Symposium (OLS '07)*
- 3 M. Ben-Yehuda, O. Goldshmidt, E. Kolodner, Z. Machulsky, V. Makhervaks, J. Satran, M. Segal, L. Shalev, I. Shimony “*IP Only Server*”. In *The 2006 USENIX Systems Practice and Experience Track (USENIX ATC '06)*
- 2 M. Ben-Yehuda, J. Mason, O. Krieger, J. Xenidis, L. Van Doorn, A. Mallick, J. Nakajima, E. Wahlig “*Utilizing IOMMUs for Virtualization in Linux and Xen*”. In *The 2006 Ottawa Linux Symposium (OLS '06)*
- 1 L. Shalev, V. Makhervaks, Z. Machulsky, G. Biran, J. Satran, M. Ben-Yehuda, I. Shimony “*Loosely Coupled TCP Acceleration Architecture*”. In *The 14th IEEE Symposium on High-Performance Interconnects (HOTI '06)*

## TECHNICAL REPORTS

- 2 B.-A. Yassour, M. Ben-Yehuda, O. Wasserman “*Direct Device Assignment for Untrusted Fully-Virtualized Virtual Machines*”. In *IBM Research Report H-0263*
- 1 The RESERVOIR Seed Team “*RESERVOIR—An ICT Infrastructure for Reliable and Effective Delivery of Services as Utilities*”. In *IBM Research Report H-0262*

## GRANTS AWARDED

IBM PI for “IOLanes: Advancing the Scalability and Performance of I/O Subsystems in Multicore Platforms”, Seventh Framework Programme (FP7) (3/2010–3/2013; part of a €4,260,426 EU project with Forth, UPM, BSC, Intel, Neurocom). Web: <http://www.iolanes.eu>.

## AWARDS AND HONORS

**2012** *EuroSys Student Travel Prize*

Received a EuroSys competitive travel award for a European student to attend ASPLOS '12.

**2011** *IBM Research Pat Goldberg Memorial Best Paper Award*

Received the 2010 Pat Goldberg Memorial Best Paper Award in Computer Science, Electrical Engineering and Math for the paper “The Turtles Project: Design and Implementation of Nested Virtualization”.

**2010** *Jay Lepreau Best Paper Award*

Received the Jay Lepreau Best Paper Award at OSDI 2010 for the paper “The Turtles Project: Design and Implementation of Nested Virtualization”.

**2010**            *Outstanding Technical Achievement Award (OTAA)*

Received an IBM Outstanding Technical Achievement Award for contributions to “Linux and Open Virtualization”.

**2009**            *Research Accomplishment*

Received an IBM Research Accomplishment for contributions to “Linux and Open Virtualization”.

**2008**            *Research Accomplishment*

Received an IBM Research Accomplishment for contributions to “Availability Manager for virtual machine availability systems management in IBM Systems Director”.

**2006–2011**    *Master Inventor*

Appointed Master Inventor in recognition of extraordinary contributions to IBM’s intellectual property.

## PROFESSIONAL SERVICE

**Workshop co-chair:**    2<sup>nd</sup> Workshop on I/O Virtualization (WIOV ’10), co-located with ASPLOS  
**Workshop co-chair:**    1<sup>st</sup> Workshop on I/O Virtualization (WIOV ’08), co-located with OSDI  
**Workshop co-chair:**    1<sup>st</sup> Annual Haifa Systems and Storage Conference (SYSTOR ’07)  
**Journal guest editor:**    ACM SIGOPS Operating Systems Review, special issue on Research and Developments in the Linux kernel  
**Journal guest editor:**    ACM SIGOPS Operating Systems Review, special issue on Best Papers from VEE and Best Papers from WIOV  
**Information Director:**    ACM SIGOPS (2011–current)  
**PC member:**            2012 ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments (VEE ’12)  
**PC member:**            1<sup>st</sup> International Workshop on Security and Privacy in Social Networks (SPSN ’11), co-located with IEEE Social Com  
**PC member:**            3<sup>rd</sup> Workshop on I/O Virtualization (WIOV ’11), co-located with USENIX ATC  
**PC member:**            2011 International Conference on High Performance Computing, Networking, Storage and Analysis (SC ’11)  
**PC member:**            2011 USENIX Annual Technical Conference (USENIX ATC ’11)  
**PC member:**            2011 ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments (VEE ’11)  
**PC member:**            Workshop on Micro Architectural Support for Virtualization, Data Center Computing, and Clouds (MASVDC ’10), co-located with MICRO  
**PC member:**            24th International Conference on Supercomputing (ICS ’10)  
**PC member:**            3<sup>rd</sup> Annual Haifa Experimental Systems Conference (SYSTOR ’10)  
**PC member:**            2<sup>nd</sup> Annual Israeli Experimental Systems Conference (SYSTOR ’09)  
**PC member:**            Workshop on Managed Many-Core Systems (MMCS ’08), co-located with HPDC  
**Publicity chair:**        2<sup>nd</sup> Annual Israeli Experimental Systems Conference (SYSTOR ’09)  
**Member:**                USENIX, ACM SIGOPS, EuroSys

## FREE SOFTWARE

Contributor to free software projects including the Linux kernel and the Xen and KVM hypervisors.

## ISSUED PATENTS

- 9 US Patent #8,095,569 “*Customized Context Menu For Files Based On Their Content*”
- 8 US Patent #8,055,951 “*System, Method And Computer Program Product For Evaluating A Virtual Machine*”
- 7 US Patent #8,001,342 “*Method For Storing and Restoring Persistent Memory Content and Virtual Machine State Information*”

- 6 US Patent #7,970,736 *“System, Method and Computer Program Product for Storing Transient State Information”*
- 5 US Patent #7,900,003 *“System, Method And Computer Program Product For Storing An Information Block”*
- 4 US Patent #7,757,280 *“Method And System For Memory Protection And Security Using Credentials”*
- 3 US Patent #7,636,800 *“Method And System For Memory Address Translation And Pinning”*
- 2 US Patent #7,603,392 *“System, Method And Computer Program Product For Storing Transient State Information”*
- 1 US Patent #7,600,093 *“Device, Method And Computer Program Product For Multi-Level Address Translation”*

## **PATENT APPLICATIONS**

- 17 US Patent Application 2011/0302,592 *“Condition Based Detection of No Progress State of an Application”*
- 16 US Patent Application 2011/0258,493 *“Detecting No Progress State of an Application”*
- 15 US Patent Application 2011/0276,578 *“Obtaining File System View in Block-Level Data Storage Systems”*
- 14 US Patent Application 2011/0178,153 *“Efficient Multi-Core Processing of Events”*
- 13 US Patent Application 2011/0138,127 *“Automatic Detection of Stress Conditions”*
- 12 US Patent Application 2011/0107,008 *“Memory Management in a Nested Virtualization Environment”*
- 11 US Patent Application 2010/0047,544 *“Para-virtualization in a Nested Virtualization Environment”*
- 10 US Patent Application 2010/0169,716 *“Managing Confidence Levels In A Computing System”*
- 9 US Patent Application 2010/0169,882 *“Multiple Layers Of Virtualization In A Computing System”*
- 8 US Patent Application 2010/0153,926 *“Operating System Aided Code Coverage”*
- 7 US Patent Application 2010/0094,990 *“Platform-Level Indicators Of Application Performance”*
- 6 US Patent Application 2009/0276,571 *“Enhanced Direct Memory Access”*
- 5 US Patent Application 2009/0204,960 *“System, Method And Computer Program Product For Accessing A Memory Space Allocated To A Virtual Machine”*
- 4 US Patent Application 2009/0164,990 *“Apparatus For And Method For Real-Time Optimization Of Virtual Machine Input/Output Performance”*
- 3 US Patent Application 2009/0083,737 *“Device, System, And Method Of Classifying A Workload Of A Software Service”*
- 2 US Patent Application 2008/0196,026 *“Device, Method And Computer Program Product For Executing A Migrated Execution Context By A Storage Controller”*
- 1 US Patent Application 2007/0064,935 *“Generating Key Information For Mutual Access Among Multiple Computers”*

## **LANGUAGES**

Fluent in Hebrew and English.

## **REFERENCES**

Available upon request.