Novel Orchestration of Virtualization to Improve Cybersecurity: Software Defined Infrastructure (SDI) as a Foundation for Clean-Slate Computing Security

Robert Ames and Lewis Shepherd[1]

[1] VMware Research Group, 12100 Sunset Hills, Reston, Virginia, USA
rames@vmware.com, lshepherd@vmware.com

Abstract. In 2017 VMware Research launched a three-year research program in collaboration with the U.S. National Science Foundation (NSF), to achieve "Clean-Slate Computing Security" by combining multiple threads of cloud-scaled and enterprise-grade virtualization and data-abstraction capabilities designed to advance the state of the art in key areas relevant to the national security and multinational cybersecurity. Specifically, this paper describes the program's use of Software Defined Infrastructure (SDI) in novel fashion to secure network and data traffic through richer and deeper situational awareness, network micro-segmentation and automation all the way from the core datacenter, to the public clouds, to the edge of IoT and back. The SDI-CSCS approach extends the VMware core technology stack across the cybersecurity and virtualization landscape, and the paper outlines the academic and lab advances already reported. Further, the paper describes new uses of Machine Learning and algorithmic exploitation of the behaviors of the virtualized infrastructure, demonstrating more responsive and predictive systems in the event of attacks. The paper also describes the extension of the SDI approach for Internet of Things and edge-computing architecture security, for Kubernetes container environments, and points to potential security uses involving SDI in trusted hardware extensions and blockchain environments.

Keywords: research; software-defined; SDI; SDN; virtualization; machine-learning orchestration; networks; network-functionality-virtualization; NFV.

1 Cyber Security Research in Collaboration with Academia

The NSF/VMware Partnership on SDI-CSCS is a three-year university research program co-sponsored by the National Science Foundation and VMware.[1] The program goal is to foster novel, transformative, multidisciplinary research that spans systems, networking, and security with the aim of exploring and creating groundbreaking new approaches to security based on the concept of SDI.

---

[1] See "NSF/VMware Partnership on Software Defined Infrastructure as a Foundation for Clean-Slate Computing Security (SDI-CSCS)," Program Solicitation for National Science Foundation 16-582, National Science Foundation Directorate for Computer and Information Science and Engineering, Division of Computer and Network Systems; https://www.nsf.gov/pubs/2016/nsf16582/nsf16582.htm

As the digital and physical worlds become increasingly intertwined, the real-world consequences of cyber-threats are becoming more pronounced. To mitigate foreseeable risks, fundamental advances in security are needed. This program was launched to explore the hypothesis that software defined infrastructure (SDI) enables realistic opportunities to revisit and improve the foundations of end-to- end computing security.

SDI is an architectural approach in which compute, storage, and networking resources are virtualized; that is, abstractions of physical capabilities are made available to applications or higher-level services in a way that is decoupled from the underlying physical infrastructure. To date, SDI has been realized most fully in the context of datacenters, but it can also be viewed as a foundation for related emerging contexts such as the Internet of Things (IoT). Novel security properties of SDI have been demonstrated, and meanwhile, compute, storage, and network virtualization techniques are rapidly maturing. The National Science Foundation understood the value of this opportunity to systematically explore and identify the full potential of SDI as a new foundation for clean-slate computing security (CSCS), and to transition research findings into practice.

This NSF/VMware partnership has combined the NSF Directorate for Computer & Information Science & Engineering long experience in developing and managing successful large, diverse research portfolios with VMware's significant expertise in SDI, virtualization technology, distributed systems, cloud computing, and other aspects of large-scale software infrastructure and infrastructure management. Over the three years of the collaboration, NSF and VMware are supporting multiple research projects with funding of up to $3,000,000 each, with NSF and VMware co-funding each project.

2 Technical Background of SDI-CSCS

Today's network and computing infrastructure rests on inadequate foundations. In particular:

- Much of the existing compute infrastructure is notorious for the omission of security as a primary design consideration;
- Client computing and human error remain weak links in the overall security chain, and as a result, application and content flaws continue to be primary delivery vectors for malicious payloads;
- The complexity and fragmentation of security technology in distributed systems has led to inadequate visibility and control over lateral movement of attacks - and these issues will become even more pronounced in the context of the emerging Internet of Things (IoT); and
- Security and compliance policy controls are still largely defined in terms of low-level entities such as processes, access control lists, firewall rules, and IP addresses.

As the digital and physical worlds become increasingly intertwined through ubiquitous "cyber-physical systems," and as the emerging IoT expands to potentially tens of billions of connected devices, the research challenges and the real-world consequences of cyber threats will become even more pronounced. Dramatic improvements are

needed, but to date there have been few ideas powerful enough to effect fundamental change. However, promising new foundations for computing have emerged under the general umbrella of software-defined infrastructure (SDI).

SDI comprises a range of technologies including: processor, storage, and network virtualization; novel separation of concerns at the systems level (e.g. software defined networking – SDN – which separates data and control planes); and new approaches to system and device management. It has become increasingly clear that SDI offers a realistic opportunity to revisit the end-to-end foundations of computing security by enabling powerful new techniques to minimize, prevent, detect, and respond to threats and intrusions.

The SDI-CSCS project operates on the assumption that certain SDI fundamentals are mature enough to begin exploring the broader impacts for cybersecurity. For example, at the network level, applications of SDI technologies include new ways to deflect and mitigate denial of service attacks. In particular, isolation of processes and subnets has been shown to be effective in contexts such as cloud computing and data-centers. At the compute level, hardware-based security capabilities can potentially be combined with virtualization to strengthen the "root of trust" within individual end systems.

The extension of architectural elements (such as hypervisors) with capabilities beyond basic virtualization (such as introspection) may provide additional contextual information that can be used as an input to monitoring/control systems to enhance security.

Looking to the future, SDI may broadly enable new capabilities such as alignment and enforcement of security policies associated with different types of virtual infrastructure, threat isolation, securing loosely-coupled micro services, and enhanced techniques for isolating, analyzing, and responding to various types of threats. At the same time, SDI enables mobility, distribution, and disaster recovery, all of which may play a role in evading, mitigating, and recovering from attacks.

3 Areas of SDI-CSCS Research Activity

In the context of an SDI framework, examples of research directions include but are not limited to:

- SDI-enabled least-privilege execution, such as by controlling access to virtual resource pools to serve the needs of a particular application and/or user (role);
- SDI as a foundation for improved visibility into the normal and abnormal behavior of highly distributed applications, and the use of such insights to determine when application behavior violates security policies;
- SDI-enabled resilience (e.g., moving target defense, adaptive response, and flexible mitigation); and
- SDI as the trust measurement interface between software and the underlying hosting platforms (processing, network, storage), e.g., mutual attestation between a virtual platform and software.

Several proposals already funded with work underway are discussed briefly below.

3.1 EP3: A Clean-Slate Software-defined Approach for Enabling Elastic Security

"A Clean-Slate Software-defined Approach for Enabling Elastic Security" is work performed by a team of faculty at Carnegie-Mellon University, designated as EP3.[2] The EP3 focus is on building a security control plane that is elastic in the processing available, its placement, and the policies it enforces; and exploring applications that exploit these new capabilities while solving necessary system performance issues.

The key research questions in EP3 are on the three related "EP" postulations each derived from virtualization's use in the security control plane: with Elastic Placement, topology does not limit security processing. With an Elastic Posture, resources can be leveraged to scale elastically based on attacks. With Elastic Performance, new security features can be deployed – and security policies defined and assigned – very quickly, in a context-aware manner.

3.2 S2OS: Enabling Infrastructure-wide Programmable Security with SDI

"Enabling Infrastructure-wide Programmable Security with SDI," or S2OS, is work performed by a collaborative team of Computer Science faculty from the University of Texas, Texas A&M, Colorado University, University of North Carolina, and Clemson University.[3]

The S2OS team has taken as its central challenge the design and construction of an SDI-defined operating system (OS) which would itself abstract security capabilities, provide fine-grained controllability, and allow security policy to be programmed throughout an enterprise infrastructure.

In this approach, a base infrastructure layer consisting commonly of virtual machines, containers, network devices, host machines, and mobile devices, would interface with a Security Microservice Management control layer, through a series of control agents (e.g. the hypervisor for the VMs, a container engine for the containers, Open vSwitch for the mobile devices, etc.). That control layer would also retain traditional enterprise security capabilities and expectations (isolation, trust management, risk management, monitoring, auditing/logging, etc.) – all accessed as SDI-managed services, allowing optimized provisioning and service flows. The control layer itself would also

---

[2] The Carnegie-Mellon University team included Srini Seshan, David Anderson, Bryan Parno, Vyas Sekar, and Justine Sherry.

[3] The collaborative S2OS team included Guofei Gu of Texas A&M, Zhiqiang Lin of UT Dallas, Donald Porter of the University of North Carolina, Eric Keller of Colorado University, and Hongxin Hu of Clemson University.

interact via an API with the topmost application layer, where SDI-wide "programmable security applications" would feature fine-grained access control, app hardening, moving-target defenses, etc., each enabled by the overall software-defined infrastructure.

SDI-CSCS researchers have characterized the known and potential future threat model(s) addressed, and attempted analytic characterizations at any changes in the overall threat surface that introduce different vulnerabilities, and therefore require subsequent attention. Experimental research that involves the creation, deployment, and evaluation of prototype cybersecurity systems has been encouraged, especially those for which system behaviors at scale can be extrapolated, as will be discussed in the full paper.

Based on these projects' findings, we envision that the resulting principles could be applied in many settings, including enterprise networking, data center networking, IoT, and/or SDN and Network Function Virtualization (NFV)-based next-generation ("5G") telecommunications. In addition, distributed scenarios of interest include, but are not limited to, client-cloud workloads, embedded-gateway-cloud (MGC) workloads, and loosely-coupled micro services. Secure content and process mobility are also of interest in settings ranging from edge/cloudlet computing to data center disaster recovery and on to content or process distribution networks.

4 Machine Learning, Data Analytics and Artificial Intelligence in SDI Environments

Across software-defined environments, VMware Research is actively prototyping cybersecurity approaches to virtualize more powerful machine-learning (ML), to make those virtualized environments more efficient, scalable, resilient and cybersecure.

VMware has a unique position in the low-level hypervisor view of the infrastructure to have visibility, insight and control of virtual servers, networks, storage, containers, and even applications. One example of our effort to exploit that unique view is our research project *PIDForest*, taking advantage of the low-latency of SDI virtualized environments to attack the cybersecurity challenge of detecting anomalies in a large dataset, for example of detected perimeter attacks. The project introduces a geometric anomaly measure called *PIDScore*, and algorithmically detects anomalies based on this definition – using the intuition that anomalies are easy to distinguish from the overwhelming majority of points by relatively few attribute values: we call this partial identification. PIDForest already performs favorably in comparison to several popular anomaly-detection methods, across a broad range of benchmarks.

Similarly VMware is investing in SDI-enabled cybersecurity capabilities that make use of the masses of data within the virtualized architecture, to bring insight through the chaos of cloud native and serverless applications, as well as the never-ending growth and diversity of mobile devices and endpoints. We have recently developed and deployed technology called AppDefense, which leverages our low-level position in the hypervisor to learn the history, intended behavior and observed behavior of applications – the "known good", and subsequently to lock that environment down. This capability

will then learn as the system runs and evolves through such things as patches and other changes.

Based on our research and findings from scaled deployments, this *Learn Lock Adapt* concept will, in theory, be more effective at identify abnormal or abhorrent behavior, because it has a deep understanding and insight of the known good behavior. It will then use the power of automation, network segmentation and low-level control to address and limit threats with speed and efficacy.

5 Virtualization and SDI for Security in Edge Computing Architectures

To address the related importance of cybersecurity in Internet of Things environments, VMware Research is also collaborating with the National Science Foundation to advance the state-of-the-art in distributed virtualized systems, with joint research on *Edge Computing Data Infrastructure* for end-to-end networked systems architecture securing edge nodes. In 2018 VMware awarded funding for two awards valued at a total of $6 million for university faculty members' research in partnership with NSF.[4]

The exponential growth in the number of networked devices and sensors, whether in industrial, commercial, or military/battlefield environments, is expected to cause an explosion in the amount of data generated at the network edge. Further, there are new latency-sensitive interactive applications that would benefit from fast access to this data; military examples include but are not limited to command-and-control and fire-control applications. Edge computing, where data is processed close to the data-generating endpoints, has been a natural evolution and already demonstrates saving on bandwidth costs and reducing latency. However these distributed architectures create a complex multi-stakeholder computing-services ecosystem; imagine allied military scenarios across land, sea, air, space, and cyber domains. Secure data-sharing with data privacy across the different players in this ecosystem will be required.

Therefore the central challenge being addressed in VMware Research efforts to extend SDI to Edge Computing Data Infrastructure is to design and develop data-centric edge architectures, programming paradigms, runtime environments, and data sharing frameworks, in conjunction with concurrent advances in VMware's existing virtualized "hyper-converged infrastructure" (HCI) software, to enable compelling new applications and fully realize the opportunity of big data in tomorrow's mobile and IoT device environments – while simultaneously addressing the implications of edge computing's multi-stakeholder context, and the need for security and privacy as first-order design and operational considerations.

---

[4] See "NSF/VMware Partnership on Edge Computing Data Infrastructure (ECDI)," Program Solicitation for NSF 18-540, National Science Foundation Directorate for Computer & Information Science & Engineering, Division of Computer and Network Systems; https://nsf.gov/pubs/2018/nsf18540/nsf18540.htm.

6 Virtualization and SDI for Security across Container Deployments and Zero Trust Architecture

Our security research and development continues to explore new areas for SDI's extension. The recent growth and popularity of Linux containers has offered particular challenges from a security standpoint. A major disadvantage of container-based virtualization compared to traditional virtual machines has been that containers share the OS kernel and other components of the host operating system, and have root access. The widespread propagation of the open-source Kubernetes platform as an automation-engine for container operations has not addressed security in an infrastructure or enterprise manner, instead only extending the fragmentation further across the enterprise.

Therefore Kubernetes clusters typically inhabit (or "sprawl across") separate security domains, with security policies applied individually per cluster, as well as upgrades, monitoring, and backups. Role-based access control (RBAC) is typically managed separately by cluster, and configuration/assurance/compliance management is impossible in a reliable enterprise way.

VMware has developed an approach to apply the fundamentals of software-defined infrastructure to cluster/container resource management enabling security benefits. Described as "Project Tanzu," the extension of an SDI control plane to Kubernetes management is intended to unify all of the various tasks needed to manage and maintain secure clusters separately – by applying policies around security, compliance, and configuration management across clusters. Thus enterprise identity can be mapped to Kubernetes RBAC across clusters; configuration and assurance policies can be defined once and pushed/manage across clusters via an SDI-enabled "Tanzu Mission Control Plane." Traffic can be monitored across clusters to ensure automated policy enforcement.

Zero-Trust Architecture: SDI promises to enable not only the convergence of traditional VMs and containers for configuration and security management, but an adaptive "attack-protection architecture" with a default zero-trust posture, narrowing cybersecurity defenses from expansive network perimeters to SDI-defined and managed "micro-perimeters" around individual or clustered resources. In a zero-trust approach, there is no implicit trust in system resources, and user/device authentication is required for granular access controls. Attack surfaces are reduced through microsegmentation; service-specific firewalls provide intrinsic application visibility and control; and security policy can be designed and applied consistently across hybrid and multi-cloud environments, protecting both on-premise and cloud-native workloads.

The microsegmentation allowed by software-defined infrastructure is ideal for zero-trust approaches, and will allow more precise and rapid cycles between cyber attack prediction, prevention, detection, and response.

7 Future Research Areas

In this paper, we describe the SDI-CSCS program designed to apply the power of virtualization to real progress in cybersecurity, in orchestration with a number of security-relevant efforts within the VMware R&D organization to ensure that the infrastructures of tomorrow will be more secure. Examples of areas of continuing additional SDI research include:

- Orchestrations of hardware-specific amplifications of software-defined infrastructure. There are emergent possibilities in the arena of virtualized "Trusted Execution Environments" for certain advanced processors, strengthening the root of trust.
- Cross-cloud and multi-cloud "Trusted Computing," extending the two-decade-old trusted computing approach pioneered in the Windows environment, and extending SDI trust across cloud architectures, establishing trust in multiple control-plane components.
- Future 5G architectures will necessarily utilize network-level virtualization (SDN – Software Defined Network) and services (NFV - Network Functions Virtualization), thus we are researching advanced security designs based on SDI for 5G capabilities.
- Exploration of SDI's relevance in scaled blockchain deployments. VMware's *Project Concord* features a core engine "concord-bft," a generic state machine replication library that can handle malicious (byzantine) replicas. This library is designed to be used as a core building block for replicated distributed data stores and is especially suited to serve as the basis of permissioned Blockchain systems. This SDI-influenced vision, optimized for decentralization, would enable high-throughput data replication, seamless core cross-ledger transaction support, and a secure world of many secure ledgers, not a single global database.[5]

There is much work to be done by our industry to ensure cybersecurity for critical infrastructures, and we anticipate further progress from our collaborations in extending software-defined infrastructure.

---

[5] Project Concord's implementation is based on the algorithm described in the paper "SBFT: a Scalable Decentralized Trust Infrastructure for Blockchains," co-authored by Guy Golan Gueta (VMware Research), Ittai Abraham (VMware Research), Shelly Grossman (TAU), Dahlia Malkhi (VMware Research), Benny Pinkas (BIU), Michael K. Reiter (UNC-Chapel Hill), Dragos-Adrian Seredinschi (EPFL), Orr Tamir (TAU), and Alin Tomescu (MIT), April 2018; accessed at https://research.vmware.com/files/attachments/0/0/0/0/0/5/4/sbft-arxiv2018.pdf.

References

1. National Science Foundation, VMware Research: “NSF/VMware Partnership on Software Defined Infrastructure as a Foundation for Clean-Slate Computing Security (SDI-CSCS),” Program Solicitation for National Science Foundation 16-582, NSF Directorate for Computer and Information Science and Engineering, Division of Computer and Network Systems (2018). Accessed at https://www.nsf.gov/pubs/2016/nsf16582/nsf16582.htm
2. National Science Foundation, VMware Research: “NSF/VMware Partnership on Edge Computing Data Infrastructure (ECDI),” Program Solicitation for NSF 18-540, National Science Foundation Directorate for Computer & Information Science & Engineering, Division of Computer and Network Systems. NSF (2018). Accessed at https://nsf.gov/pubs/2018/nsf18540/nsf18540.htm.
3. Gueta, G., Abraham, I., Grossman, S., Malhi, D., Pinkas, B, Reiter, M., Seredinschi, D-A., Tamir, O., Tomescu, A.: SBFT: A Scalable Decentralized Trust Infrastructure for Blockchains. Arxiv (2018).
4. VMware Research Homepage, http://research.vmware.com/, last accessed 2019/10/03.