February 18, 2011

***Comments of EMC Corporation on*:**
The Preliminary FTC Staff Report on Protecting Consumer Privacy in an Era of Rapid Change: A Proposed Framework for Businesses and Policymakers (December 2010)

EMC commends the FTC staff for its thoughtful draft Report on protecting consumer privacy in an era of rapid technological change. One of the areas on which the FTC has sought comment focuses on how stakeholders can develop and deploy privacy-enhancing technologies to address and minimize privacy concerns. (See Draft Report at Appendix A-2). Industry-led technological developments can and should be a key driver of sensible and consumer-friendly solutions in this important area.

These comments provide a brief overview of one such technological development -- "virtualization." "Virtualization" is a flexible and extensible technology that allows privacy and data security controls to be implemented directly into a computing environment, thereby making *privacy by design* a reality and helping to achieve the FTC's goal of minimizing privacy concerns even as cloud computing facilitates and encourages the transfer, processing, storage, and maintenance of personal data beyond a company's own infrastructure.

What is "Virtualization"?

Virtualization, in the world of computing, means the creation of a virtual (as opposed to a physical) version of something, such as a hardware platform, operating system, or network resource. More specifically, virtualization is a method of resource tracking and allocation that permits multiple operating systems to run concurrently on a single physical computer and share hardware resources with each other. For example, virtualization allows an Apple computer that is running a Mac operating system to host a "virtual machine" (or "VM") that looks and operates exactly like Micrososft Windows, and allows Windows applications to run simultaneously on that same physical computer alongside Mac applications. Likewise, with "server virtualization," the same physical computer can be used to act as the server running a company's email/calendaring system as well as its data storage/backup system.

In short, virtualization represents an important paradigm shift from a *device-centric* to a *data-centric* computing environment.

Virtualization works by inserting a thin layer of software directly on the computer hardware or on a host operating system. This software contains a virtual machine monitor or "hypervisor" that allocates hardware resources dynamically and transparently. For example, if applications running under one virtual machine or one virtual server experience a spike in demand, the physical computer is able to temporarily borrow and allocate resources from the other virtual machine or server to supply the necessary resources in an efficient and adequate manner. The overall effect of this resource tracking and allocation across multiple independent and simultaneously running virtual computing environments is far greater efficiency and flexibility than with the traditional single operating system paradigm. In particular, even though some basic multi-tasking is possible in the traditional computing environment, most machines' resources are vastly underutilized. In contrast, by allowing a user (or users) to run multiple virtual machines with varying operating systems and applications on a single physical machine, with each virtual machine sharing all the resources of that one physical computer across multiple environments, virtualization maximizes the use and efficiency of a physical computing

platform's resources, thereby requiring fewer machines to be purchased, maintained, and managed.[1]

Given the efficiencies that virtualization is able to achieve, it has been a significant enabler of cloud computing, as cloud providers can implement this technology to efficiently and securely track and allocate computing and network resources to various customers.

Virtualization and Data Protection

Virtual machines offer three significant advantages over the "physical" machine paradigm for enhancing privacy and security: superior visibility, control, and flexibility. First, in virtual environments, it is possible to see all of the bits as they travel from one application to another; there are no data flows, network traffic, or related processes that cannot be evaluated by the virtualization infrastructure. By contrast, for example, traditional physical computing environments are vulnerable to "rootkit" infections that mask the presence of malware. The improved visibility of virtualization, which is sometimes referred to as "introspection," is unaffected by any of the masking used in rootkits, giving it a significant advantage over conventional security tools.

Introspection also gives rise to a second enhanced capability, that is, greater control over data-protection-enhancing features not found in the traditional physical computing environment (which has a more limited ability to evaluate and control data and data flows). Because of the greater visibility to data flows, features like anti-virus, access controls and authentication,[2] data loss prevention,[3] intrusion detection, encryption, and other data protection controls can be built directly into the virtual environment (*i.e.*, designing data protection directly into the virtual layer) in a transparent manner without reconfiguring the underlying physical systems. Data deduplication is another technique found in virtualized environments which, although aimed primarily at reducing data storage requirements, also may provide security benefits.[4]

Lastly, unlike traditional physical environments, which rely on static security configurations to maintain privacy and data security, virtual environments have built-in flexibility that facilitates

---

[1] Real-world benefits of virtualization also include migrating legacy operating systems and software applications to virtual machines.

[2] Data security should not only make data harder to steal, but also make it harder to use in the event data is compromised. Data-centric security makes data harder to use in the event it is stolen, which is particularly important from an access and authentication standpoint.

[3] Data loss prevention (also referred to as "data leak prevention" or "DLP") enables organizations to discover where personal data is located within an organization's environment and monitor how it is used within such environment. DLP can also enforce security controls in the event data is being used in a manner that violates data protection policies (*e.g.*, block a transfer of sensitive personal data to a USB stick or similar portable data storage device).

[4] Deduplication is a specialized data compression technique used to eliminate coarse-grained redundant or duplicate data (*i.e.*, the same data that keeps getting stored over and over again) within an organization's existing environment. Deduplication looks at data at a very granular level, identifies repeating patterns, and ensures that such patterns are not stored again, thereby storing only one copy of the unique data.

and maintains dynamic data protection as the environment changes. Privacy and data security controls seamlessly follow virtual machines, allowing users to leverage key virtualization features like load balancing and disaster recovery while being assured that data protection policies are "Always On" and will follow and control virtual machines.[5]

By associating data protection controls with ***logical*** boundaries in a dynamic, virtualized environment, it is possible to make data in this environment more secure than data stored in a traditional physical environment.

We encourage the Federal Trade Commission to continue to explore how virtualization can be used to enhance data protection even in an increasingly cloud-focused environment where rapid technological development is the norm.

We hope that these comments are helpful. Please contact me if you have any questions.

Sincerely,

Demetrios Eleftheriou
Senior Counsel - Privacy
EMC Corporation
176 South Street
Hopkinton, MA 01748
508-293-6327 (direct)
508-497-6915 (fax)
demetrios.eleftheriou@emc.com

[5] Virtualization allows the user to recover to any machine, not just specific duplicate hardware, thereby reducing hardware and maintenance costs and the complexity of maintaining a backup site.