

## Evaluating Thin-Client Security in a Changing Threat Landscape

Equivalent security controls can be, and are, implemented on PCs—without giving up the functionality that is sacrificed with the thin-client model.

### Executive Overview

**Intel IT's security team continually analyzes our computing model to determine how it needs to evolve in response to an ever-changing threat landscape. Recent cyber-attacks on a number of high-profile targets provided added impetus to re-evaluate the security offered by thin-client models and whether using thin clients could help defend against similar attacks.**

We identified five attributes that are often perceived as security benefits in thin clients: prevention of physical data loss, removal of administrative privileges, limitations on installed applications, client integrity, and ability to roll back to a known good state.

We determined that while these controls can contribute to a more secure environment, they would not have prevented the recent cyber-attacks from being successful.

Furthermore, we also observed that these controls are not unique to thin clients: Equivalent controls can be, and are, implemented on PCs—without giving up the functionality that is sacrificed with the thin-client model. Where such controls have not been implemented holistically on either thin clients or PCs, the reason is often because they place unacceptable restrictions on user productivity, not because of the client architecture.

We also considered other restrictions and costs of thin clients, including the inability to support mobile computing, highly interactive or compute-intensive applications, and rich media such as video. Thin clients also require significant additional server capacity and network bandwidth. Some thin-client models can also increase the risk of business disruption if an outage occurs within the central network resources upon which the thin client depends.

Based on our analysis, we see thin clients as suitable for some niche uses. However, Intel's environment is 80 percent mobile, and most of Intel's users require the functionality and flexibility of mobile business PCs. Mobile business PCs also position us to take advantage of emerging technology trends and service delivery models.

**Toby Kohlenberg**

Senior Information Security Specialist, Intel IT

**Omer Ben-Shalom**

Principal Engineer, Intel IT

**John Dunlop**

Enterprise Architect, Intel IT

**Jerzy Rub**

Information Risk and Security Manager, Intel IT

## Contents

Executive Overview.....	1
Background.....	2
Client Security Analysis.....	2
Security Controls Typically Associated with Thin Clients.....	2
Mitigation Value of These Controls.....	3
Availability of Security Controls on PCs.....	4
Thin-Client Security Concerns.....	4
Security In An Evolving Threat Landscape.....	4
Meeting Intel's Enterprise Needs ....	5
Future Positioning.....	6
Dynamic Virtual Client.....	6
Conclusion.....	7
For More Information.....	7
Acronyms.....	7

## IT@INTEL

IT@Intel is a resource that enables IT professionals, managers, and executives to engage with peers in the Intel IT organization—and with thousands of other industry IT leaders—so you can gain insights into the tools, methods, strategies, and best practices that are proving most successful in addressing today's tough IT challenges. Visit us today at [www.intel.com/IT](http://www.intel.com/IT) or contact your local Intel representative if you'd like to learn more.

## BACKGROUND

**Intel IT supports a very large enterprise environment, with about 83,500 employees spread across 61 countries. Intel depends on its employees for business innovation that enables the company to grow and continue to create a competitive advantage. To foster this innovation, Intel IT provides about 80 percent of employees with mobile business PCs.**

The cyber threat landscape has been rapidly evolving, with an increasing shift towards zero-day attacks, which target previously unknown vulnerabilities within hours of discovery. Some recent attacks on other high-profile Web sites have exploited previously unknown vulnerabilities in common client software such as Web browsers.

Intel IT's security team continuously monitors this threat landscape and regularly analyzes how our business and compute models match up against it. Recent cyber-attacks on some high-profile targets caused us to re-evaluate the security offered by the thin-client model. We also wanted to investigate whether thin clients could help defend against similar attacks in the future.

After analyzing these thin-client attributes, we examined whether similar controls can be applied to PCs. Then we considered our findings from the perspective of the overall Intel IT client strategy, taking into account enterprise user needs and emerging technology trends.

## CLIENT SECURITY ANALYSIS

**We analyzed security controls commonly associated with thin clients along with their value in the evolving threat landscape. We assessed whether they would have prevented or mitigated targeted zero-day attacks such as the recent exploits at other high-profile Web sites. We then analyzed whether similar controls can be applied to PCs.**

### Security Controls Typically Associated with Thin Clients

We identified five main attributes that are commonly perceived as security advantages in thin clients:

- **Physical data loss prevention.** With thin-client models, storage is restricted to the data center, reducing the risk of physical data leaks. In addition, thin clients often lack ports for attaching external media or USB memory sticks, further reducing the ability to copy data directly from the client.
- **Non-privileged users.** User administrative privileges are removed, reducing the ability of exploits to change system files and settings.
- **Restrictions on user-installed applications.** Users are not able to install additional applications that would enlarge the attack surface of the client machine or infect the system with malware.

- **Client integrity.** All clients are maintained in a consistent state, based on a known configuration baseline. New patches can be rapidly and consistently applied by patching server-based images.
- **Ability to roll back to a known good state.** With thin clients, this can often be achieved by rebooting or reloading previous versions of a single virtual container file.

## Mitigation Value of These Controls

We found that while these controls can contribute to a more secure environment, they would not have provided significant overall protection in the recent zero-day attacks.

- **Centralized data storage.** Traditionally, data theft involved using a device to copy data physically stored on the system. However, today's thefts typically take place over networks, as shown in Figure 1. The restrictions imposed by thin clients do nothing to prevent this. All thin clients have fast network connections and most have Internet connections. Attackers can use these fast networks to rapidly transmit data from the server through the firewall.
- **Non-privileged users.** Removing users' administrative rights may reduce the impact of an infection and make it more difficult for malware to spread to new systems. However, it cannot always prevent the initial compromise, and unless extreme restrictions are imposed, attackers may still be able to avail themselves

of users' remaining privileges to gain access to other systems and any data to which the user has access. Furthermore, removing users' administrative rights has no protective effect if the exploited service is running as a system process rather than as a user process.

- **Restrictions on user-installed applications.** These restrictions can help if they prevent installation of non-essential applications that could be targeted. However, recent exploits have targeted ubiquitous, essential applications such as Web browsers. Furthermore, it is much more difficult to restrict users' access to malicious Web sites or prevent them from running undesirable Web services than it is to prevent them from installing software on a business PC.

- **Client integrity.** Applying consistent, up-to-date patches is generally easier with a centralized image, as used in the thin-client model. However, it would not have prevented the attacks, because zero-day attacks target previously unknown vulnerabilities.
- **Ability to roll back to a known good state.** Rolling back the client system using a server-based client image would not have helped in recent attacks. The initial compromise provided access to Web-based services and accounts, and rolling back the system after the compromise would not have removed this access. In addition, since the attack exploited an unknown vulnerability, the system could have been just as easily recompromised.

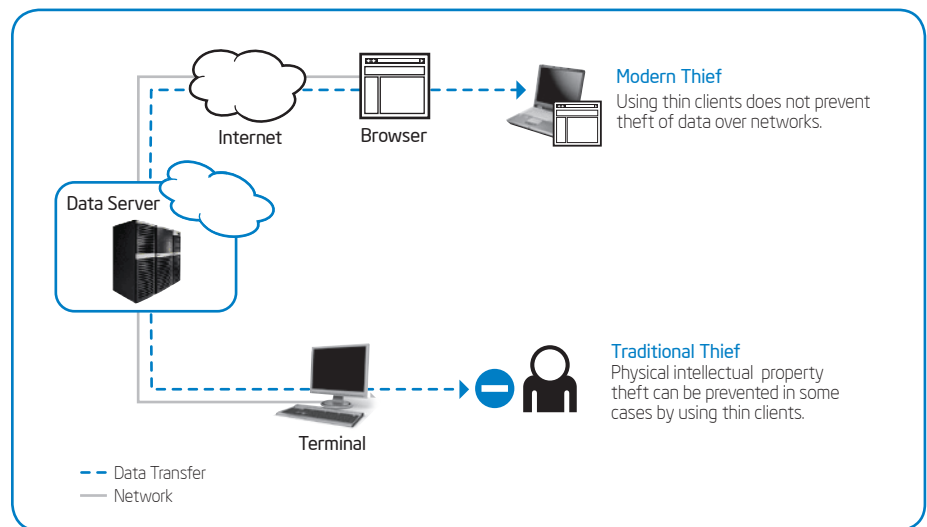


Figure 1. Thin clients can help prevent physical theft of data from clients; however they do not prevent theft over networks.

## Availability of Security Controls on PCs

We determined that we could, and indeed do, provide equivalent controls on PCs—without giving up the functionality lost with the thin-client model. We have not implemented these controls holistically, but we have used individual controls where appropriate.

To prevent physical data theft, Intel uses full disk encryption and enterprise rights management tools. These methods can be supplemented with global domain policies or physical system modifications to lock down, encrypt, or restrict the use of USB memory and other external storage devices. In addition, folder redirection can be used to store data in the data center, either exclusively or as a mirror of the data on the client. The latter approach supports mobile computing.

The decision about whether to assign administrative rights to users is not specific to thin-client models. Many tools are available to remove the administrative privileges of mobile business PC users. If requirements dictate that PC users receive administrative rights, third-party software packages can be used to manage these rights and to restrict users from installing applications or carrying out other activities unless each action is specifically permitted by IT administrators. Intel uses multiple methods to restrict administrative access and users' rights.

Centralized management of a common OS or application image is not unique to thin-client computing. For example, streaming allows centrally managed OS and application images to be shared by multiple desktop PCs, and anticipated future capabilities include centrally managed virtualized clients that are downloaded to users' PCs. In addition, technologies from several companies can be used to implement system rollback on

business PCs. Intel maintains standard, centrally managed system and application images. These are updated regularly and used whenever rollback is necessary.

Intel IT currently is able to quickly deploy patches to maintain client integrity, and the most time-consuming aspect of the process is testing and validation of a new patch, not actually deploying it; the time required to test and approve would not change in a thin-client environment. With PCs, patches can be rolled out in waves to increasing numbers of users to mitigate potential problems introduced with the patches; updating all clients at once can itself be a risk.

## Thin-Client Security Concerns

There are also security concerns related to the thin-client model. Centralizing applications and data also centralizes the threat: A network of thin clients provides many access points to servers storing shared data and applications, with the associated risk that compromise can affect the entire IT infrastructure. Some thin-client models can also increase the risk of business disruption if an outage occurs within the central network resources upon which the thin client depends.

Centralizing all data, including personal data, introduces new privacy concerns, with an increased risk of infringing government regulations in some countries. Thin-client models can also have unintended consequences for data leakage; for example, we have found that users of thin clients are more likely to print paper copies of information that can then be disclosed to unauthorized parties.

To respond to a security breach, security professionals need to know which systems were compromised and when the compromise first occurred. In fact, preservation of such

evidence is mandated in some cases. Rolling a thin-client system back to a known good build, by rebooting from a server-based image, may actually have the negative effect of destroying this vital evidence on the client.

---

## SECURITY IN AN EVOLVING THREAT LANDSCAPE

**As the threat landscape mutates toward targeted zero-day attacks, we need to adapt by taking a different approach to security. In this landscape, IT security professionals need to assume that clients are vulnerable—whether they are PCs or thin clients.**

Frequently, attackers use custom malware; this is unlikely to be detected or prevented using traditional methods. Attacks typically focus on ubiquitous components that exist on both PCs and thin clients, such as essential business applications, the OS, or parts of cloud-based services.

Detection requires more sophisticated behavioral analysis for user access and rights utilization, including a more balanced mix of detective and corrective controls.

Both types of controls may actually be easier to implement using new platform technologies—such as Intel® Virtualization Technology (Intel® VT-x), Intel® Virtualization Technology for Directed I/O (Intel® VT-d), and Intel® Trusted Execution Technology (Intel® TXT)—or separate physical hardware rather than a shared virtualized server-based environment, which is what the thin-client model essentially uses.

For example, an attack originating within one virtual machine (VM) running on a server may target another VM on that same server in what is known as a VM escape.

Reliable preventative and detective controls, analogous to network or host-based intrusion prevention systems (NIPS or HIPS), do not yet exist in the VM management layer to help protect against this sort of attack. By distributing the execution of the VMs to separate client devices, or using technologies on the platform to provide the hardware isolation, the risk of this sort of attack is greatly diminished.

## MEETING INTEL'S ENTERPRISE NEEDS

**After determining that thin clients would not have prevented the recent zero-day industry attacks—and that similar security controls can be implemented with PCs—we analyzed which clients best fit Intel's current and**

**future enterprise needs. This analysis was based on business requirements and technology trends as well as security concerns.**

We found that mobile business PCs support the widest range of uses and therefore meet the requirements of most Intel employees.

Locally installed applications, combined with local processing power, provide users with increased flexibility to work anywhere, including locations without network access. This also means users retain some computing capabilities in the event of a disaster. Mobile PCs support compute-intensive applications and rich media for communication, including video and Voice over IP (VoIP), collaboration, training, and research. Users can run a wider choice of software applications, including business applications as well as

some personal applications (Intel, like many other companies, permits reasonable use of corporate resources in this way).

Another important advantage of mobile business PCs is that they support all service delivery models. Our strategy includes a growing number of services delivered from internal and external clouds, and we are exploring new delivery models such as application streaming. Equipping users with mobile business PCs means they can run any mix of these models while continuing to use conventional locally installed applications, as shown in Figure 2.

In contrast, thin clients have significant limitations. Often, there is limited support for mobile computing or working offline, and users cannot effectively run bandwidth-, graphics-, or compute-intensive applications.

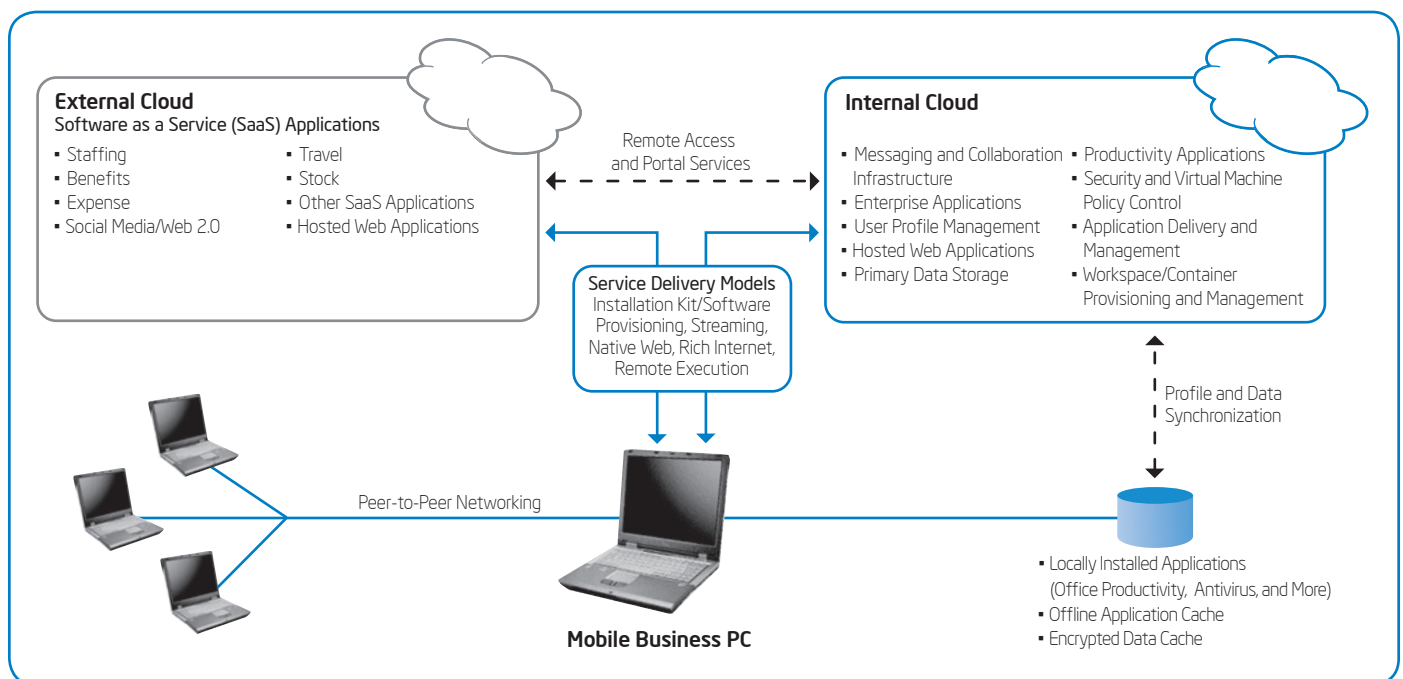


Figure 2. Mobile business PCs can support all emerging service delivery models while continuing to run conventional locally installed applications

Our analysis has also shown that thin clients require significant additional server capacity, as well as increased bandwidth across the network supporting the connected users.

Because of these limitations, we have determined that thin clients may be suitable for only specialized use cases, such as call center terminals, shared kiosks, or manufacturing controllers. However, for these use cases, our analysis suggests that OS streaming to PCs may be a better fit for us than using thin clients. This is because PCs execute their workloads locally, providing greater responsiveness and a better user experience; also, less server and network capacity is required.

FUTURE POSITIONING

Looking forward, our client strategy must continue to keep pace with security requirements, enhance employee productivity, and reduce IT costs. At the same time, we need to manage added complexities such as an increasing number of devices per user, a diversification of form factors, and the adoption of consumer technologies within the enterprise.

We plan to take advantage of new technologies, computing models, and service delivery methods to reconcile these seemingly conflicting requirements. We are planning a segmented approach that includes a range of

client solutions that fit the needs of different enterprise users.

Dynamic virtual client (DVC) is one option we are considering. See “For More Information” at the end of this paper for discussions of other solutions.

Dynamic Virtual Client

DVC—a virtualized PC environment delivered on demand to clients running native (Type 1) high-security hypervisors—is a key element of our client strategy. This solution is shown in Figure 3.

We believe DVC will deliver several advantages, including device-independent mobility. Users should be able to download and run their virtual containers on a variety of client hardware, with better separation of business and personal workspaces on the same system.

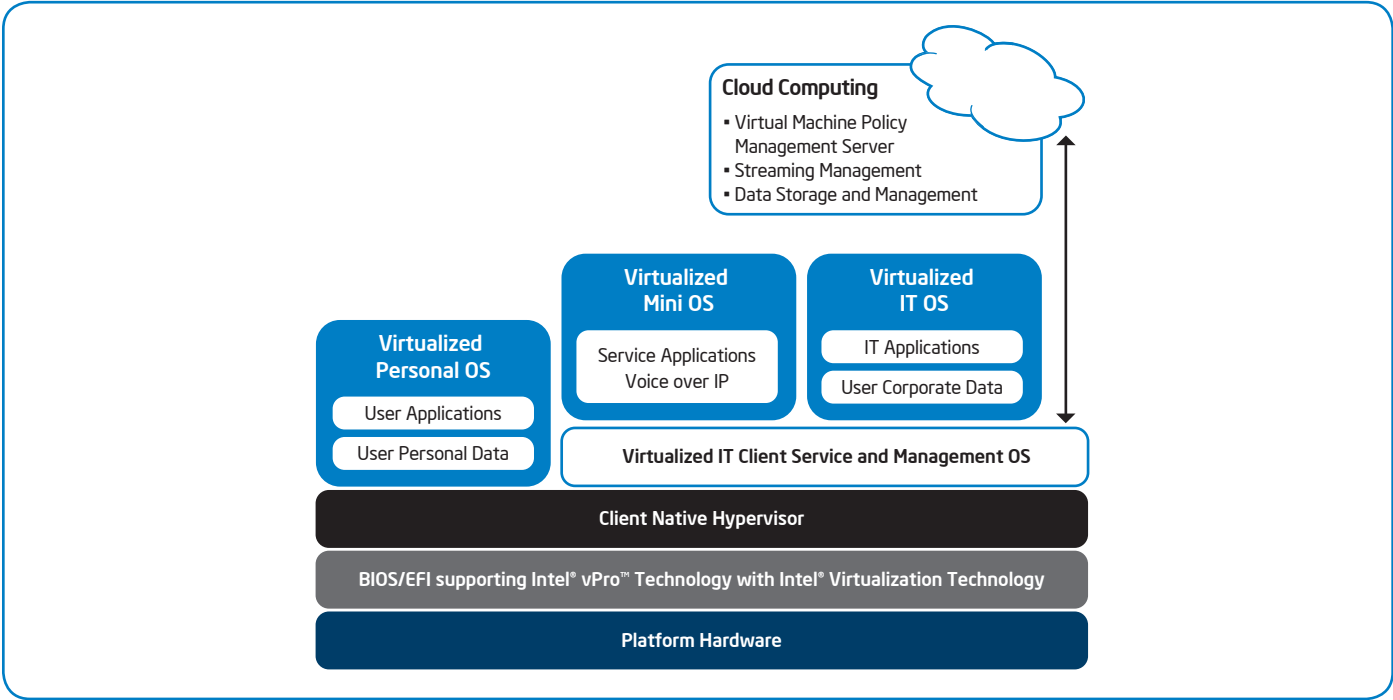


Figure 3. Intel IT is investigating a dynamic virtual client (DVC) solution.

For specialized niche uses where it is feasible to implement all the required controls, we are considering OS streaming to provide high-availability, fully managed, and tightly controlled desktop PCs. Streaming to PCs allows us to reap the benefits of centralized management, with server-based OS images shared among many desktops, without the performance sacrifices associated with thin clients. With PCs, the workload executes locally, providing better performance; also, less server and network capacity is required.

## CONCLUSION

**We determined that while the controls often associated with thin clients can contribute to a more secure environment, they would not have provided protection against recent zero-day attacks. In addition, these controls are not unique to thin clients: They can be, and are, implemented on PCs using other methods, without giving up the functionality that would be sacrificed with the thin-client model.**

Taking into account functionality and performance as well as enterprise security, we have found that mobile business PCs support the widest range of uses and therefore meet the requirements of most Intel employees. Because mobile business PCs support all emerging service delivery and computing models, they also position us for the future. The limitations of thin clients make them suitable only for specialized niche uses in our environment.

## FOR MORE INFORMATION

Find additional IT@Intel white papers at [www.intel.com/IT](http://www.intel.com/IT).

- "Enabling Device-Independent Mobility with Dynamic Virtual Clients"
- "Better Together: Rich Client PCs and Cloud Computing"
- "Developing an Enterprise Client Virtualization Strategy"
- "Improving Manageability with OS Streaming in Training Rooms"

## ACRONYMS

DMA	direct memory access
DVC	dynamic virtual client
HIPS	host-based intrusion prevention system
NIPS	network intrusion prevention system
SaaS	software as a service
Intel® TXT	Intel® Trusted Execution Technology
Intel® VT-d	Intel® Virtualization Technology for Directed I/O
Intel® VT-x	Intel® Virtualization Technology
VM	virtual machine
VoIP	Voice over IP

**For more straight talk on current topics from Intel's IT leaders, visit [www.intel.com/it](http://www.intel.com/it).**

---


This paper is for informational purposes only. THIS DOCUMENT IS PROVIDED "AS IS" WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, NONINFRINGEMENT, FITNESS FOR ANY PARTICULAR PURPOSE, OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, SPECIFICATION OR SAMPLE. Intel disclaims all liability, including liability for infringement of any proprietary rights, relating to use of information in this specification. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted herein.

Intel, the Intel logo, and Intel vPro are trademarks of Intel Corporation in the U.S. and other countries.

\* Other names and brands may be claimed as the property of others.

Copyright © 2010 Intel Corporation. All rights reserved.

Printed in USA  
0410/JLG/KC/PDF

 Please Recycle  
322970-001US

