DataDivider Protect Zone Security Controls

Introduction

DataDivider provides the ability for Merchants taking voice payments to have their agents capture cardholder data without exposing this data to their local desktop or network. This allows Merchants to remove their agent desktops, networks and backend systems from PCI DSS scope. DataDivider is able to capture cardholder data without exposing this data to the agent desktop by providing access to its Virtual Keypad and protecting this access within a Protected Zone which runs on the agent desktop.

This paper discusses how the Virtual Keypad works and how it is not possible to bypass the Protected Zone in order to get access to cardholder data on the agent desktop or compromise the agent desktop to otherwise harvest cardholder data. Removing the agent desktop and network from PCI scope vastly reduces the cost of PCI compliance and the risks of any breach on phone payments. These savings can be over $1,000 per agent machine per annum. DataDivider is a Level 1 PCI Certified Service Provider.

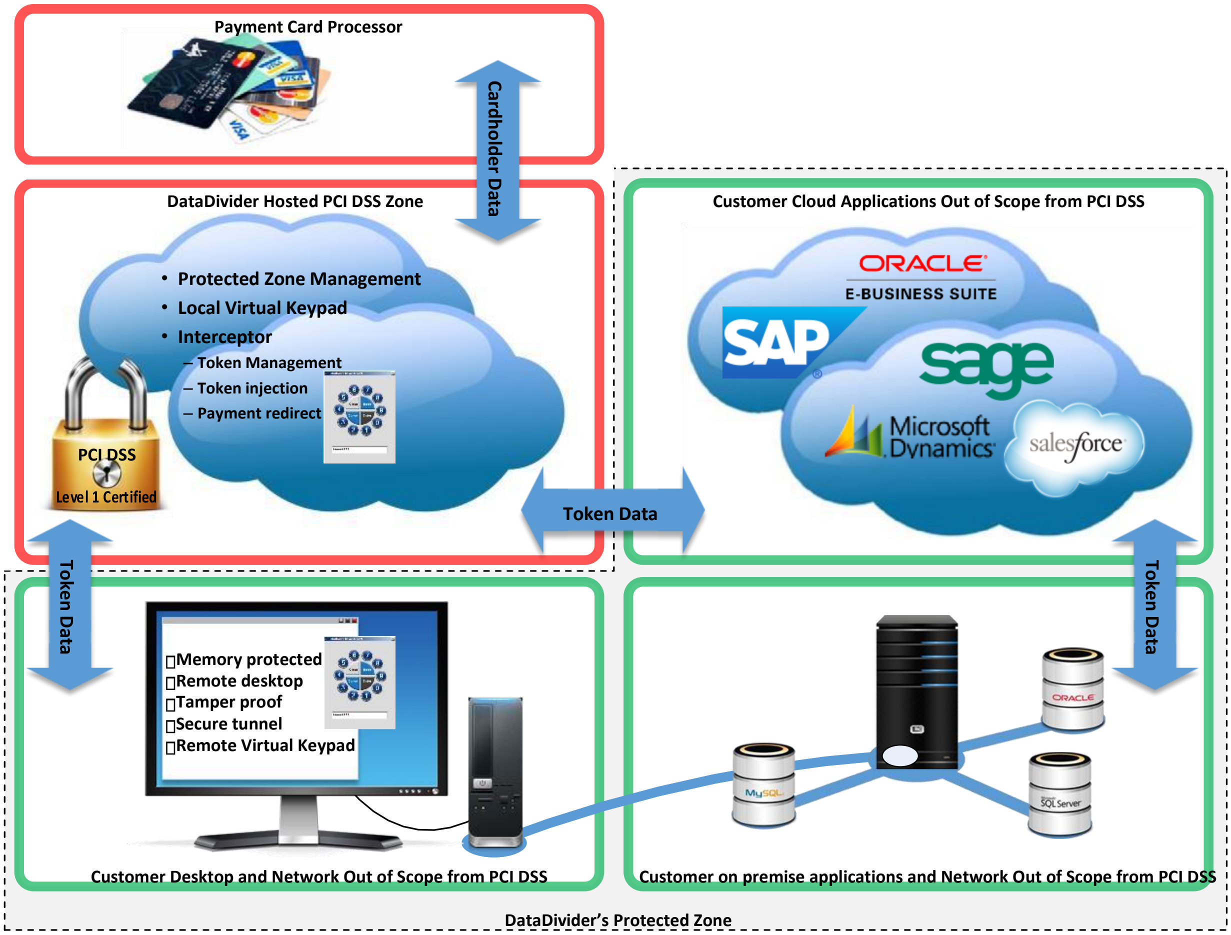

Figure 1 - DataDivider Voice Payments **DataDivider Virtual Keypad**

The Virtual Keypad runs in DataDivider's PCI Level 1 Certified hosted environment. Access to the Virtual Keypad is through a Remote Desktop Protocol (RDP) session run under the protection and control of the DataDivider Protected Zone. The Virtual Keypad presents the agent with a circular rotary keypad where the digit zero is initialized in a random position. The only form of data entry on the Virtual Keypad is via mouse clicks on the agent desktop. On entry of each cluster of digits (i.e. every four digits for MasterCard and Visa cards) the Virtual Keypad randomly rotates clockwise or counter-clockwise 1 to 5 digits.

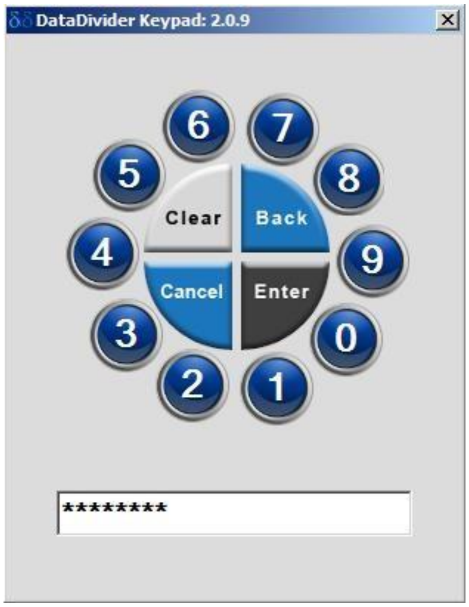

Figure 2 - DataDivider Virtual Keypad

DataDivider Protected Zone

It is not possible to correlate a mouse click coordinates to digits entered on the Virtual Keypad without the image of the Virtual Keypad at the time of the mouse click. It is therefore one of the roles of the Protected Zone to ensure it is not possible for any cyberattack on the agent desktop to get access to any visual images. Furthermore it is the role of the Protected Zone to guard against any attacks that could compromise the payment process and provide access to cardholder data. DataDivider works on the premise that the agent desktop has been compromised, however, through its protections and controls negates the ability of any malicious software gaining access to cardholder data or interfering with the genuine payment process. Following the deployment of the Virtual Keypad and the Protected Zone it will only be possible for malicious software to get access to mouse click coordinates which have no correlation to any entered digits. In effect this takes the agent desktop, its network and the Merchant's back-end systems out of scope from PCI DSS.

Establishing the Protected Zone

DataDivider establishes a secure encrypted session within layered tunnels on a merchant's desktop to ensure that its Virtual Keypad cannot be compromised. Multi-factor authentication validates credentials, location and legitimate profile. The session is allocated a protected area of local memory. During the session if this memory is inappropriately tampered with, the session will block tampering or terminate with informative security alarms.

A secure Remote Desktop Protocol (RDP) session is then established to the DataDivider PCI Level 1 Certified hosted environment. Within this session DataDivider's Virtual Keypad can operate such that it can capture sensitive cardholder data without exposing data to the local merchant environment.

The DataDivider Protected Zone takes a fundamentally different approach to securing and isolating the desktop. Instead of implementing the typical security controls such as AntiVirus, patch management, scanning, logging, personal firewalls and network segmentation, the DataDivider Protected Zone launches an on demand security layer around the secure session that connects to the DataDivider hosted PCI Zone. This layer is controlled by the multiple layers of encryption and a specific policy that prohibits man in the middle attacks and data leakage. It requires no extraordinary requirements of the end user or local system and does not require IT staff to install software, minimizing operational overhead and desktop support complexity.

The remainder of this white paper provides a technical review of the DataDivider Protected Zone's security features.

Secure Session

The DataDivider controls include a secure site to site login portal via whitelist that is protected by 2048 bit certified encryption. The portal provides authorized access through proxy to the secure browser policy, which then loads the DataDivider portal login and subsequent HTML5 RDP session; neither of which are public facing. Once authenticated to the Protected Zone via IP the client security policy executes on the client which consists of a series of DataDivider enforcers that provide protection of the secure session. These DataDivider agents are executables which use a variety of techniques such as process injection and hooking of system APIs to ensure the session security is maintained.

The DataDivider agent process launches a new, protected, instance of Internet Explorer (IE) using the IE version already installed on the machine. This instance and any subsequent child processes launched from within it have the DataDivider agent's API hooking, process control logic, antivirus and antispyware, and anti-tampering processes injected into it. This new instance of IE, running within a locked down, TLS encrypted tunnel, provides many of the controls securing the "Protected Zone" and its look-and-feel (borders, toolbar, logos, etc.) and operational behavior are controlled by the DataDivider agent logic as well as a policy file that is delivered as part of the launch process. The Protected Zone uses various techniques to harden the DataDivider agent against hostile code attempting to compromise the processes. These controls are extended to certain child processes launched from within the remote desktop session. All protected processes are denoted by a colored border.

Security Rules

The DataDivider browser and Protected Zone uses a form of sandboxing to isolate the session and supported applications from interactions with untrusted, potentially malicious software running on the agent's machine. The protected processes are closely guarded using numerous techniques to detect and prevent outside attempts to extract data from or alter the behavior of these protected processes. The behavior of this sandboxing can be adapted, via policy, to allow for certain trusted inter-process interactions such as returning tokens to local applications and ERP systems.

This sandboxing is accomplished using many techniques in concert. These techniques include:

- Protection of binaries on disk: The DataDivider agent is protected from an a priori alteration in two ways: 1) the executable and associated DLLs are signed, and their signatures validated at load time; 2) the cached binary package, if left behind via the caching option, is check summed to detect tampering attempts or a stale cached file, and either case would cause a new DataDivider agent binary to be downloaded.
- At start of execution: the DataDivider agent checks the signature of the binaries and all DLLs loaded by the process.
- Throughout the session, the DataDivider agent uses a variety of techniques to monitor attacks on the protected processes' space, including: ○ Monitoring all new thread creations and validating their DLL's signature, as well as starting location, to block code injection
    - Validation of signatures of all process images attempting to load into the protected process
    - Validation of Winsock LSPs ○ Detection of any attempt to debug ○ Blocking all hook libraries from loading ○ Blocking AppInit from loading any DLLs
    - Using a "honey pot" to detect and identify hooking modules ○ Monitor all Asynchronous Procedure Calls and validating that the APC function is contained in a valid loaded module
- Throughout the session the DataDivider agent proactively injects into all other processes, to not only catch key loggers and screen capturing software, but to also assist in detecting applications performing untrusted operations on the protected processes.

- In the event of attempts to disable or subvert the protected session, the DataDivider agent can be configured to take various actions, including allow, block and notify user, block silently, or exit. The events will generate logs and alerts to the DataDivider administrators.

Security Model

Unlike traditional antivirus products, The DataDivider Protected Zone does not use blacklists of known malicious software to make determinations on how to interact with other software running on an end user's machine. Instead, all external processes and software are considered untrusted by default, and are prevented from interacting with or extracting information from the secure browser and its supported child applications. If required, certain trusted third-party software can be added to a whitelist within the security policy to allow that software some level of access to the protected session. Other Protected Zone security mechanisms – SSL certificate whitelisting, hostname resolution, allowed network destinations – follow the positive security model, allowing administrators to leverage their knowledge of the operation of the protected applications to limit critical protected session actions to a site enumerated set of resources (e.g., to mitigate SSL proxy attacks, an administrator can define an allowed set of SSL certificates – defined as common name (CN) and thumbprint pairs – that can be used within the protected session). Although it is rare that access to applications are required from within the protected session these can always be accommodated without any compromise to any security measures.

Temporary and Process Specific Controls

Security mechanisms are in effect on an end point only basis while a session is active and most only apply to the session itself, leaving the rest of the desktop applications fully available to the agent. No permanent alterations are performed on system settings. This enables DataDivider to apply security controls on a targeted, temporary basis without altering the agent's device.

Malware Vs. Unauthorized Applications

An important concept in DataDivider's security approach is that the Protected Zone does not attempt to discern between malicious software and benign-but-untrusted software. Every instance of unrecognized, third party software is considered to be untrusted, and the Protected Zone makes equal, maximal efforts to prevent all untrusted software from accessing protected data or altering the behavior of protected applications. Only software that has been explicitly designated as being trusted is allowed to interact with the secure session. This enables DataDivider to integrate with local applications and ERP systems while preventing all leaking of Primary Account Number (PAN) and other sensitive data.

Protected Zone Security Features

Keylogger/Frame Grabber Defense

The Protected Zone provides zero-hour protection of keyboard input and displayed data continuously throughout the session, thus providing protections of authentication credentials, as well as other session keystrokes and rendered display data. It uses a patented method of real-time behavioral and heuristic detection of suspicious processes. Any suspicious processes, whether malware or legitimate applications, are blocked from accessing the hardened browser; no other processes are affected. The DataDivider security policy can include a “whitelist” of permitted key logger and frame grabber processes. The whitelist utilizes three different methods for enumeration: executable name, checksum or vendor signature. One or more of the methods can be used. Processes not white listed are blocked automatically. Through the DataDivider security policy the administrator has multiple options as to how to react to (e.g., user notification, event logging, session termination, etc.) in the event of detection.

Furthermore, all PAN and CVC2 fields are protected from keyboard access. Entry into these fields will only initiate the Virtual Keypad which again prohibits keyboard access and only allows data entry through mouse clicks.

Key Loggers: Prior to starting the protected browser session, and then continuously throughout the session, The DataDivider secure session executes a key logger / screen grabber detection system. This is a multi-pass process that identifies suspicious applications and then utilizes multiple techniques – runtime behavioral analysis, transparent honeypot, whitelisting, and heuristics – to refine and minimize false positives and false negatives (as described in detail above). The process does not make a value judgment as to whether detected applications are malware or legitimate; the intent is to detect any application that could potentially capture session inputs or displayed output. Once detection is made, an attempt is made to block the suspicious programs’ ability to capture keystrokes or graphical display (in the case of frame buffers). If the attempt to block fails (e.g., due to privilege escalation required), the DataDivider secure session has multiple administratively configurable action options (e.g., warn user, shut down session, or continue session but log event). Default options are always set to shut down the secure session.

Screen Capturing: Screen grab threat defenses operate in a similar manner to the key logger defenses. In addition to defeating screen capturing applications, the DataDivider secure session can block PrintScreen keyboard commands.

ScreenCapture protection is employed with the following techniques:

Process Injection techniques:

1) The DataDivider secure session injects into all processes running on the host at the same privilege level and lower

2) The DataDivider secure session hooks the following functions in GDI32.dll that capture screen/draw screen:
   - i) BitBlt
   - ii) StretchBlt

Whenever one of these functions is called we find all protected windows in the screen capture and clear out its contents.

3) The DataDivider secure session hooks the following DirectX API:
   - i) Inside dxgi.dll -> AcquireNextFrame

These are blocked and not allowed.

Global PrintScreen Keyboard Hook (Ctrl+PrtSc, Alt+PrtSc, etc):

We set a global message hook for the VK_SNAPSHOT virtual key and block it.

Detection and blocking of Mirror Drivers (VNC, WebEx, etc):

All display devices are enumerated and validated against the Mirror Driver service DLL, devices detected as mirror drivers are blocked unless whitelisted by policy.

Session Privacy

In order to defeat unauthorized users or applications, be it malware or legitimate but otherwise unauthorized application, from "digital dumpster diving" on data created or downloaded during the secure session, the Protected Zone provides real-time encryption of the disk-based session data, and performs a post-session secure deletion of these files.

The Protected Zone uses ECDH/DH AES cipher suites, and the fallback is AES for clients that do not support ECDH/DH ciphers.

Process Integrity

Injecting hostile code into a session's process space is a common exploit vector for malware. The Protected Zone provides several steps to prevent malware from getting loaded into the protected session to compromise data. This includes the following features:

- Launch of a new, clean instance of IE, watching for and preventing injection by external processes during startup. Validation of IE and of DLLs loaded at browser startup is done. Process injection protection is in effect for the duration of the session for all protected processes
- Preventing debuggers from attaching to protected processes.

- Administrative white listing of allowed applications, Browser Helper Objects (BHOs), toolbars, and other plug-ins - only those browser extensions required to operate the web applications. The result is that unnecessary BHOs including potentially hostile, data compromising extensions will not be loaded.
- Only the secure session is affected. There is no impact on other or subsequent browser launches.

SSL Certificate White Listing

The Protected Zone session will only allow HTTPS (i.e., SSL/TLS) connections to be established with servers whose certificate matches one of the whitelisted certificates, or has a CA certificate in its chain that matches one of the white listed certificates. This feature enables Merchants to utilize their inherent knowledge of their web applications to use only an enumerated set of SSL certificates and ensures the protected browsers' certificate store will not be bypassed. This defends against the possibility of spoofing the certificate validation process, as part of a multi-stage HTTPS Man-in-the-Middle attack.

Invalid Certificate Override

SSL connections are not allowed with a server whose certificate is not trusted by IE on the agent's machine (i.e., for self-signed certificates, expired certificates, mismatched domain name, etc.). It is a less restrictive option than certificate white listing for improving HTTPS security on the endpoint. This can be useful in preventing social engineering where HTTPS proxies attempt to fool users into accepting invalid certificates that enable intercepting / decoding of HTTPS traffic.

Network Destination Controls

Many browser attacks are vectored through malicious code planted on a legitimate site, which redirects a browser – at the HTTP / HTML level, for malicious purposes, unbeknownst to the user. Certain types of cross-site scripting (XSS) and cross-site request forgery (CSRF) are common examples. Others silently capture information and forward it to a malicious server using the browser's communication features. The Protected Zone networking controls prevent these attacks from making connections to unauthorized sites, even if the protected site itself is compromised. It does this by a white list of allowed network destinations in the policy. Before any attempt to open a URL is allowed, the destination is checked against this white list. As a side benefit, this feature also prevents users from manually browsing to arbitrary, potentially harmful sites.

Hostname Resolution Bypass

DNS has been revealed as a major attack vector over the last couple of years. If an attacker can poison a user's DNS server (or the user's local "hosts" file), they can make the browser (and the user) think they're connecting to a legitimate site (e.g., online banking site), when they're actually communicating with a malicious server that looks like the real thing. The Protected Zone can include a predefined list of hostname-to-address mappings for the customer's protected servers, which allows the protected browser to bypass DNS and local hosts file lookups altogether, precluding the effectiveness of these types of redirection attacks. In this manner the Protected Zone can limited access to only pre-authorized genuine websites and servers. For example, it would not be possible for malicious malware to spoof a Payment Service Provider's Virtual Terminal.

Information Controls and Auditing

Users are prohibited from extracting data from inside the protected session and making this available to the desktop environment (although no sensitive cardholder data is exposed to the protected session). The Protected Zone controls users' ability to download files, use clipboard functions including drag and drop, print, print screen and other related operations. This is achieved by trapping all relevant file I/O and vetting this against the DataDivider Protected Zone security policy. Only allowed content and/or operations can proceed. All attempts at such functions, even when restricted, are also logged.

For example, where an unmasked format preserved Token is returned from the Virtual Keypad into the agent's payment screen, it is not possible for the agent to cut and paste the Token (alternately the Token can be masked only visualizing the last four digits).

Virtual Machine and Remote Desktop Controls

There is a significant information security hole in applications running on an operating system that is in turn operating in a "hosted" system, such as a virtual machine instance or a terminal services (remote desktop) session. If the host OS has been compromised, or is under the control of a malicious user, key logging, screen scraping, and other methods of stealing content will easily bypass security software running on the virtualized system and fall outside of PCI compliance. At startup of the secure session a probe of system settings will detect any such hosted session. If detected the secure session is discontinued and the user is notified.

Session Time Limits

A 15 minute timeout is pre-defined and enforced. This will force a session to close if no activity within the session occurs in the specified period. In addition, a maximum time limit can be applied to sessions.

Blocking COM Snooping

Many web browsers provide programming interfaces (APIs) that allow external process to query and modify data in the browser's process. This is another vector for malware to steal information or misdirect the user. The Protected Zone DataDivider agent blocks all of these APIs for the protected session, effectively preventing this type of attack.

Zero Bypass Options

To ensure it is not possible to bypass the Protected Zone and all of its protections, access to all payment processing is limited exclusively to the DataDivider hosted environment. This is enforced by white listing the DataDivider hosted environment with the specific payment gateways for each client and removing the Merchant's IP range from their current white list. Additionally corporations may wish to block direct access to Virtual Terminals via their firewalls. These measures protect against any malware re-directing the agent to a bogus hosted payment page where a fraudster would attempt to harvest cardholder data.

Mitigation of Common Browser Attacks

Below is a list of some common attacks, and how the Protected Zone blocks or otherwise mitigates each. This is not an exhaustive list of web-centric attacks, nor is in an exhaustive list of the types attacks that can be mitigated by the Protected Zone. It's just a sample of some of the more common attacks which the Protected Zone guards against.

DNS Poisoning: To defeat DNS or Host file poisoning threats, the Protected Zone can be configured to bypass the customer computer's hostname resolution configuration (HOSTS file, DNS server settings). Affecting only the protected browser's hostname resolving, the Protected Zone pushes a virtual "hosts" file, or redirect DNS queries to DataDivider hosted DNS server.

Man-in-the-Browser (MITB) Attacks: MITB threats require the insertion of attack code into the target's browser process space, either via code injection techniques, loading of a hostile BHO (i.e., plug-in or extension), or exploiting vulnerabilities in a legitimate BHO/plug-in. The Protected Zone provides comprehensive anti-code injection mechanisms to deter this threat vector. In addition, the Protected Zone controls which BHOs are running within the protected session by use of whitelisting.

Drive-by Downloads: This type of attack leverages vulnerabilities in browsers, browser plugins, and document viewers to silently install malicious software on a user's machine when they visit an infected site, often without requiring any further interaction from the user. The Protected Zone's advanced bidirectional process isolation technology prevents malicious code delivered by these vectors from entering the confines of the secure session process and infecting the rest of the protected session. This technology is even effective against malware that leverages zero day vulnerabilities in browsers, plugins, and document viewers.

Man-in-the-Middle (MITM) / Session Hijacking: MITM and session hijacks must re-direct HTTPS traffic to a hostile intermediate site as well as induce the user to accept the site's SSL certificate. The Protected Zone uses multiple defense mechanisms to defeat such threats. These include: URL destination controls; SSL certificate whitelisting; and hostname resolution controls (see DNS Poisoning, above), which protect against network layer attempts to redirect HTTPS connections.

Key Loggers: Prior to starting the protected session, and then continuously throughout the session, the Protected Zone executes a key logger / screen grabber detection system. This is a multi-pass process that identifies suspicious applications and then utilizes multiple techniques – runtime behavioral analysis, transparent honeypot, whitelisting, and heuristics – to refine and minimize false positives and false negatives. The process does not make a value judgment as to whether detected applications are malware or legitimate; the intent is to detect any application that could potentially capture session inputs or displayed output. Once detection is made, an attempt is made to block the suspicious programs' ability to capture keystrokes or graphical display (in the case of frame buffers). If the attempt to block fails, the secure session will terminate.

Screen Capturing: Screen grab threat defenses operate in a similar manner to the key logger defenses. In addition to defeating screen capturing applications, the Protected Zone can block PrintScreen (PrtScn) keyboard commands. System wide attempts to capture the screen (e.g., print screen keyboard command) from outside the secure session are completely blocked. When another processes outside of the secure session tries to capture the screen image of the secure session (e.g., snipping tool), the actual operation is allowed but, the protected session images are rendered opaque. Therefore any attempt to access the video data related to the secure session from the Operating System is prevented.

Emerging / Zero Day Attacks: The Protected Zone does not rely on signatures or blacklists to provide its security mechanisms, and has proven to be effective in providing immediate defenses against new attacks (e.g., SpyEye) or variants of existing attacks (e.g., Zeus variants).

Non Browser-Based Defenses: The Protected Zone provides security mechanisms for the protected session. These mechanisms extend to processes launched from within the protected browser, such as browser tabs and sub-windows, and external. When the protected session ends the security mechanisms end and all session data (cookies, cache files, history file) is securely destroyed.

DataDivider Performance Impact

Using the DataDivider Protected Zone will have a minimal incremental impact on existing performance and session start-up. The SSH DataDivider agent launch process requires the downloading of a small software package to the agent's device. The time it takes to do this is directly related to network access speeds, but at typical broadband rates, takes a few seconds. Even this relatively minor impact can be mitigated by storing the DataDivider agent binary in a temp file area. If this option is enabled, the launch process will check the existence and checksum of the binary; if it matches this file is used instead of doing the download. The DataDivider agent, upon start of execution, performs its key logger and frame grabber scan of all existing processes. Generally speaking, on a moderately loaded system with typical broadband speeds and not caching the DataDivider agent, startup times are typically around 3-5 seconds and the agent must then log into the launched RDP session. Once logged in however, the agent will be connected to that session until the idle timeout or maximum timeout are exceeded. Depending on the functions performed by the agent, these timeouts can be configurable.

Summary

The combination of the DataDivider Virtual Keypad, Protected Zone and the Level 1 certified hosted platform ensures that cardholder data can be captured without exposing this data to the Merchant's local computing infrastructure. Furthermore, the merchant can complete their payment processes without further exposure to cardholder data and without any opportunities of the payment process being compromised. This substantially reduces the scope of PCI DSS for telephone payments.