Secure Medical Research Workspace

TR-11-01

February 10, 2011

Phillips Owen
powen@renci.org
Renaissance Computing Institute

Michael Shoffner
shoffner@renci.org
Renaissance Computing Institute

Xiaoshu Wang
xiao@renci.org
Renaissance Computing Institute

Charles Schmitt
cschmitt@renci.org
Renaissance Computing Institute

Brent Lamm
brent_lamm@unc.edu
University of North Carolina at Chapel Hill

Javed Mostafa
jm@email.unc.edu
University of North Carolina at Chapel Hill

TR-11-01
Renaissance Computing Institute

Secure Medical Research Workspace

Abstract

SMRW, the **S**ecure **M**edical **R**esearch **W**orkspace, is a comprehensive solution to protect Electronic Health Records (EHR). The SMRW utilizes virtualization technologies to facilitate the setup, data provisioning, management, and tear down of protected virtual workspaces. The virtual workspace will incorporate Data Loss Prevention (DLP) technologies and techniques to prevent unauthorized use and transmission of data in order to maintain compliance with Institutional policies and HIPAA data regulations.

1. Background – UNC Case Study

Healthcare professionals and researchers requesting Protected Health Information (PHI) from the UNC Health Care System (UNCHCS) clinical, operational, and administrative systems are provided data extracts sent via secure transfer methods (e.g. secure email) by the Medical Information Management (MIM) organization, after obtaining proper authorization. While this model served to ensure PHI data was provided only under sanctioned circumstances, and triggered the proper disclosure processes, there was no way to control the security of the data once provided to the requestor. Once the data was transferred to the requestor's control there is risk of the exposure of data due to lost or stolen portable storage devices (e.g. laptop computers, USB drives).

Users with legitimate access to sensitive data present a complex challenge in managing data loss risk. "Insider threats" collectively represent many potential vectors of deliberate or unintentional data loss (see Figure 1).

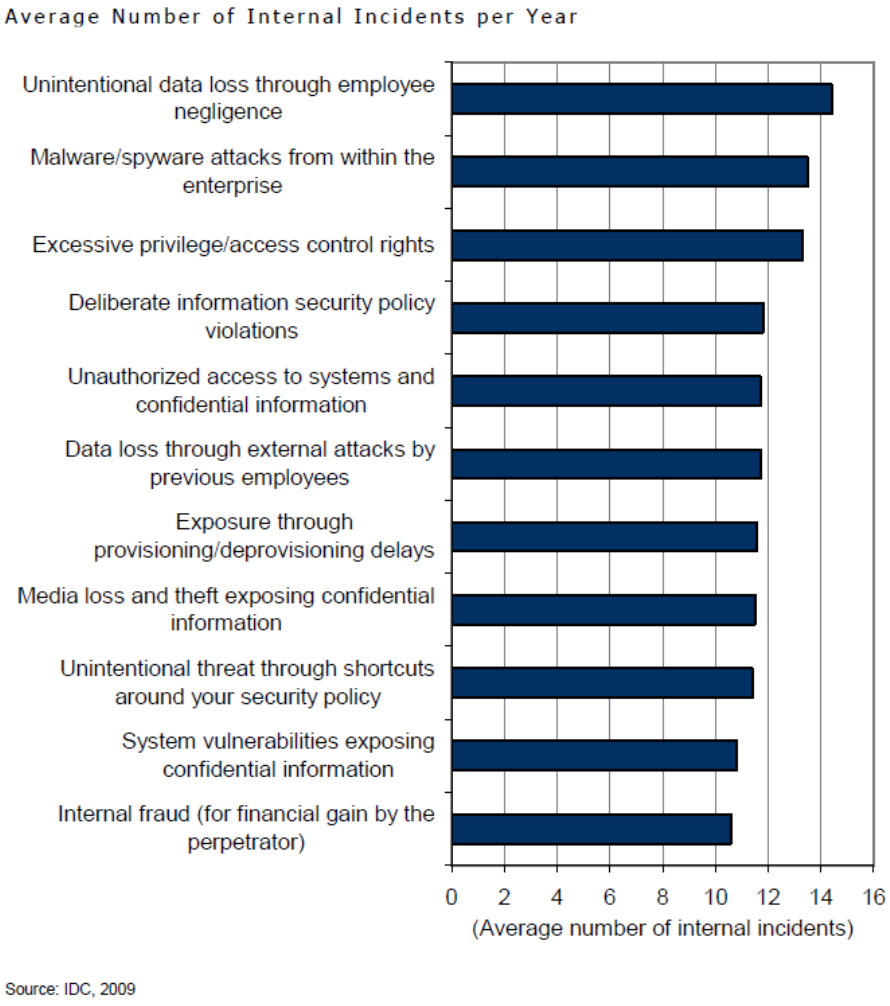

Figure 1 - Average number of incidents per year[1]

[1] 2009 IDC/RSA Whitepaper "Insider risk management a framework approach to internal security" page 11.

Unfortunately in the current UNCHCS data distribution model there is no known way to allow requestors access to open, manipulate, and analyze the data provided via the network-based file storage without also providing the ability to make copies of the data on portable storage. For example, requestors have the ability to copy/paste data files to other drives located on their local computing device, and they have the ability to open data files, extract information from the files, and save to new files on their local computing device, and printing to a local or remote printer. Requestors must agree that they will not perform these actions prior to being provided data, however there is currently no way to prevent accidental or purposeful violations of the agreement once the data has been provided to the researcher.

The Carolina Data Warehouse for Health (CDW-H) is an extremely rich resource available to UNC researchers for conducting studies in a wide variety of fields including medicine, public health, informatics, economics, and public policy. Currently, the system contains about 2 million patient records covering the period since July 1, 2004. With the inception of the CDW-H as the strategic central repository of data for quality improvement and health-related research needs, a new governance structure has been established to ensure all data access requests meet the necessary regulatory compliance requirements. In an effort to quickly improve the process of provisioning PHI data to approved requestors, this governance body identified, approved, and implemented an enhanced data provisioning model. This current model provides requestors with access to their approved data on a secure network-based file server, where they can utilize typical Windows file access techniques, such as mapping a network drive, to open, view, manipulate, and analyze their data within a virtual workspace (see Figure 2).

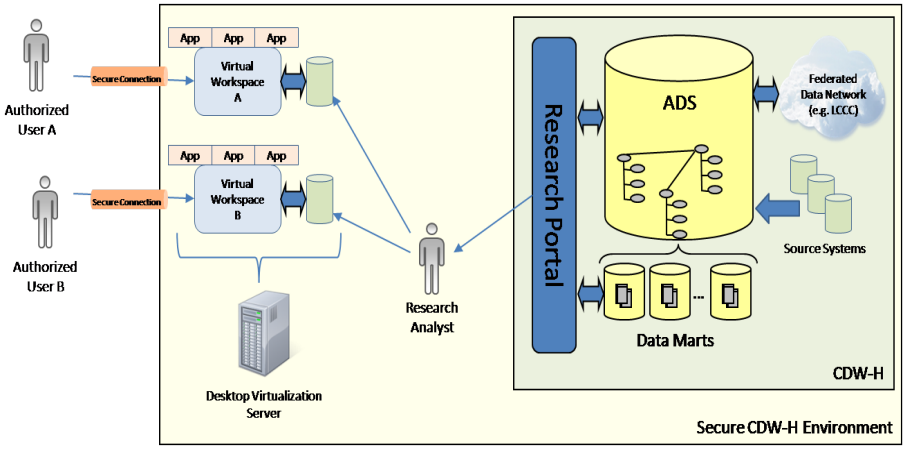

Figure 2 – Enhanced data provisioning model

This enhanced data provisioning model is a significant improvement over the previous methodologies since it does not require physical transfer of the PHI data to a requestor controlled storage device. However, there is a major limitation of this model, the inability to prevent copying of the data away from network-based file storage.

The need exists for a solution which will provide approved requestors with the flexibility of data access and analysis of the current security model, coupled with the ability to prevent physical removal of PHI data from a central and secure storage environment. In discussions with other entities, we have found that this need applies to other institutions that provision restricted data.

2. Objective

This document presents a model for a Secure Medical Research Workspace (SMRW) environment that reduces the risks of unauthorized access, loss of, and damage to patient information used by researchers. This document describes the model both in general terms and in its application to the UNCHCS. We also believe that this model applies broadly to any scenario where sensitive data is provisioned from a centralized data store to end users.

Information provisioned to researchers is usually a subset of a much larger data set. The root data sets by their very nature are strictly monitored and protected as a whole by numerous policies and techniques. However, data subsets provisioned to a researcher today do not have an equivalent level of auditing and monitoring to insure comprehensive protection outside of the main data warehouse.

Data security classifications, the corresponding risks, and business requirements have been identified and used to create data compliance policies. An important goal of UNCHCS is to strike a balance of the enforcement of these data compliance policies while also providing a minimum of inconvenience to each researcher. A listing of the underlying business requirements can be found in Appendix C.

The research environment will also provide business analysts and administrators the tools necessary to facilitate the distribution and management of virtual workspaces that contain sensitive Electronic Health Records.

To summarize the objective for all roles,

The workspace environment will provide the researcher with:

- A protected, standardized, and virtualized computing environment.
- Provisioned healthcare data gathered from various sources.
- Tools to work with the provisioned healthcare data.

The workspace environment will provide business analysts and administrators with:

- Security auditing and forensic capabilities.
- Research data usage metrics.
- Automated methods to facilitate the creation and distribution of virtual computing environments.

3. Technological Overview

3.1. Workspace Virtualization

Computer virtualization technologies offer many benefits to an organization. With respect to the SMRW environment this strategy offers the following advantages:

- Pre-tested virtual workspace pools
- Pre-configured and standardized workspaces ready to deploy
- Ability to quickly respond to "special environment configuration" requests

Base-line virtual image configurations are created and placed into a pool. Virtual workspaces in the pool are pre-packaged with a set of common software tools targeted to known research or administrative roles. In addition, workspaces are also loaded with protective measures such as firewalls, anti-virus/malware, and Data Loss Prevention (DLP) software packages. The base-line workspaces are rigorously tested prior to deployment to insure software package inter-compatibilities as well as meeting all applicable data security policies and regulations.

There are a wide range of base-line virtual workspaces in the pool so that an administrator can chose one best suited for the end user. Every effort is made to identify and create a virtual workspace pool that contains workspaces that need little or no customization prior to deployment. Providing a pre-configured pool of hardened virtual workspaces facilitates an agile and reliable administrative environment.

3.2. Virtual workspace experience

Virtual workspaces in the prototype program are pre-loaded with the Microsoft Windows 7 Enterprise operating system. Each virtual workspace is placed within a Windows Active Directory domain. Within this Active Directory domain multiple Group Policy Objects (GPO) are created and deployed to manage the Researcher's virtual workspace experience. GPOs are used to insure compliance with items such as:

- Acceptable internet URLs
- Desktop experience
- Adding or removing programs
- Launching of non-approved applications
- Approved peripheral devices

The implementation of an Active Directory/GPO structure results in a simplified model of remotely administering and propagating policies across users or virtual workspaces, either individually or in groups, quickly and efficiently.

3.3. Data Loss Prevention - DLP

To secure a workspace requires imposing measures that can be used to guard against potential security breach targeted toward the protected workspace. Here, the word "workspace" is used in a very general sense. It can refer to a physical computer, a virtual environment, or a corporate network where data is hosted and/or computed. Depending on the direction of information traffic to be guarded, security software is divided into two groups (See Figure 3).

Software that guards against in-bound traffic is intrusion protection system, which can be further divided into Intrusion Detection System (IDS) and Intrusion Prevention System (IPS). The former system detects, whereas the latter responds to, security threats. On the other hand, software system that guards against the security breach from out-bound traffic is called Data Loss Prevention (DLP) software, which is the main focus of this overview.

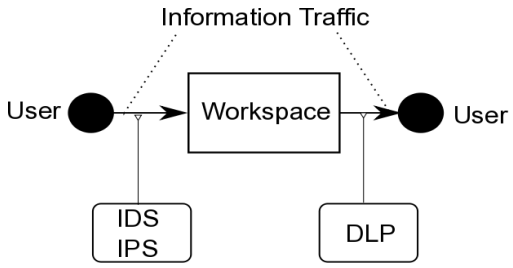

Figure 3 - Information Traffic flow

Although similar in nature, intrusion prevention and data loss prevention systems differ in purposes. Intrusion prevention system aims at preventing an outside attack, such as denial of service, and the penetration of active content with malicious intent, such as viruses or malware. DLP systems, on the other hand, are aimed at protecting sensitive data from leaking out of an organization. The protected content is mainly of two kinds: Private identifiable information (PII) and Intellectual property (IP). PII is characterized as private information. The concern of protecting PII is often aimed at complying with corporate compliance audits. IP, on the other hand, is more of secretive in nature. Compared to PII, IP's definition can be very broad and often has a straightforward definition, such as credit card numbers or a patient's medical billing code. Protecting IP often requires a more sophisticated detection algorithm and more complex management capabilities from a DLP product.

3.3.1. Main Features of DLP Software

Two major features distinguish DLP software from other security software. The first feature is that DLP's capabilities are always based on a set of central policies. In addition, these policies can be controlled or configured by appropriate personnel with regard to what kinds of data needs to be identified, what kind of activities need to be monitored, and how their usage should be protected. This controllable policy-based feature distinguishes DLP products from other products, such as an enterprise firewall whose security policies, once set, rarely needs to be modified by an IT administrator.

The second feature of DLP software is in its content-awareness. DLP software must have the capability to inspect the content of a data, often in real-time, in order to identify, classify, and subsequently react in accordance to a predefined policy. This content-awareness distinguishes DLP products from the traditional device control technologies which rely on data context (i.e., location and users) and are indirectly regulated by an external access control mechanism. Compared to context-based protection, DLP technologies enable data protection to be carried out on a much finer level of granularity.

Nevertheless, DLP should not be considered to be an antagonist to, or a replacement for, context-aware technology. In fact, these two technologies complement each other because whether the manipulation of a sensitive data violates a security policy cannot be solely determined by the nature of data. The enforcement of a DLP policy, in fact, can only be meaningfully applied when it leverages data context, i.e., who is using it and where the data comes from, and where it is sent to.

3.4. DLP product market scope

Comprehensive protection of a workspace requires security measures to be enforced at multiple levels. Data content inspection is just one of the many facets of the entire security requirement. The purpose of focusing on the above two features, i.e., content-awareness and policy-based action, is to narrow the scope of the market landscape. Thus, if a software product does not make its entry into in this review, it does not imply that the product is not useful in terms of securing a workspace. Take IPSwitch's product MOVEit DMZ as an example. As MOVEit DMZ allows customized policy to be set upon the network transfer of a particular file, it can obviously be used to protect a designated workspace. However, because MOVEit DMZ does not perform content inspection, it is therefore not considered to be a DLP product. Gartner's group, for instance, considers IPSwitch a leader in the market for managed file transfer (MFT) (see Figure 4) and does not mention it on the magic quadrant for DLP (see Figure 5).

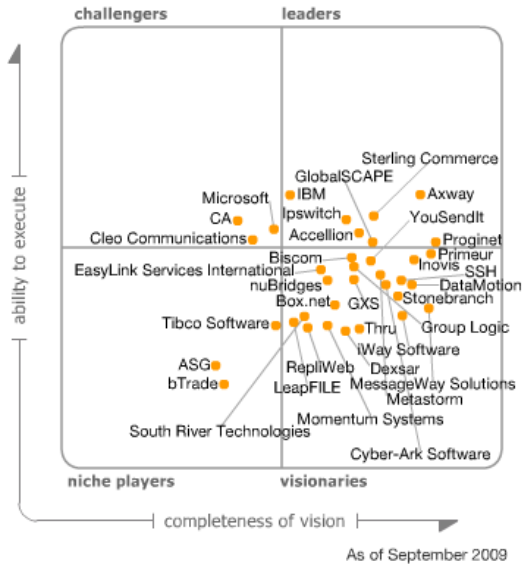

Figure 4 - Magic Quadrant for Managed File Transfer[2]

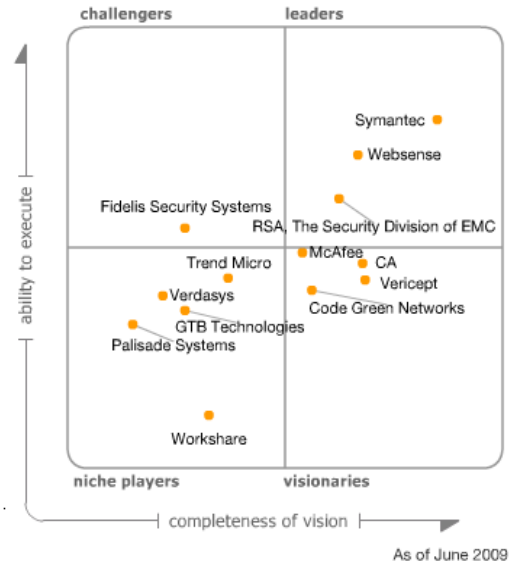

Figure 5 - Magic Quadrant for Content-Aware Data Loss Prevention[3]

Another prominent omission in the review is the VMSafe technology developed by VMWare. The reason for its omission is due to the different levels of protection. DLP offers protection at data- or application-level. VMSafe, on the other hand, offers protection at platform-level. Obviously, platform protection can only help data- or application-protection. As an operating system's security API, VMSafe can be an extremely useful mechanism for engaging a DLP product.

Naturally, there could be some other software packages that have been omitted – but not by intension rather by ignorance. As DLP software market becomes more competitive in recent years (with eight acquisitions since 2006), some vendors start to market their DLP capability as an added feature to their products that is used for a different purpose. For instance, Clearswift's products are both policy-based and content-aware. However, the company brands its product as a content-aware secure web/email gateway, as opposed to a DLP product. Because of this branding strategy, it is very likely this review has missed some relevant products.

[2] Source: Gartner, September 2009

[3] Source: Gartner, June 2009

4. Our approach

In developing a solution, we have considered similar approaches taken by other groups at UNC, vendor suggested solutions, and solutions taken by other entities. In general, most solutions we have seen do not offer adequate protection or configurability in how data is protected from leaving the work environment. Of those solutions that do offer a protection mechanism, most use a secondary storage location onto which data from the environment is copied. Data is then retrieved from this secondary location by the end user. While on the secondary data location, the data is available for auditing by environment administrators. Please refer to Appendix A for details on UNCs governing policies related to the workspace and the environment.

4.1. Secure Medical Workspace conceptual overview

The Conceptual overview of the environment is illustrated in Figure 6. For this solution, we define a Secure Medical Workspace (SMW) environment that represents all actors in the information data path.

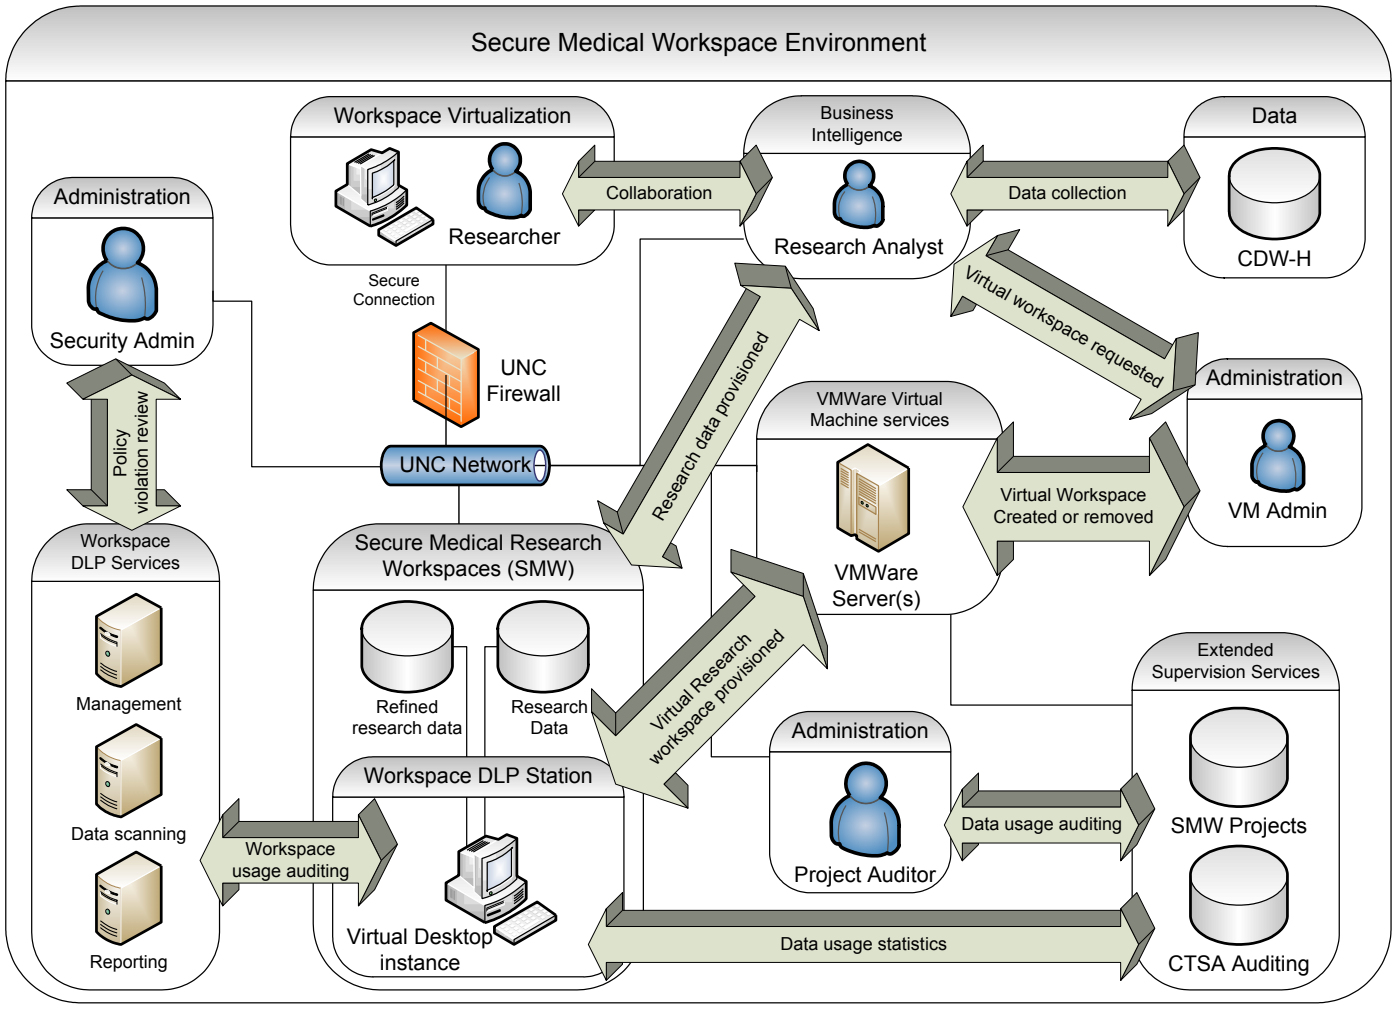

Figure 6 - The Secure Medical Workspace Environment

4.2. Process Actors

4.2.1. Researcher

A Researcher is a Principle investigator (PI), or a member of PI team, who has authority on the research project. This person also accesses and works in the virtual environment to perform research using the sensitive data that has been provisioned.

4.2.2. Research Assistant

A Research Assistant is a graduate student (or other non-PI team member) who accesses and works in the environment to perform research under the direction of the PI team.

4.2.3. Research Analyst

A CTSA Research Analyst is one who serves as contact point for the Research team (as well as other roles) for the provisioning of data subsets. This role will be managed by NCTraCS.

4.2.4. Virtual Machine (VM) Administrator

A Virtual Machine Administrator assembles sets of base images for use by researchers, hardens the images, and coordinates with the Security Administration for testing. This role also provides for sharing of images on an as-needed basis, for example between projects. This role will be managed by NCTraCS.

4.2.5. Project Auditor

A Project Auditor works with Researchers to set up various Research team accounts on the system. This role will be managed by NCTraCS and will also be responsible for monitoring the usage of the provisioned data.

4.2.6. Security Administrator

A Security Administrator audits and assesses the environment both at design and implementation levels (e.g., "attacks" the deployed system to assess vulnerabilities) by using a set of criteria and a testing suite. This role is a combination of initial setup and ongoing components that works with policies and personnel of NCTraCS, ISD, and UNC Information Security Office (ISO).

4.2.7. Outside Compliance Auditor (not depicted)

An Outside Compliance Auditor represents a regulatory agency who reviews the actions of the Research Team in the SMW environment. This person also monitors other aspects of the system's security and usage. This role will be conducted by ISD.

4.2.8. System/Network Administrator (not depicted)

A System/Network Administrator sets up, configures, and maintains overall environment infrastructure (e.g., networking, data environment). This role is a combination of personnel from NCTraCS and the Information Services Division (ISD) of UNCHCS.

4.3. Prototype "Commercial Off-The-Shelf" products (COTS)

The following software products were selected for the prototype evaluation:

- Desktop virtualization services - VMWare VSphere ESXi 4.1.0.
- DLP services - Trend Micro DLP version 5.5
- Anti-Virus/Anti-Malware services - Symantec Endpoint Protection version 11.0
- Researcher workspace - Windows 7 Enterprise (x86)
- Researcher tools – Microsoft Office Suite, SAS, etc.

4.4. Workspace monitoring

Comprehensive monitoring of virtual workspace operations are accomplished with a combination of the built-in abilities of the DLP software services as well as standard Microsoft Windows Active Directory functionalities.

4.4.1. Auditing

Data policy violations are captured by the DLP services software. Depending of course on the DLP product selected, various levels of details of the violation event are captured. The minimum level of event details captured must include:

- The name of the policy violated,
- The virtual workspace that was used in the violation,
- The data that violated the policy,
- The timestamp of the violation,
- The action taken at the time of the violation,
- And finally, the individual responsible for the violation.

4.4.2. Forensics

Data leakage forensics will use captured audit events for analysis in the event of a data leakage security breach. By looking at what data was leaked and analyzing the various auditing events the "who/when/where" of the source of the event can be traced back to the root cause. The analysis of these events can be leveraged to develop policies that will prevent similar events in the future.

Data usage and policy violation events must be archived over long time spans for the potential of forensic analysis at a much later date. In addition, data usage trends should be analyzed periodically to determine if a particular policy should be created, or in the case of existing policies, bolstered or relaxed.

4.4.3. Windows auditing

Workspace usage characteristics are captured by the Microsoft Windows Domain Controller and contain items such as log-in and log-out times as well as when peripheral items (such as printers) are utilized. Other events related to the Windows Domain are available for tracking and require targeted setup involving the creation and implementation of Group Policy Objects (GPO).

4.5. DLP Evaluation criteria

There are a number of criteria to consider in the evaluation of a third-party DLP software product. In our evaluations high priorities were given to comprehensively audit the activities of the end user as well as the ability to provide intricate DLP detection customizations.

Requirement
Auditing of all events, including logging of data that was moved (or attempted to be moved) into or out of the system and any user justification reasons
Ability to customize interactions with user, in order to prompt user for authorization to continue or inform user of denial/success of operation
Ability to enable and disable web access
Ability to restrain web access to prevent data downloads
Ability to restrain web access to prevent data uploads
Logging of all web sites accessed and web forms submitted

Ability to enable/disable data transfer in/out via file copy
Ability to enable/disable data transfer in/out via system clipboard, email, and IM
Ability to enable/disable data transfer in/out via other applications not listed (e.g. ftp)
Usable applications for data transfers in/out
Active filtering of data transfers in/out
Ability to define and deploy new filtering algorithms
Support for DMZ zone for file transfers in and out of the environment

In addition to the basic requirements listed above, the behavior of the DLP product in certain situations should be observed and well understood. Below are some of the basic conditions used in our evaluation.

Action	**Expected result**
Copy a local file that has policy violation information to a protected location.	File not copied.
Copy a local file that has NO policy violation information to a protected location.	File copied successfully.
Open a file located in a protected drive (with Microsoft Excel) and added a policy violation. Exit with changes and indicate to Excel to save changes before closing.	File not saved.
Open a file (with Microsoft Excel) that is located in a protected location that already has policy violation information. Make no changes to the file and exit.	File not saved with DLP interaction.
Open a file located in a protected drive (with Microsoft Excel) and added a policy violation. Exit with changes and indicate to Excel to NOT save changes before closing.	File not saved with no DLP interaction.
Cut/Paste sensitive information from a file to a protected destination.	Operation aborted with DLP interaction.
Print sensitive information to a protected destination.	Operation aborted with DLP interaction.

5. Conclusion

In this document we have presented an explanation of the scope and challenges in protecting sensitive patient health information used by researchers. In addition, we have summarized what data leakage prevention means and looked at various data leakage protection methodologies used to deter, detect, and defend against the unauthorized or unintentional movement of data. Finally we have discussed the notion of a Secure Medical Research Workspace as a virtual workspace where a researcher can perform their studies in a managed environment.

The notion of implementing a virtual workspace has many benefits for researchers and administrators. The researchers are provided with a standardized desktop pre-configured with everything they need to perform their studies. The administrators are equipped with the ability to quickly create reliable environments that have been exhaustively tested prior to deployment. Administrators are also equipped with the data and forensic tools necessary to monitor all environment activity as well as to analyze and prevent data policy violations.

The usage of information provisioned to researchers must be supervised. A derived data set distributed from a data warehouse must be given the same security considerations as the data housed in the warehouse itself. Understanding the capabilities of DLP products, coupled with a thorough understanding of how data is used and who is using it within an organization, is critical in the proper design and implementation of a environment that will be able to proactively protect against data policy violations.

Appendix A: Policies governing operations of the SMRW

Data Export Policies: The policies are used to define and enforce decisions concerning the movement of data from within a SMRW to external systems, such as disk space on a user's personal computer. This covers what data should be controlled from being exported and how should it be controlled. Policies need to address:

- User-level restrictions and permissions
- Restrictions on types of data that can be exported, including active filtering of the content (e.g., should data be scanned, if so what algorithms should be used with what parameter settings)
- Conditions under which data can be exported, including:
    - Data Use Agreements and IRBs that users must acknowledge
    - Time constraints (e.g., results cannot be exported until 6 months after data receipt)
    - Preapprovals from groups
- Auditing requirements by compliance officers for exported data
- Notifications that must be provided regarding exported data (e.g., to the data providers, to data providers external to the SMRW environment)
- Restrictions on data exports based on receiving resources (e.g., mapped drive, USB, another SMRW within the same secure environment, another SMRW in a different secure environment, DMZs)
- Restrictions on mechanics of data export allowed (SCP versus FTP, encryption level)
- Restrictions on printer usage (e.g., remote printing allowed, print material requires scanning)
- Restrictions on screen captures
- Actions to be taken under violations of policies

Data Storage Policies: These policies are used to define and enforce decisions concerning the storage of data within a SMW. Here we focus on standard access control rights around data and standard user-level access control.

Data Import Policies: These policies are used to define and enforce decisions concerning the movement of data into the SMRW. In general, we are more concerned on the implications the import of data has on the policies enforced on a SMRW.

- User-level restrictions and permissions
- Restrictions on types of data that can be imported, including active filtering of the content (e.g., should data be scanned, if so what algorithms should be used with what parameter settings)
- Based on type of data imported, what other SMRW policies are triggered (e.g., does importing of data with PHI from another SMRW automatically disable data export from this SMRW)
- Restrictions on ability to import data based on applications (e.g., disable export from database engines or business intelligence systems).
- Allowing for import of SMRW policies alongside the import of data.
- Auditing requirements by compliance officers for imported data
- Notifications that must be provided regarding imported data
- Restrictions on mechanics of data import allowed (SCP versus FTP, encryption level)
- Actions to be taken under violations of policies

SMRW Access Policies: These policies are used to define and enforce decisions concerning the ability of users to access a SMRW.

- User-level restrictions and permissions to use SMRW, including time and resource based restrictions/permissions
- Policies regarding the resource from which the SMRW is accessed (e.g., from mobile devices, from IP addresses external to a campus network, access protocol such as the versioning of RDP)

SMRW Control Policies: These policies are used to define and enforce decisions concerning the administration of the SMRW configurations.

- Standard user level admin policies (e.g., admin control, application usage policies, security patching)
- Mappings from SMRWs to VMs (e.g., specific mappings from a SMRW to a specific pool of VMs)
- VM re-imaging policies (e.g., VMs are re-imaged from a base image once a month to eliminate rootkits)
- Policies around scanners (malware, spyware, and anti-virus scanning)
- WWW specific policies (e.g., web browser policies, use of white and black lists of site) that go beyond data export/import policies
- Ability of user to alter policies
    - Notifications that must be provided user modification of policies
    - Auditing requirements for user based modifications of policies
    - Actions to be taken under violations of policies

Appendix B: Business Requirements

The abbreviated business requirements for the proposed solution are:

1.) Research teams that are granted access to data provided by the UNCHCS are provided an environment that allows the team to freely work with sensitive data, but that restricts the ability of the research team to remove any sensitive data from the environment.
2.) Research teams making use of the provided environments will be provided with a description of the policies governing their use of the environment and detailing the terms under which data can be removed from the environment.
3.) The environment must allow for removal of data from the environment under the following conditions:
    a. Data that is not restricted by the policies should be removable. This will typically consist of results derived from analysis of sensitive data that does not include sensitive data. A mechanism should be provided to easily allow removal of such data.
    b. Data that is restricted by policies will be removable only on a case by case basis and only after approval by the UNC Tracs Institute. A mechanism should be provided to allow Tracs system administrators to remove such data from the environment.
4.) The determination of whether data is restricted or not for removal purposes will be determined by one of the following methods.
    a. The environment administrators will determine the appropriate method on a case by case basis.
    b. The research team after review of their data use agreement and after an electronic confirmation is made by the research team.
    c. Automated procedures for determination of whether data is sensitive. In the case of automated determination, auditing of the data transfer and the process used to make the determination is required.
5.) The environment should provide at least one mechanism to remove data from the environment. The environment administrator should determine which mechanisms are enabled on a case by case basis.
6.) An audit trail should be maintained of any data removed from the environment. In addition, any data removed from the environment should be copied to a secondary location approved by the Tracs Institute for review for compliance with the policies. The audit trail should include the user name, date of removal, and policies in place for the research team. Any agreements by the user of their actions should be noted in the audit trail.
7.) The environment should provide the resources critical for researchers, including software tools, operating system, and external resources such as printers. Resources will be determined by use cases.
8.) The environment should have capabilities for managing the setup, maintenance, and shutdown of all resources.
9.) The environment should fit into the workflow being developed by the Tracs Institute and UNC for research on biomedical data, which generally includes the request for IRB, work with researchers to

identify data, extraction and provisioning of data, publication of results and relevant data, and deletion and/or archival of research data.

10.) The environment should provide security measures aimed at preventing unwanted intrusion by either other humans or by programs (viruses, malware).

11.) The environment should provide an audit trail of all activities within the environment, including data movements into and out of the environment, programs run within the environment, and user access times to the environment. The audit trail should also keep a copy of all data removed from the environment and any authorizations made to allow removal of data from the environment. Audit trails should be reviewable for compliance purposes by authorized personnel (to be determined by the Tracs Institute) and audit trails should be kept for a period of 7 years after the project end.

12.) The environment should provide a way to associated all resources associated with the environment with the research team's project.

13.) The environment should use existing UNC authentication and authorization mechanisms when possible.

14.) The environment should allow import of external data. However, import of external programs should be disallowed unless the research team has approval from NCTraCS.

15.) The environment should receive an annual review to ensure compliance with UNC and UNCHCS security policies.

Appendix C: Abbreviations and Acronyms

Abbreviation or Acronym	Value
API	Application Programming Interface
CDW-H	Carolina Data Warehouse for Health
CSV	Comma Separated Values
CTSA	Clinical and Translational Science Awards
DLP	Data Loss Prevention
DMZ	De-Militarized Zone
FTP	File Transfer Protocol
GUI	Graphical User Interface
HTTP	Hypertext Transfer Protocol
IAM	Identity and Access Management
ICAP	Internet Content Adaptation *Protocol*
IDS	Intrusion Detection System
IPS	Intrusion Prevention System
IRB	Institutional Review Board
ISD	Information Services Division
MIM	Medical Information Management
NCTraCS	North Carolina TraCS Institute
PDF	Portable Document Format
PHI	Patient Health Information
PI	Principle Investigator
RENCI	Renaissance Computing Institute
SCP	Secure Copy
SDK	Software Development Kit
SMRW	Secure Medical Research Workspace
SOA	Service Oriented Architecture
SPAN	Switched *Port* Analyzer
UNC	University of North Carolina
UNCHCS	University of North Carolina Healthcare System
USB	Universal Serial Bus
VM	Virtual Machine
XML	Extensible Markup Language