tiCrypt Use Case Examples

Last Updated:Thursday, July 6, 2017 in
partnership with the UF Research
Computing Department

www.terainsights.com
info@terainsights.com
thomas@terainsights.com

tera insights

tiCrypt Use Case Examples

In this document, we describe a number of workflows that are common among researchers using restricted data and we show how the work can be accomplished efficiently using the secure environment. We will try to show that the work is almost the same as in any traditional environment for working with data. The advantage of working with tiCrypt is that the risk for data compromises to individuals and to the institution are much reduced, compared to traditional implementations.

We will explain a few examples of increasing complexity to convey how a user can translate the process they use in a traditional infrastructure to the process in tiCrypt. Most researchers are familiar with the way research workflows are supported by desktop computers connected over a secured network to server computers that provide network shares with the restricted data to the workstations. Confidentiality of the data is ensured by assigning roles to the users that translate in the membership of groups. The groups then control access to specific folders on the shared drive.

Another secured way to support the workflows is by using virtual desktop infrastructure (VDI). Here the desktop computer connects to a VDI server which presents a desktop view to the user, but all work is done on the VDI server and data never leaves the server. The VDI server connects to the file server over a secured network and permissions to see or modify data are handled by roles, as in the previous case.

Example 1: Single researcher with data use agreement

Many research projects involve work on data that has been given under a data use agreement (DUA) from an external organization by an Institutional Review Board (IRB) agreement. The data is then deposited into the tiCrypt environment by a data custodian in the form of one or more files in the vault. The files can be organized in a hierarchy of directories.

The data custodian can then share the data with the researcher who will do the work with the data. The user now has authorized access to the data.

The user creates an encrypted drive to hold the data for analysis within a virtual machine. The user defines a VM of the appropriate type with the required software for the work and selects to attach the encrypted drive to the VM. This VM setup is then remembered by the system, so that a simply click will start the VM. The user will connect to the VM by the mechanism specified by tiCrypt. Different VMs have different mechanisms, including RDP, web page, or console. This can be done in one session or multiple session spanning several weeks of work.

Once the VM is running, the user can use the transfer window to move data between MyVault and the encrypted drives on the running VM. Note that data can only be moved in and out of an encrypted drive when it is mounted on a running VM.

As needed the user can make backup copies of result files by copying them from the encrypted drive to the vault. Files in the vault are replicated to another disk storage device and backed up to tape, fully encrypted, once per day.

When the work is complete the researcher can share the results in the vault, with the data custodian or whomever needs the results and delete the encrypted drives. Results can also be downloaded to the local computer, though the research is required to acknowledge that they assume liability for properly managing restricted data and all downloads are logged.

The data custodian can then revoke the sharing access to the data. All the sharing and copying activities are logged with the unique public key as identifier for audit purposes. There is no way for the data custodian to verify that the encrypted drives have been deleted or that an extra copy of the data was not taken out of the encrypted drive into the Vault in a different folder.

The workflow is very similar to sharing data in special folders created on a file server and working on the data with a virtual desktop session as far as the actions needed from the researcher and the data custodian. Because of the high degree of confidentiality and strong separation of data and work paths the process inside the tiCrypt environment significantly reduces the risk of unauthorized data access and disclosure.

Example 2: Faculty supervising multiple projects

Often the research group led by a faculty member is involved in multiple research projects. Groups can be set up with the correct members to limit access to any data set to only authorized individuals, just like with groups on a file server. These groups reflect the IRB approvals in place for each project. The researchers can then work on the data they are authorized to access as described in example 1.

Example 3: Multiple researchers working on shared data

Some research projects involve collaboration of multiple people working on the same data set. If the data is read-only then this workflow can be easily be accommodated with high level of confidentiality in the tiCrypt environment. The shared data can be copied into a special encrypted drive that can be mounted in any VM used by any authorized researcher to carry out their part of the work. The tiCrypt environment logs all accesses to the virtual drive. Thus audit logs will contain the necessary information to investigate any potential security incident. The traditional approach of shared write access to data on file servers is inherently insecure and does not ensure compliance. As such this method of data sharing is not an option within tiCrypt. It must be kept in mind that many of workflows were developed a long time ago.

It is therefore important to carefully assess the need for shared write access and look for alternative workflows that accomplish the same research objectives. It must be kept in mind that many of these workflows were developed a long time ago when the cost and capabilities of computer systems and software was very different from what it is today.

Example 4: Researcher developer working on data

A researcher developing algorithms in Python, R, SAS, and even coding in C/C++ can be supported in the tiCrypt. Users do not have admin access to VMs, but most software development environments allow users without administrator privileges to add modules and scripts and to compile source code. The only caveat being that the user must make sure that all installed additions are installed in the home drive and not in the system drive. In most cases installation in the system drive will fail, but there may be a rare case where it appears to work and then the installation will not work correctly after the VM is restarted.

Example 5: Staff maintaining VMs

Research computing staff members will create and maintain the VMs that are available to authorized users of tiCrypt. The process of creating and maintaining such VM is carefully controlled to ensure the integrity of the tiCrypt environment. The VMs are assembled in a special system, called vmbuilder, outside the secure perimeter of tiCrypt. When the VM is ready for deployment it is tested by an authorized administrator and the image is packaged and digitally signed. Then it is inserted into tiCrypt where it will replace an existing VM or will appear as a new VM on the menu.

Research Computing has a process whereby certain individuals supporting researchers can be trained and certified to build and maintain VMs for that research group. This way changes to VMs can be implemented in the required timeframe and Research Computing staff does not become a bottleneck for research.

Example 6: Researcher creating restricted data

Some research generates data that is restricted as opposed to using data granted by a data use agreement (DUA) or Institutional Review Board (IRB) protocol. The basic workflow is the same as in the other examples, with the main caveat being that the researcher should make sure that the newly created data is properly backed up. Data that resides in the encrypted disk will be backed up, but there is the possibility that the state of the virtual disk is not consistent when the backup copy is taken if a VM with write access to the virtual drive was running. The safest way is to copy the file from the virtual drive to the vault; then it will be replicated and backed up to tape without problem.

Example 7: Exporting data

Data can be exported from tiCrypt using the web browser and other mechanisms for large files. User accounts can be restricted so that download is not possible. This is a decision to be made by the faculty sponsor of the user.

It is the responsibility of the faculty research supervisor and all users to make sure that all legal requirements are met before downloading any data from the secure environment into a less secure or unsecure environment.