Statement by Hamilton County Sheriff's Office
Regarding
Arbitrator Server Crash

The Hamilton County Sheriff's Office Network

The Hamilton County Sheriff's Office ("HCSO") Computer Network ("Network") is a vast collection of interdependent systems, comprised of hardware and software, interconnected by a number of network switches dispersed throughout numerous physical facilities. The heart of the HCSO Network resides in the Administration Office, inside a dedicated server room. The HCSO manages its own Network domain, which is a named grouping of hosts and servers with managed login, access to resources, and permissions.

The HCSO's Network domain is housed in a series of large, rack-mounted servers, which function as the "brains" behind the network. These servers have a specialized operating system ("OS") installed and running on them. This OS is provided by a third-party vendor referred to as VMWare. This specialized OS allows the HCSO to install and run multiple operating systems (such as Windows) on our network.

The word "server" typically leads one to visualize a physical computer that sits in a network rack. In today's modern enterprise IT environment, however, physical space is not always plentiful. One of the considerations building a network is a need to maximize the physical space inside a network closet. The answer to this conundrum is "virtualization."

"Virtualization" is the practice of creating managed slices of server resources that work exactly like separate stand-alone physical servers with their own separate operating systems. From the user's standpoint, there is no functional difference than working on a physical server. The appearance and functionality are the same, thus the term "virtual." This allows the organization to compartmentalize the various informational platforms it uses. These operating system instances are referred to as virtual machines ("VM"). The advantage of a VM is clear - .less physical rack-space, multiple servers.

Virtual machines communicate with a large amount of disk storage referred to as a storage area network ("SAN"). A SAN is a dedicated, high-speed network of disks, network cards and processors. A SAN is primarily a way of creating virtual hard drives and virtually attaching them to any physical or virtual computer via the network so that those virtual hard drives work exactly like a physical hard drive physically attached. A SAN uses server and network resources to make virtual slices of disk space available to other physical and virtual servers as needed. SANs were adopted by the HCSO to improve application availability to its employees and performance of the applications by segregating storage traffic from the rest of the network traffic. Storage traffic is a term used to describe the network load that is involved when a disk is being tasked with sending information that requires a large amount of data to be transferred from that disk repository, such as files, etc. Network traffic, which usually consists of very small amounts of data, is not as intensive. The benefits of segregating these types of traffic through the switches is that a network engineer gains more efficiency and speed.

SANs also enable the HCSO to more easily allocate and manage disk/storage resources, achieving better efficiency. Instead of having isolated storage capacities across different servers, a pool of disk capacity can be shared across many different workloads and "carved" up, allocating various sizes of storage to a resource as needed.

As a matter of course, this system of disks is quite large and must be formatted in such a manner as to ensure that if one drive fails, other drives can continue to provide operability. To that end, the HCSO SAN is formatted in a method referred to as "redundant." This is a generalized term to describe a method of formatting disks, more technically referred to as a Redundant Array of Independent Disks ("RAID").

In a RAID, data is spread across multiple disks, to ensure that if one or more disks physically malfunction, the data is saved in such a way that a technician can simply replace the malfunctioning drive. The system then adapts and copies the data back to the new drive and the entire operation continues unaffected. The HCSO has over 50 drives in its SAN, comprised of multiple RAID configurations.

Redundancy is not "back-up." Back-up is a term used to describe a number of methods, both hardware and software based, which make actual copies of large portions of disk space to use in the event of a disk failure or when network engineers make changes that cause unexpected results to a network's behavior. In cases such as these, they would then "roll back" to the state that the network was in prior to the changes they made, in order to restore functionality to the desired result.

The 3-2-1 Backup Protocol

The "3-2-1" backup protocol is an easy-to-remember acronym for a common approach to keeping data safe in almost any failure scenario. The protocol states: keep at least three (3) copies of your data, and store two (2) backup copies on different storage media, with one (1) of them located offsite. The 3-2-1 protocol is not required by any entity that certifies any procedures of the HCSO. Rather, it is a practice within the information technology sector.

The HCSO does observe this protocol with its core data. However, backup of video storage is different. Core data consists of anything other than video data, including, but not limited to: all user files, email, records management systems, SQL data, booking records, booking data, etc. Video data consists of over 100 TB of only video footage. The HCSO designates the video data separately from the core data because of its volume.

Because of the volume of in-car videos and the length of time that they are retained by the HCSO, the requirements to deploy a fail-safe method of back up for these particular files have been too cost-prohibitive for a government agency of the HCSO's size. Backing up video using a 3-2-1 protocol (over 100 TB of files) would require a storage solution costing tens of thousands of dollars, as well as the employment of additional IT staff. It would require at least two (2) similar allotments of disk space in addition to the production disk (including room for growth), the

software and hardware to manage these back-ups and the manpower to dedicate to the increased volume of tasks associated with maintaining this level of network infrastructure. A "production disk" is a term used to identify the disk that is actually running the service and to which users connect. All other versions are back-ups. The back-ups are reserved for their explicit purpose of backing up data, rather than being used in day-to-day operations.

The software that the HCSO uses to perform back-ups is produced by a company called VEEAM. The HCSO uses this third-party software to back up user files, case management files, server OS/configurations, email, etc. These are core components of our network and though voluminous, they are maintained in compliance with the 3-2-1 backup protocol.

The misconception proliferated in the media is that the HCSO provided no back-up solution for video footage. The HCSO does have a method of "back-up" in place for the Arbitrator VM system; however, the decision was made at the original implementation phase that the 3-2-1 protocol backup of this video storage would not be possible due to cost prohibitions. The HCSO proceeded with a partial back-up protocol and procedures within the software to monitor the health of the video storage regularly.

One of the prime features of the VMWare system is that it has the ability to take what are referred to as "snapshots" of the various disk file structures along with varying other aspects of the disk. A snapshot is a catch-all term used to describe a collection of data, as it appeared at the time that the snapshot was created. When another snapshot occurs, that particular "instance" is another, unique collection of data, as it existed at that time. It could consist of anything, but in the case of the HCSO, the contents of the snapshots in question were the files that make up the Arbitrator virtual machine. Once snapshots are on the disk, a third-party backup application entitled VEEAM takes over and compares this snapshot with information in its back-up repository, in order for VEEAM to determine what changes have occurred over time. Once VEEAM makes this determination, it then consolidates all of these "snapshots" and folds them back into the disk as one entity, discarding the unnecessary snapshots.

Unbeknownst to HCSO Networking, in spite of the HCSO's decision to set up the VEEAM software to bypass the potentially destructive nature of "snapshotting" the Arbitrator video files, VMWare was actually taking snapshots of these video files covertly, with no inherent ability for their software to notify or even indicate that snapshotting was taking place. The fact that snapshots of such a large quantity were taking place covertly and then failing to consolidate, caused them to pile up and consume the allotted size of video disk space, which in turn, caused the Arbitrator system to begin to slow down and be unresponsive.

The Arbitrator System Crash

On the morning of January 10, 2020, around 9 AM, one of our network technicians discovered that the Arbitrator video server was running unusually slowly. A technician began to inspect the VM (server) in question and noticed the machine was displaying a need for disk consolidation in VMWare.

The technician then shut down this VM and attempted a disk consolidation task, which failed. He attempted several more times to consolidate the disk and get the VM to boot back up. Each attempt was unsuccessful. Having exhausted this option, the next step was to call the vendor, VMWare, to troubleshoot and attempt to resolve the issue.

VMWare's first step was to look at the internal VMWare VMFS filesystem on the SAN to assess the status of the Arbitrator virtual machine. The initial assessment was that the VM was running with 92 snapshots piled up on the disk. Per VMWare's specifications, any machine holding more than 30 snapshots would be running a high risk of data corruption.

VMWare then:

1. Attempted to consolidate the VM disk through the program's graphical user interface (GUI), which failed, leaving only one snapshot consolidated;
2. The VMWare technician then moved the corrupted VM to another datastore (a totally different disk SAN) – This was performed by VMWare in order to attempt to get the VM to boot up. This was also unsuccessful; and
3. The final attempt by VMWare to correct the issue was to clone or copy the VM to another brand new VM (in effect a fresh copy of the Arbitrator server). After about 8 hours of this attempted cloning, this method also failed.

At this point, VMWare had exhausted all of the options available to them. The decision was made by HCSO Network technicians to begin performing two simultaneous tasks. One was to establish an initial cause of potential data loss, and the other was to begin investigation into a third-party data recovery option.

Investigation by the HCSO of the Arbitrator System Crash

The investigation by HCSO Networking began with VMWare. The following questions were asked of the vendor:

> **Q**.) We setup our VMWare user interface to warn us if any of our VM's had a snapshot in their environment. Why were we not warned about these snapshots?
>
> **A**.) Per VMWare's design, third party Application Programming Interface ("API") snapshots (VEEAM) are not shown to the end-user GUI of VMWare without writing a customized alarm in VCenter (component of VMWare).
>
> **Q**.) Why was the HCSO not informed of this requirement?
>
> **A**.) VMWare stated simply that this was the "way the program was designed."

At this point, HCSO IT ended its questioning of VMWare. We did request and receive assistance from VMWARE with creating a custom event notification, wherein an email message is sent to

the HCSO that there are snapshots that are not consolidating. This event notification has been tested and found satisfactory, and is currently in use.

The HCSO Networking Investigation continued with VEEAM. The following question was asked of VEEAM:

> **Q**.) Per VEEAM's logs, snapshots have not been successfully consolidated on this virtual machine since December 31st 2019. Why was this not brought to our attention via a warning or error?
>
> **A**.) Per VEEAM's design, this event is considered informational only, and not a remarkable error or warning, and therefore is never brought to the user's attention.

The screenshot of the support log below illustrates where the error occurred. Line 5 is the failure point:

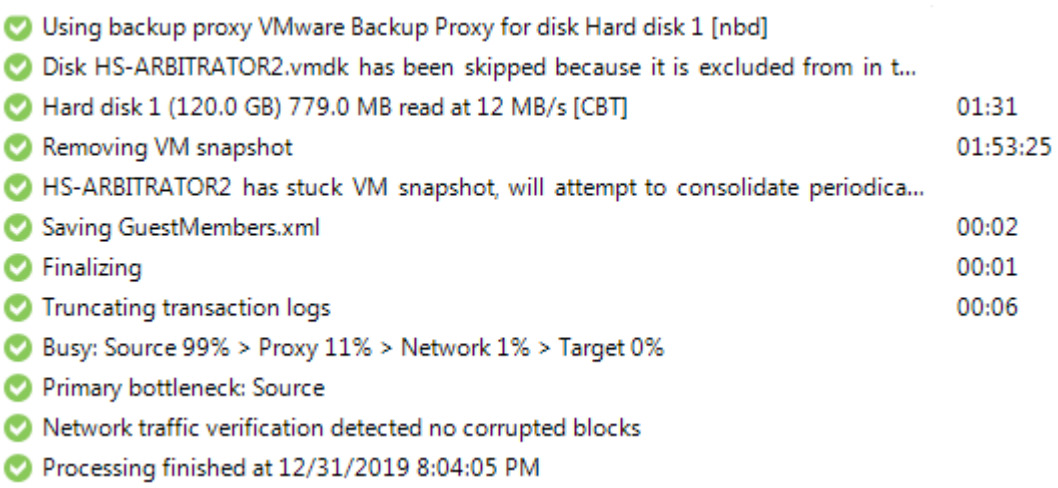

At this point, our questioning of VEEAM ended. It was determined that any further questioning of the company's engineering practices on such a critical point would be fruitless.
The HCSO concluded that the data loss was caused by a combination of several design factors of the two separate vendors, primarily:

- VMWare's design decision to not show existing snapshots in the GUI that are taken through their API; and

- VEEAM's design decision to treat unsuccessful disk consolidation as an informational event instead of a warning or error.

Attempted Data Recovery

The HCSO contacted three different vendors for data recovery. The one that the HCSO determined would be most capable was DriveSavers. They instructed the HCSO to purchase a disk drive large enough to get the VM's disk file and remaining snapshots copied to it and overnighted to them for investigation.

At this point, the HCSO Network staff had two mandates: (1) perform the requested copy for DriveSavers; and (2) re-format the existing video drive space on the SAN in order to set up a new clean and working version of the Arbitrator virtual machine and storage space for the continued operation of our officers' in-car video footage that was still being generated on a daily basis and waiting for upload.

The copy process took approximately ten (10) days. On the morning of February 3, 2020, the device (with the data) was shipped off to DriveSavers. Once copying the data was completed, the re-format and new virtual machine set-up was initiated successfully. On the morning of February 17, 2020, Adam Marthaler spoke with DriveSavers and received the news that 99% of the data was full of 0x00 characters (all zeros in the block-level data sense) and was essentially useless. With these findings, data recovery was deemed not possible.

DriveSavers shipped the drive back to the HCSO, where it was promptly placed into secure storage inside our Property and Evidence facility for safekeeping, where it has remained since that time.

Applicability to Wilkey Videos

The Wilkey in-car videos were preserved well before this data loss. At no time was there a question of whether the Wilkey videos were securely preserved and documented.

Almost a year prior to the final incident in July of 2019, the HCSO had identified the weaknesses inherent with the nature of the back-up methods used for this type of evidence and had begun the process of researching, selecting and negotiating the migration to a different system for the collection and storage of both in-car video and of body camera video. This was largely possible at this time due to the changes in the ability to store data securely in a cloud format. Not only has cloud storage become more secure, and therefore considered an appropriate manner of storage by various oversight agencies, but it has also become more affordable over the years as opposed to previously being cost prohibitive, as it was when the HCSO first chose the Arbitrator system for the collection and maintenance of in-car video.

There have also been a number of published misrepresentations with regard to the HCSO's production of the Wilkey videos. One example of the inaccurate media narrative is that the HCSO could not produce the videos in a timely manner, thus requiring the assistance and resolution implemented by the TBI "within 24 hours." At no point did the TBI participate in the pulling of videos. In fact, with the assistance of the vendor, the HCSO Network staff was able to utilize a mass export solution, so that an initial copy of all videos requested was delivered to the DA's office by the time that the TBI agent had arrived onsite at the HCSO. The TBI's only contribution to the project was supplying six (6) extra laptops to the District Attorney's Office, as well as to assist in getting the contents of the original removable drive copied to said laptops, in order to facilitate review by more than one individual at the DA's Office. These laptops were provided by TBI only because the HCSO did not have six extra laptops on-hand to loan to the DA's Office for an indefinite period of time.

Regarding the copying of videos, the Arbitrator software contained a feature that allowed HCSO Network personnel to set parameters for video to be exported based on the date range and badge number. HCSO Network personnel performed the mass export by inputting the date range requested by the DA's Office of January 1, 2019 through July 11, 2019, narrowed by Wilkey's badge number. The video footage pulled for the HCSO IA at Lt. David Sowder's request, was pulled on December 18, 2019, using the same procedure.

Budgetary Considerations

The HCSO has an annual budget of $59 million as of fiscal year 2019. From this budget, it pays salaries, operates a jail and provides oversight for the Silverdale Detention Center, and provides a superior level of service to our community with these taxpayer funds. Of that $59 million dollars, $1.35 million (approximately 2%) is allotted for capital expenses. Of that fractional amount, the HCSO purchases patrol cars, bulletproof vests, office supplies, cell phones, computers, software, and hardware. It further provides training to its employees and maintains facilities from those funds among other things. To say that budgeting for these needs is a challenge would be using charitable language at its most optimistic. Despite the budgeting challenges, the HCSO's commitment to excellence for our citizens does not waiver.