opentext™

OpenText Tableau Forensic TX1 Imager

User Guide

opentext™

Contents

1 Preface

This guide presents a wide range of technical information and procedures for using the Tableau Forensic TX1 Imager, a product of OpenText. It is divided into the following chapters:

- **Overview:** Provides general information about TX1 as well as unpacking, starting up, and navigating TX1 menus and reading the LEDs.
- **Configuring TX1:** Provides system overview information about TX1 as well as procedures for configuring and connecting it.
- **Using TX1:** Provides detailed information and procedures for TX1 operation.
- **Adapters:** Describes the adapters that extend the drive acquisition options and destination drive capabilities of TX1.
- **Specifications and troubleshooting:** Provides a brief list of potential problems and solutions. For more complete and current troubleshooting information as well as answers to frequently asked questions (FAQ), visitOpenText My Support (https://support.opentext.com/).

1.1 Drive capacity and transfer rate measurement conventions

The computer industry generally adheres to two different conventions for definitions of the terms megabyte (MB) and gigabyte (GB). For computer RAM, 1 MB is defined as 2^20 = 1,048,576 bytes and 1 GB is defined as 2^30 = 1,073,741,824 bytes. For drive storage, 1 MB is defined as 10^6 = 1,000,000 bytes and 1 GB is defined as 10^9 = 1,000,000,000 bytes. These two conventions are known as powers of two and powers of ten respectively. Microsoft deviates from the hard drive capacity measurement convention and uses the powers of two convention for its operating systems.

Tableau products report drive capacities and transfer rates according to the industry standard powers of ten convention. In TX1 screens, reports, and documentation, a 4 GB hard drive stores up to 4,000,000,000 bytes; a hard drive with a 150 MB/sec transfer rate transfers 150,000,000 bytes per second.

2 Overview

Tableau TX1 is a powerful, yet intuitive, forensic imager that offers superior local and networked imaging performance with no compromises. The touch screen user interface is easy to use and provides a familiar user experience similar to modern tablets and smartphones. TX1 is custom built for forensics and provides many standard and advanced features that serve the specialized needs of digital forensics and incident response practitioners, including:

- Acquisition of PCIe, USB 3.0, SATA, SAS, FireWire, IDE, and network shares (iSCSI and CIFS). **Note:** PCIe and IDE adapters (sold separately) are required to image these drive types.
- Output to USB 3.0, SATA, SAS, and network shares (iSCSI and CIFS).
- The ability to target file-based evidence with a powerful logical imaging function, including an intelligent, easy to use search engine with wildcard support and industry standard file outputs (lx01 and metadata csv files).
- The ability to save, import, and export logical imaging search criteria.
- The ability to automatically acquire any drives connected to the source ports on TX1 based on predefined job settings.
- The ability to enable and configure 802.1X network authentication to strengthen network access security.
- The ability to export locally attached media as an iSCSI share for remote, network-based acquisition.
- Support for cableless, toolless, SATA or SAS destination drive connections via the optional drive bay (TX1-S1), which also provides drive cooling.
- The ability to duplicate a source drive to up to four destination drives.
- Destination drives that can be a mix of directly connected drives and network shares.
- The ability to run multiple jobs in parallel of any type (clone, physical image, or logical image) to any available destination media.
- The ability to run up to two forensic jobs simultaneously at full speed, and to add additional jobs to the queue and reorder them.
- The ability to pause and resume imaging jobs, including resumption from power loss and certain types of job failures.
- The ability to prevent damage to disk drives by spinning them down when they are ejected from TX1 prior to physical removal.
- Browser based remote user interface to any number of network connected TX1s, with the ability to directly download selected files to the remote system.
- User management – create, delete, and manage user profiles.
- Superior network imaging performance through a 10 gigabit Ethernet interface, with auto-negotiation for use with slower 1 gigabit Ethernet networks.
- Superior data transfer rates, even while performing calculations of MD5, SHA-1, and SHA-256 hash values on two active jobs.
- Viewing extensive drive detail, including partition and filesystem information and raw hex data.
- Detection and notification of many popular encryption types (whole disk and volume based), RAID types, proprietary self-encrypting drives, and Apple device Core Storage volumes.
- The ability to unlock Opal-compliant self-encrypting drives (SEDs) and BitLocker-encrypted drives/partitions to enable unencrypted source media acquisition, and as an alternative to Tableau encryption for destination/accessory port media.
- The ability to unlock APFS-encrypted volumes to enable unencrypted source media acquisition (source ports only).
- Browsing drive filesystems, with the ability to view image and text files directly in TX1 user interface.

- Extensive filesystem support - APFS, ExFAT, NTFS, EXT4, FAT(12/16/32), and HFS+.
- Whole disk, open standard, destination drive encryption using XTS-AES.
- Automatic blank checking of source and destination drives.
- Comprehensive destination/accessory drive wiping capabilities, including NIST 800-88 compliant wipes and the ability to specify a custom wipe pattern.
- HPA, DCO, and AMA support for the detection and handling of hidden/protected data areas on source drives. This includes standalone HPA/DCO/AMA disablement, DCO/AMA “shelving”, and trim support for the creation of a destination DCO or AMA.
- The ability to update system time via an NTP server.
- Detailed forensic logs for case documentation in text and HTML formats.
- The ability to filter the forensic log list to only show logs of interest based on specific case and/or drive information. The filtered logs can also be exported or deleted.
- Always free firmware update support via Tableau Firmware Update (TFU) on a host system, via TX1 file browsing of any mounted filesystem (local drive or network share), or via the remote user interface.
- Clearly labeled and color-coded source (write blocked) and destination (read/write) ports.
- User interface localization support for German, English, Spanish French, Korean, Portuguese, Turkish, and Chinese languages, including virtual keyboard support for user inputs.

TX1 can operate as a standalone device, as shown above.

Optionally, TX1 can operate with the destination imaging bay (TX1-S1), as shown above.

The left source (write blocked) side of TX1.

The right destination (read/write) side of TX1.

2.1 TX1 kit contents

TX1 ships in a boxed kit with custom foam that includes the following items:

Item	Model #	Description
	TX1	TX1 Forensic Imager
	TX1-S1	Optional destination drive bay for up to two 2.5” or 3.5” SATA or SAS drives
	TP6	Provides power to TX1, the optional TX1-S1 Drive Bay, and most common combinations of source and destination drives. Uses a universal 3-prong style AC line cord and is compatible with 100-240V AC line voltages worldwide.
	TC4-8-R3	Unified SATA/SAS signal and power to 8” SATA/SAS signal and 8” 3M power cable (4 pieces).
	TC-PCIE-8	8” PCIe adapter cable. Must be used with a Tableau PCIe adapter.
	TPKG-VCT-5	Five piece Velcro cable tie kit
	TPKG-CLOTH	Microfiber screen cleaning cloth

Item	Model #	Description
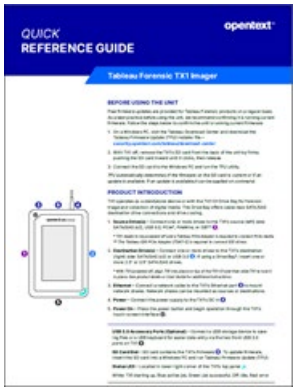		Quick Reference Guide

Do not discard the TX1 foam packaging, as it is designed to fit several industry-standard hard sided carrying cases (for example, the Pelican 1500). If you received the TX1 kit in the cardboard box shipped by OpenText, you can reuse the stacking foam inserts in your own hard sided case.

2.2 Navigating TX1

Use the TX1 touchscreen display to navigate from module to module and choose or modify options. Use the touchscreen keyboard or a USB keyboard to enter alphanumeric text when prompted. See “USB keyboard support” on page 5.

2.2.1 Home screen

The Home screen of TX1 displays icons for initiating the following functions:

- Duplicate
- Logical
- Verify
- Hash
- Browse
- Restore

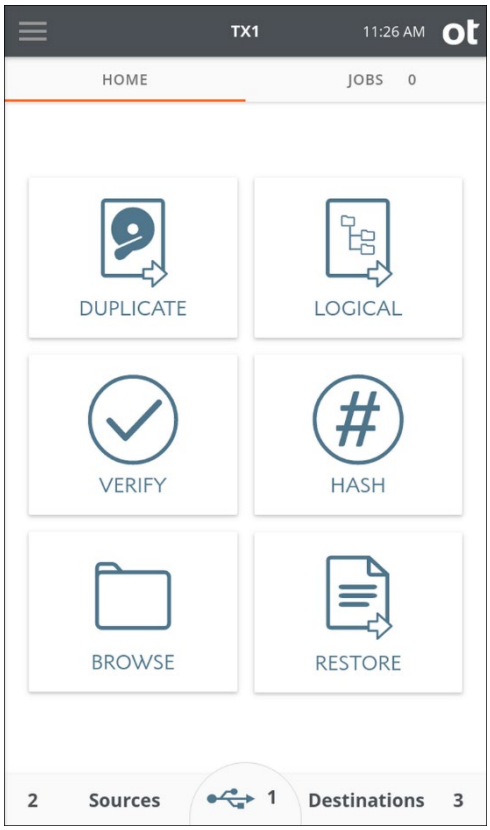

Tap one of these icons to begin a job and enter the job setup screen. A job setup screen provides a stepper-based flow from which you can view default settings, enter job notes for your case, change settings, and start the job. Tap the left arrow in the job setup tab to navigate back to the previous screen or to the **Home** screen.

Across the top navigation bar there are buttons to quickly access the side navigation menu , the **Home** screen, and view the current time. The TX1 model name in the top navigation bar links to the **Home** screen.

Note: In the event of abnormal cooling conditions, a warning triangle will be shown in the top navigation bar, to the right of the time. Such a warning will never be seen under normal operating conditions. Please refer to “Troubleshooting common problems” on page 175 for more information.

There are two tabs on the main screen: **Home** and **Jobs**. The **Home** tab shows the Home screen or one of the many other screens. The **Jobs** tab always shows the **Jobs** summary screen and the current number of total active and queued jobs.

The three buttons at the bottom of the **Home** screen are for **Sources**, **Accessory drives**, and **Destinations**. The number in each area represents the number of attached and detected drives (including mounted network shares). In addition to providing the drive

count, these buttons can be tapped to display a summary list of available drives and allow access to further drive details, screens, and operations.

Note: The middle area of the bottom row of the Home screen (for USB Accessory drives) is only shown when a USB Accessory drive is connected.

2.2.2 Side navigation menu

The menu icon in the upper left corner of the top navigation bar displays the side navigation panel that provides a menu of additional options and information. For additional information on the items in this menu, see “Configuring TX1” on page 18.

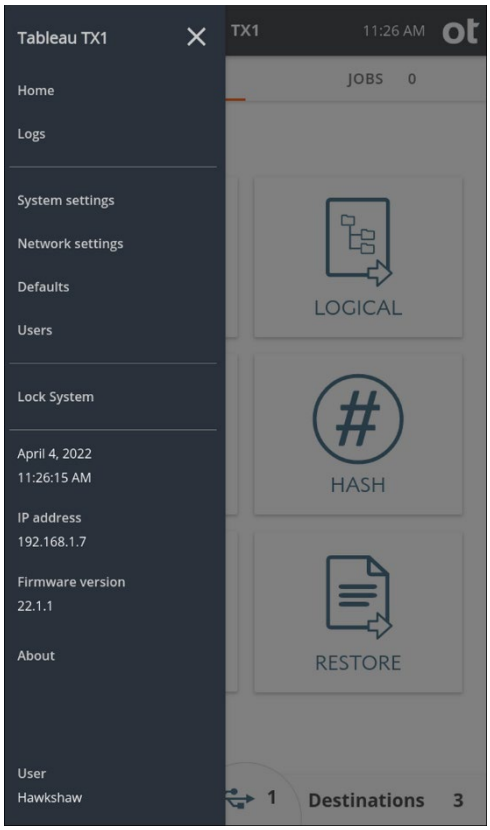

2.2.3 Jobs tab

The **Jobs** tab, which is found on the right side of the tab navigation bar, provides a convenient way to keep track of queued jobs, active jobs, and recent work done by TX1. It includes a counter in the top tab area. When TX1 is first powered on, the jobs counter is at 0. Once jobs are underway, each active or queued job will increment the counter.

The **Jobs** screen has four areas: **Automated Acquisition**, **Active Jobs**, **Queued Jobs**, and **Recent Jobs**.

The **Automated Acquisition** area will indicate if that mode is enabled. When enabled, the count of any duplication jobs run during that automated session will be displayed. See “Using automated acquisition” on page 74 for more details on this valuable, time-saving feature.

Any active jobs will be shown in the **Active Jobs** area. After two active jobs are running, additional job requests are placed into the queue. **Queued Jobs** start as soon as one of the two active jobs is complete, or when the source or destination resource the job is waiting for becomes available. Likewise, as jobs are completed, they are moved to the **Recent** area, so that a user can see what has already completed.

Jobs are queued in the order they were entered. You can reorder queued jobs by dragging individual jobs to place them in the order you prefer. On the TX1 screen, press and hold the drag icon ⠿ of the job you want to select, then drag the job to the desired position in the queue and release.

A fifth area, **Media Utilities**, displays media utility operations and only appears when a media utility operation is active. Media utility operations will move to the Recent area when they are complete. Media utilities cannot be queued like forensic jobs.

Here are two examples of the **Jobs** tab in action:

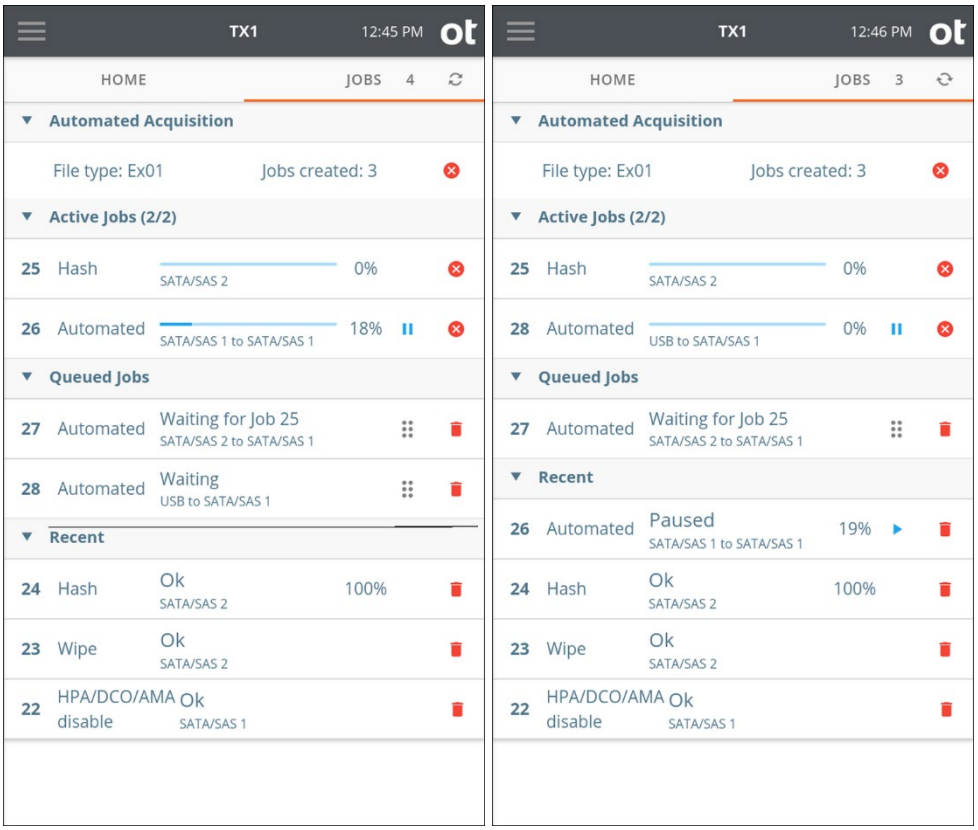

In the first example on the left, there are two active jobs – a hash and an automated duplication job. There are two automated duplication jobs queued and waiting for resources to be able to start. Note the textured grab areas in the queued job tiles which can be used for drag-and-drop job reordering. In the Recent area there are three jobs, each of which completed successfully. Since jobs pending includes both active and queued jobs, the **Jobs** tab counter reads 4.

In the second example, Job 26 has moved from the **Active Jobs** area down to **Recent** (due to being paused), which allowed Job 28 to start as shown in the **Active Jobs** area. See “Pausing and resuming a duplication job” on page 94 for more information regarding that feature.

There are several actions that can be taken on jobs from the **Jobs** tab, depending on both the state and type of the job. **Queued Jobs** and jobs shown in the **Recent** area can be deleted by tapping the **Delete** button on the right side of the job row. Jobs in the **Active Jobs** area can be canceled by tapping the **Cancel** button . Imaging jobs in the **Active** area that are capable of being paused and resumed (of type e01, ex01, dd, and dmg) can be paused by tapping on the **Pause** button , at which point the job will be moved to the **Recent** area with a status of Paused. Tapping the **Play** button or tapping the **Resume Job** button in the header of the **Job Status** screen will display a **Resume Duplication** screen. See See “Pausing and resuming a duplication job” on page 94 for more information about pause and resume.

2.2.4 Job status

After a job starts, its **Job Status** screen is automatically displayed. This status screen shows the details of a given job, including a header with its status, overall data rate, time remaining, and percent complete. The lower area of the **Job Status** screen shows additional job details, including sub-step progress (for example, **Duplication** separate from **Verification** in a duplication/verification job), a settings summary, and a listing of the drives involved in the job. Tapping a drive tile opens a drive details screen, which provides a quick view of all the information available for the drive. An example of an active **Job Status** screen is shown below.

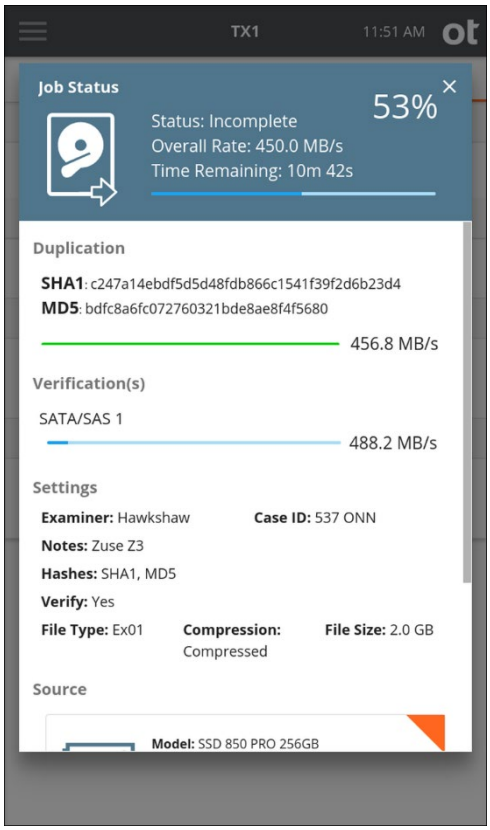

Note: The solid orange triangle in the upper right corner of the drive tiles indicates that the drive is currently being used as part of an active job. This is shown regardless of the location of the drive tile and makes it easy to spot drives that are in use. Similarly, drives that are part of a queued but not yet active job will show as an orange border triangle with white in the middle.

Once a job has completed, the **Job Status** screen is displayed and shows the final status of that job. Completed jobs have a link to the log file for that job on the right side of the header area at the top of the status screen.

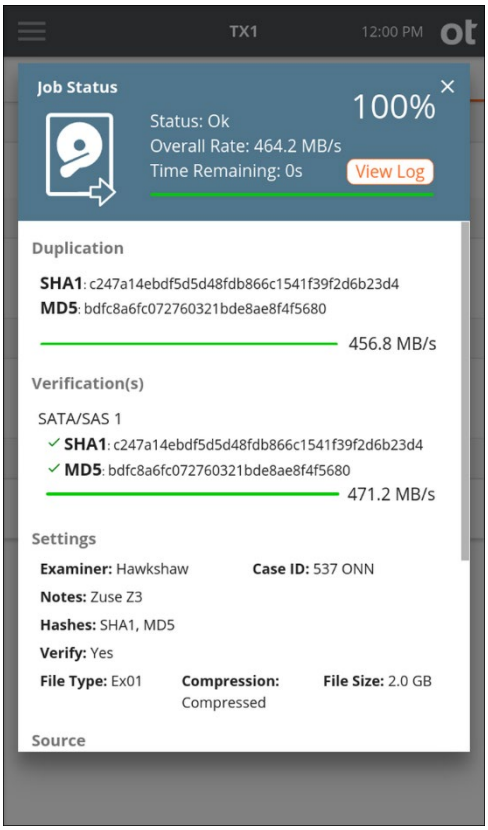

Clicking the **View Log** link immediately displays the detailed text log for that job. You can easily click back to the **Job Status** screen or close the log and return to the Home screen. The link back to the **Job Status** screen works regardless of the method used to get into the log details window. For example, if you had navigated to the log details for a given job through the side navigation menu **Logs** view, there will be a link at the bottom of the log to view the **Job Status** screen (**View Job**).

Note that the **Job Status** screen can also be viewed for completed jobs that are shown in the Recent area of the Jobs tab. Simply tap on the job row from the Recent jobs area and the final **Job Status** screen for that completed job is shown, including a link to the job log.

2.2.5 Quick Reference Guide

TX1 ships with a Quick Reference Guide that illustrates the drive connections, Status LED, power button, cable recommendations, and tips for getting started. Keep this card with TX1 as you familiarize yourself with its operation.

2.3 Reading the status LEDs

On/Off indicator LED: The illuminated power switch is located in the top-left corner of TX1 and it displays a white LED when the unit is on.

DC In LED: The TP6 power supply cable has a blue LED ring near the end of the barrel connector that indicates the TX1 power supply is receiving adequate DC input power.

Activity LED: The multi-color activity LED is located in the lower-right corner of TX1. It is white when the unit is booting up, blinking white when a power issue is detected, off when the unit is idle, blue when an operation is in progress, blinking green when an operation completes successfully, and blinking red when an operation fails.

Network Interface LED: The 10 gigabit Ethernet connector is located on the back of TX1 and it has two LEDs. The table below provides details for interpreting the status of these network interface LEDs.

Status	Green LED (Indicates link status and blinks for activity)	Green/Yellow LED (Indicates link speed)
No Link	Off	Off
100 Mbps Link	On / Blink	Off
1 Gbps Link	On / Blink	Yellow
10 Gbps Link	On / Blink	Green

2.4 Interpreting audio feedback

TX1 plays one of two sounds that indicate status at the end of a job. A chime sound plays for a successful job, and a buzzer sound plays for a failed job. You can change the volume of the sounds or disable them from the **System Settings** submenu in the side navigation menu.

2.5 On-screen warnings

TX1 will provide, as necessary, on-screen warnings within various settings and operations screens of the user interface. Yellow warnings call the user's attention to a potential risk, but do not impede operations. Red warnings mean either that a selected setting cannot be accommodated, an operation has failed, or a potential exists for forensic evidence to be missed, such as when a DCO or AMA is detected and not removed. Users are encouraged to pay attention to and read any displayed warnings when they appear and proceed accordingly.

2.6 USB keyboard support

You can plug a standard USB keyboard into either USB accessory port on the front of TX1. Some users find it more convenient to use an external keyboard to enter data instead of using the virtual touchscreen keyboard. Language localized keyboard support is limited to the virtual keyboard only.

Note: Wireless keyboards will also work with TX1. To use a wireless keyboard, simply plug the USB wireless adapter into either of TX1's front USB accessory ports, and it should automatically pair with the keyboard and start working. Note that there are many vendors of wireless keyboards and some may not be compatible with TX1. If you prefer to use a wireless keyboard and yours is not working with TX1, please contact Customer Support for keyboard recommendations.

2.7 Configuring TX1

This chapter describes the steps to configure TX1 prior to using it on a regular basis.

2.8 Startup sequence

When turned on, TX1 displays an initialization screen during the boot sequence. Once booted, TX1 displays the Home screen and then sequentially powers on connected drives.

3 Configuring TX1

This chapter describes the steps to configure TX1 prior to using it on a regular basis.

3.1 Startup sequence

When turned on, TX1 displays an initializatoin screen during the boot sequence. Once booted, TX1 displays the **Home** screen and then sequentially powers on connected drives.

3.2 Configuring TX1

TX1 default settings are defined using sensible, best-practice values. There are many options and settings you can configure and customize to your specific needs. Tap the Side Navigation menu icon in the upper left corner to access the following options:

- **Home:** Return to the **Home** screen.
- **Logs:** Access the forensic **Logs** screen.
- **System Settings:** Access the **System Settings** screen.

- **Network Settings:** Access the **Network Settings** screen.
- **Defaults:** Access the **Operations Defaults** screen.
- **Users:** Access the administrator level **User Management** screen.
- **Lock System:** Lock the screen with a PIN to prevent access while unattended.
- **About:** Access the **About** screen to view additional information such as the firmware build Id, serial number, network MAC address, copyright, and licensing information.
- **User:** Access the **User Management** screen for the current user.

3.2.1 System settings

Tap **System Settings** to display the **System Settings** page.

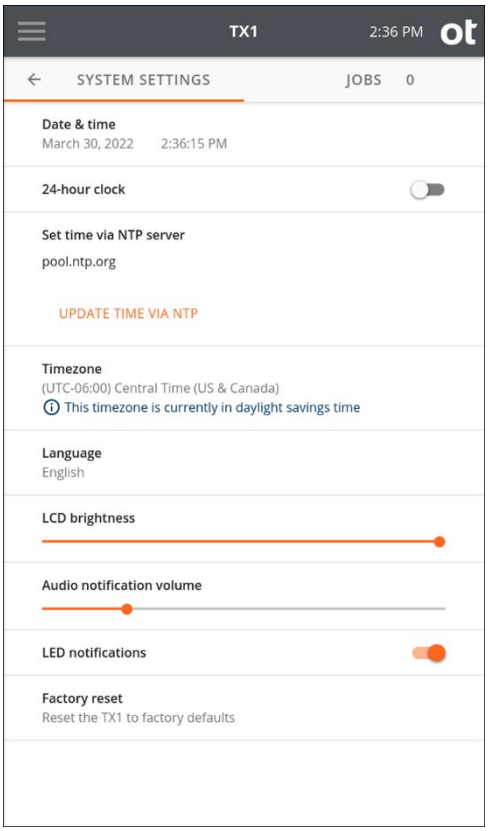

The **System Settings** page provides selections to configure basic TX1 system options including **Date & time**, **24-hour clock**, **Set time via NTP server**, **Timezone**, **Language**, **LCD brightness**, **Audio notification volume**, and **LED notifications**. You can also perform a factory reset.

Note: A factory reset will restore all system settings to their factory defaults and delete all existing job logs. Exporting all logs to an external device is recommended before initiating a factory reset.

Tap the toggle buttons to enable or disable a setting such as the **24-hour clock**. To define a slider setting value, such as the **LCD brightness**, tap and hold the slider selector, then slide to the desired value. For the **Timezone** setting, a multi-value selection box will be displayed to allow selection from a predefined list of settings.

Tap the setting row to reveal additional settings such as **Date & Time** (as shown in the screenshot below). Once the area expands, tap a setting value to reveal and select from a drop-down menu.

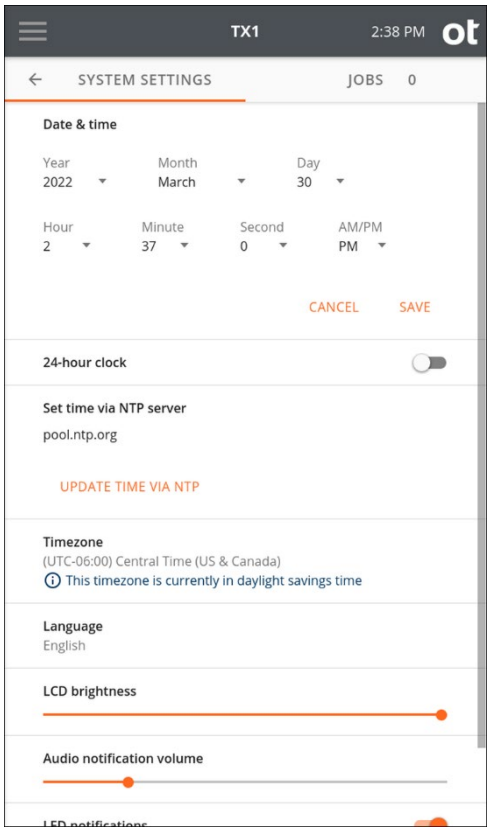

To enable setting the TX1 system time via an NTP server, a valid network connection and at least one NTP server source is required. The default NTP server is a public internet option (pool.ntp.org (http://pool.ntp.org/)). This can be changed to one of the other public internet options or a local network NTP server. Tapping on **Set time via NTP server** displays a list of options for NTP server connectivity. Note that the **UPDATE TIME VIA NTP** button is inactive if there is no network connectivity or no working NTP server available. Contact your local network administrator for assistance in setting up your NTP server.

Note: Due to the forensic implications of changing TX1's system time during an active job, the ability to use an NTP server to set the time has intentionally been limited to a manual call to the NTP server that is only available when no jobs are in progress. Also,

should the NTP server update routine fail for any reason, a warning message to that effect will be shown instructing you to manually set the time and date.

The user interface display language can be set to German, English, Spanish, French, Korean, Portuguese, Russian, Turkish, or Chinese. Changing the language automatically configures the virtual keyboard to match the new language selection.

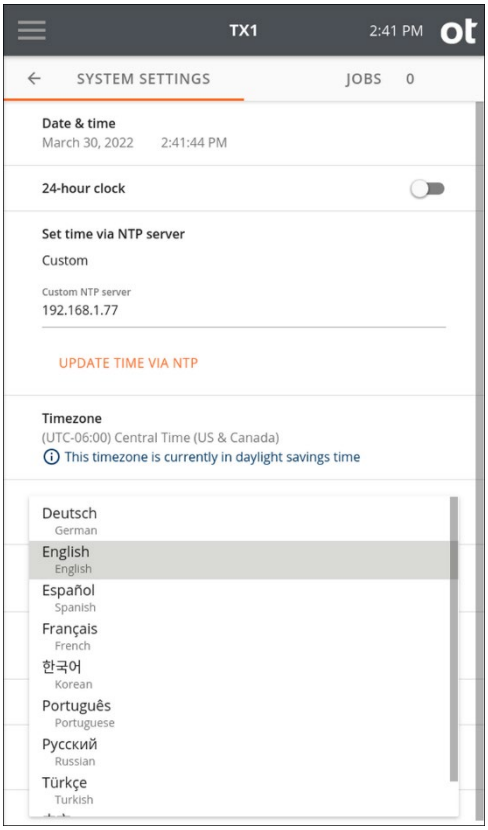

3.2.2 Network settings

Tap **Network Settings** to display the **Network Settings** page.

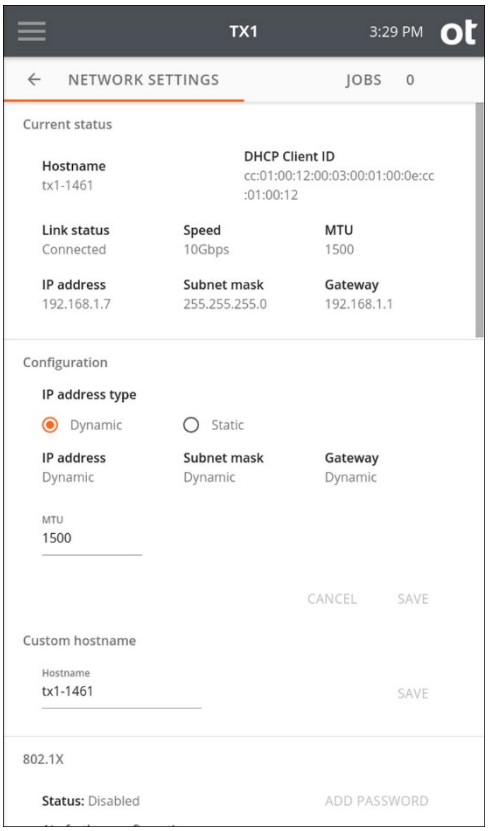

The **Network Settings** screen displays network related information and the current connection status in the top area, followed by a **Configuration** area for setting the IP address, MTU (maximum transmission unit) value, and custom hostname. Following the network configuration area are areas for 802.1X configuration, and an HTTPS certificate area.

Note that, for static IP address assignments, the DNS address and Domain fields can be entered to make mounting CIFS shares easier.

The default MTU value is 1500, but if your TX1 is attached to a network that supports jumbo frames, change the MTU value to 9000, which may enable much faster network transfer speeds.

Note: Maximum allowed MTU settings exist for the various standard network link speeds. Attempting to manually set the MTU in the **Configuration** area of this screen to a value that exceeds the maximum allowed link speed MTU will result in no change to the actual setting. This prevents slow network performance due to mismatched MTU settings.

The maximum allowed MTU settings for each link speed are as follows:

- 100 Mbps: 1,500
- 1 Gbps: 9,000
- 10 Gbps: 65,280

Each network device in the end-to-end communication path should use the same MTU value to achieve optimum and reliable performance. Contact your network administrator to verify the network configuration.

3.2.2.1 Configuring 802.1X network authentication

TX1 can be configured to connect to a network using IEEE 802.1X port-based authentication. This standard is designed to provide greater control over which physical devices are allowed on a given network, which greatly improves overall network security. A typical 802.1X network consists of an authentication server (RADIUS), an authenticator (LAN switch), and supplicants (network client devices).

To begin setting up your TX1 for use on a network with 802.1X authentication, tap the Edit button in the bottom right of the 802.1X settings area to choose one of the three EAP types (**TLS**, **TTLS**, or **PEAP**).

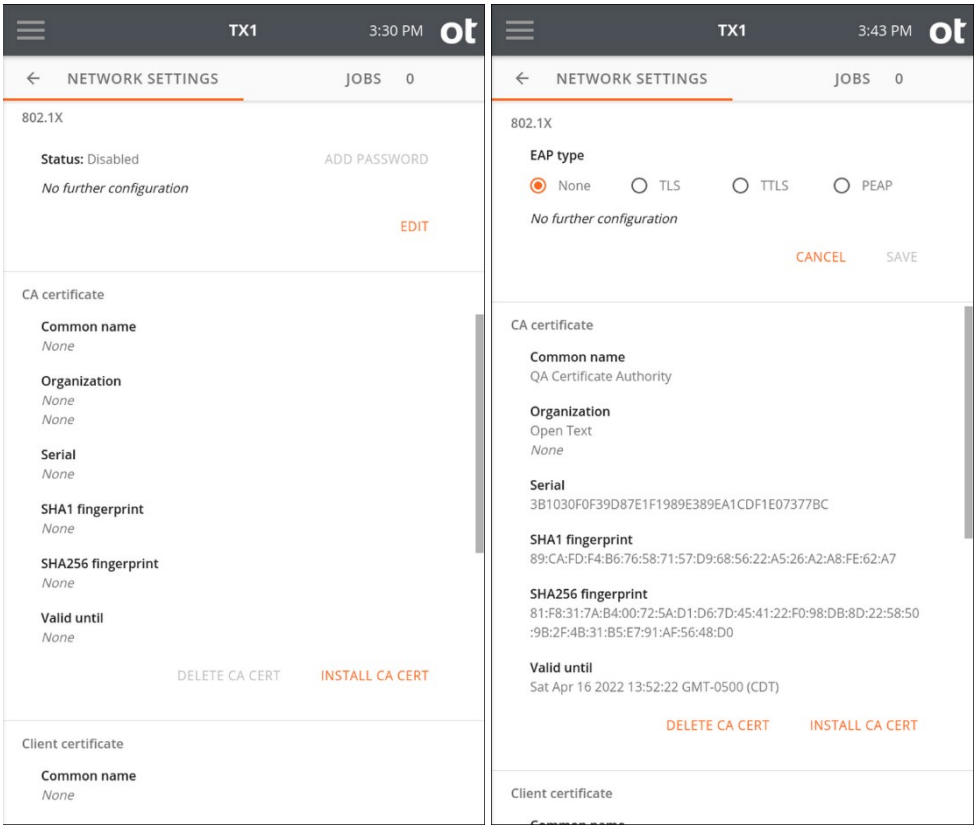

Then tap **Identity** to enter your 802.1X identity (required). Note that each EAP type has additional requirements and configuration settings depending on the type selected, as covered in the sections below.

Note that one or more certificates (depending on the EAP type and other settings) may need to be loaded onto your TX1 before attempting to authenticate on the network. The certificate loading process is straightforward. First, store the required certificates on a USB memory device, and then insert that device into a TX1 USB Accessory port. In the **CA** and/or **Client certificate** areas in the **Network Settings** screen, tap the appropriate certificate installation button (**Install CA Cert** or **Install Client Cert**). A browse window will appear, allowing navigation to the appropriate memory device and certificate file. Simply select the desired certificate file from the browser and tap the **Install** button.

TLS: TX1 and the authentication server authenticate each other by mutually verifying their certificates. A CA (Certificate Authority) certificate, and a client certificate, issued by the certification authority, must be installed before authenticating using this method.

Tap **Save** to show the selected EAP type and status in the settings summary.

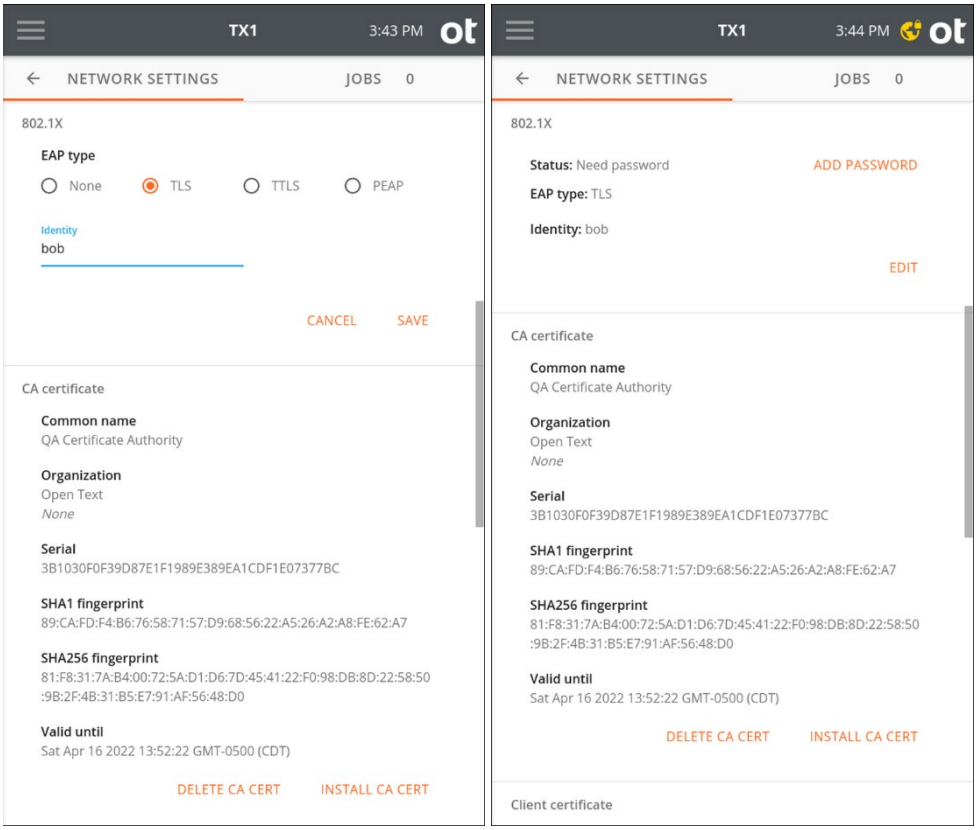

TTLS: Select a **Phase two** internal protocol (**EAP-MSCHAPv2**, **MSCHAPv2**, **MSCHAP**, **CHAP**, or **PAP**). A CA certificate must be installed on TX1 to enable server authentication. This method uses an identity and password for client authentication. A client certificate is not required.

Tap **Save** to show the selected EAP type and status in the settings summary.

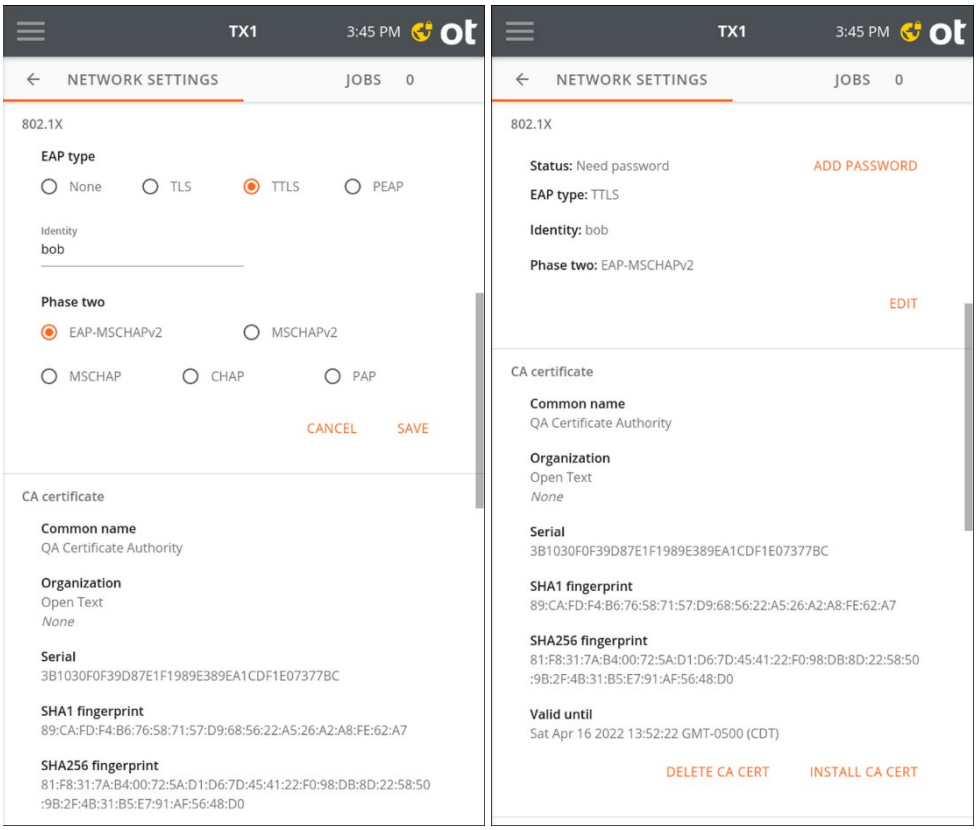

PEAP: Select a supported **Phase two** internal protocol (**EAP-MSCHAPv2** or **MSCHAPv2**). A CA certificate must be installed on TX1 to enable server authentication. This method uses an identity and password for client authentication. A client certificate is not required.

Tap **Save** to show the selected EAP type and status in the settings summary.

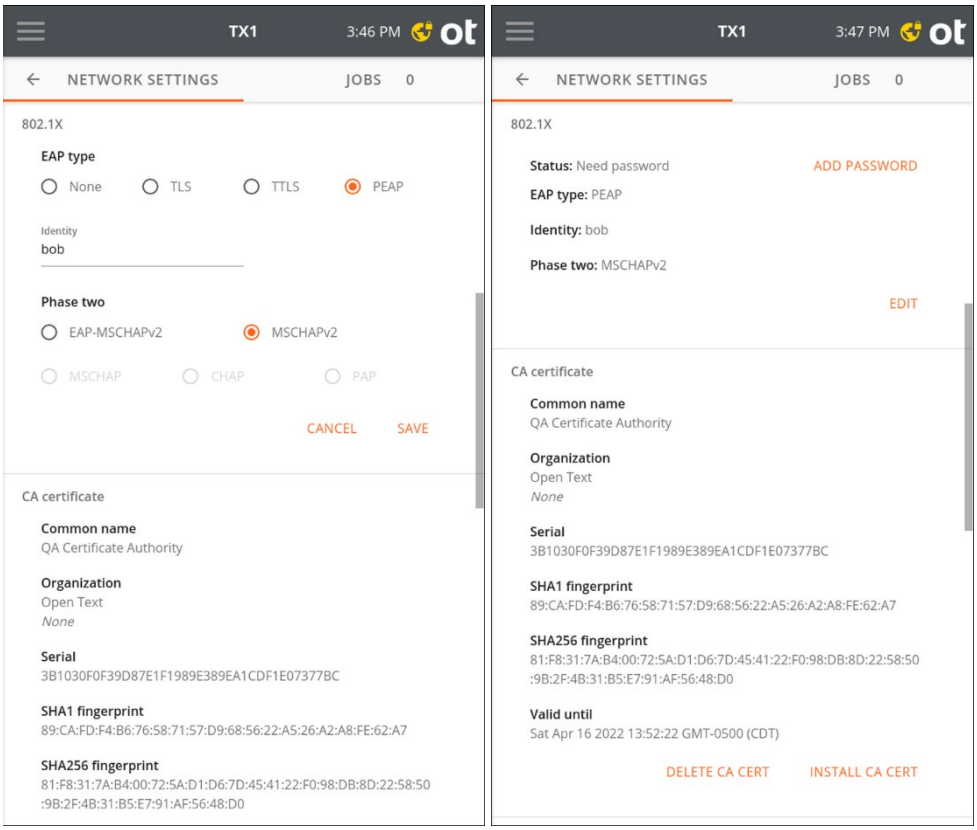

After saving the selected EAP type and **Phase two** internal protocol settings, a yellow icon appears in the right side of the top navigation bar, and an **Add Password** button becomes active in the settings summary area. Tap the navigation bar icon or Add Password to enter an 802.1X passphrase/password. Tap Submit to begin the authentication procedure.

Upon successful authentication the network lock icon will disappear and status in the settings summary will report Authenticated.

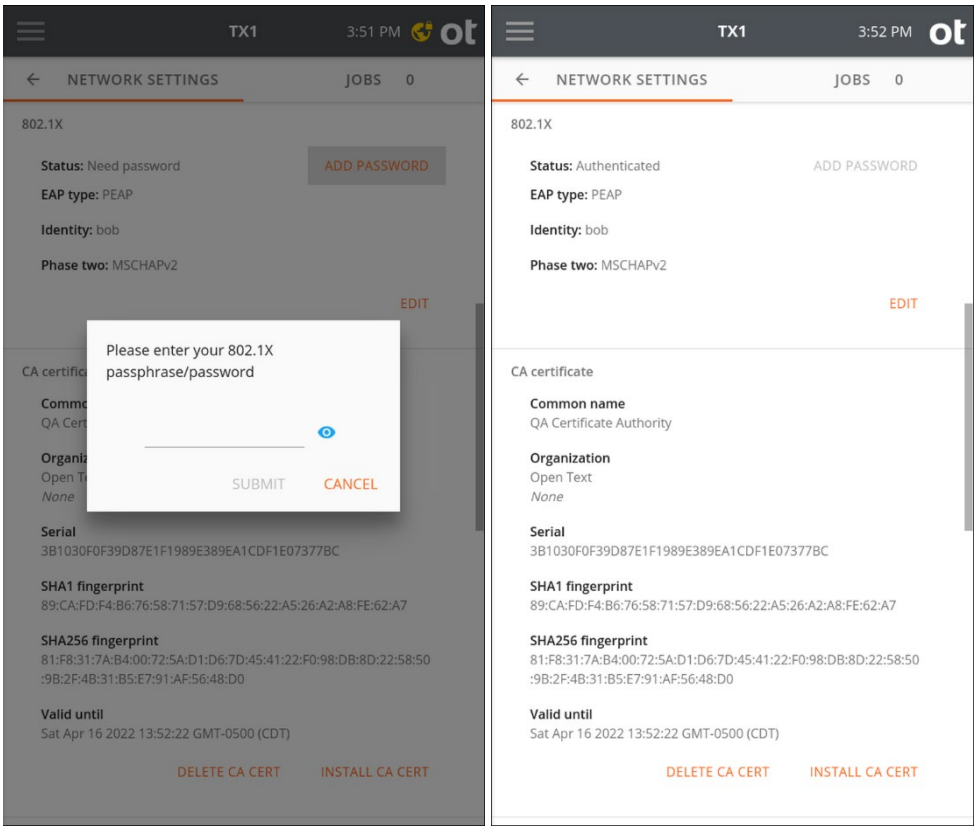

Note: 802.1X passphrases are required to decrypt encrypted private keys. These passphrases can be between 4 and 1023 characters in length.

3.2.2.2 HTTPS certificate setup

TX1 generates an SSL certificate on startup. You can use this certificate, manually generate a new certificate, or install your own certificate. The bottom area of the **Network Settings** screen shows the current SSL certificate information and provides options for manually generating a new TX1 certificate or installing a custom certificate. See “Remote web interface” on page 163 for more details regarding SSL certificate options.

3.2.3 Default settings

Tap **Defaults** to display the **Operations Defaults** page.

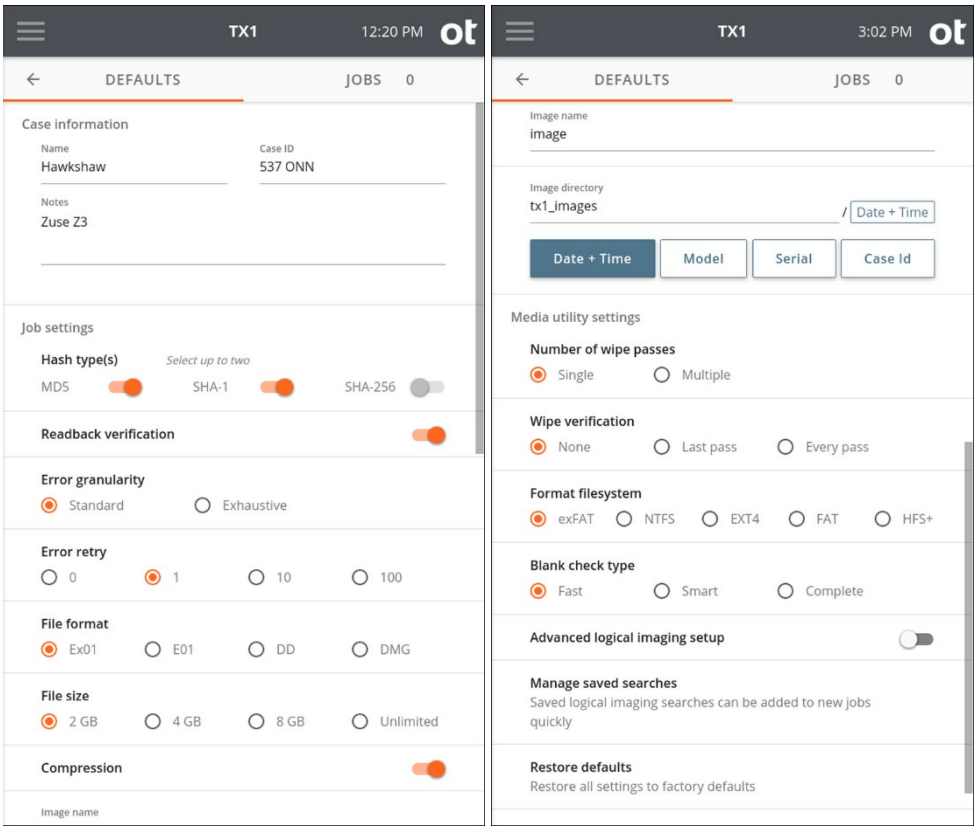

Several different entry methods are used for the settings on this page, including direct data entry, sliders, radio buttons, and an **Image Directory** name builder area. As shown above, the **Date + Time** directory path element boxes have been selected (in that order), so therefore the image directory path is `tx1_images/<Date+Time>`.

Note: The source drive model number, Serial number, and Case ID are not used in this example.

Tapping a selected directory path variable box deselects it. Changing the order of selection changes the order in which each element is incorporated into the image directory path.

Note: Certain combinations of image directory and image name settings can result in filename duplication on the destination filesystem for a given job. TX1 checks for this situation at job startup, and, if it detects a file name conflict, it automatically appends sequential numbers to the subsequent filenames, as shown in the file browse example below. Though this feature exists, it is recommended that the image directory and image name values be set such that no conflicts occur.

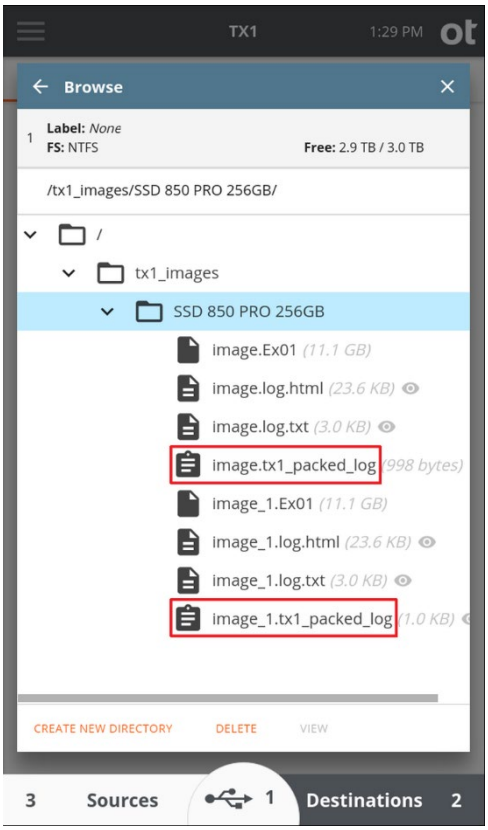

The **Advanced Logical Imaging Setup** switch sets the mode of operation for the **Logical Imaging** setup screens. The default value is off, which provides a basic and easy-to-use search method for targeting forensically valuable information on a given source drive. See “Logical imaging” on page 105 for more details on basic and advanced setup modes.

The **Manage Saved Searches** item allows for management of logical imaging searches that have been previously saved and the addition of new searches via direct entry or importation. Using saved searches makes lx01 job setup more straightforward and efficient. In this area, saved searches can also be exported to any mounted destination (including network-based filesystems) which allows for easy sharing of desirable/standard searches across all your TX1s. See “Logical imaging” on page 105 for more details on search setup and utilization.

Default settings can easily be restored to their factory set values. Simply select the **Defaults** item from the side navigation menu and then scroll to the bottom of the screen. Select the **Restore defaults** item and all settings will be immediately returned to their factory set values.

3.2.4 User management

In some forensic work environments, it may be desirable to set up distinct users with unique passwords to limit access to the available TX1 units. Also, with the addition of remote access capability, the ability to set user credentials has become a security requirement. TX1 ships with a default user named User1, which has administrator rights but no password and no remote access. User1 also has Auto Login enabled (which is only possible for users with no set password). These defaults make TX1 easy to use in environments that do not require user management or remote access.

If your work environment would benefit from setting up individual users with credentials, or to simply add remote access capability to any user (even default User1), tap the **Users** item from the side navigation menu. A list of users who are set up are displayed on the left tab. An example of the initial user list screen (with only default User1) is shown below. From this initial user list screen, a user with administrator rights can delete or modify any existing user and create new users.

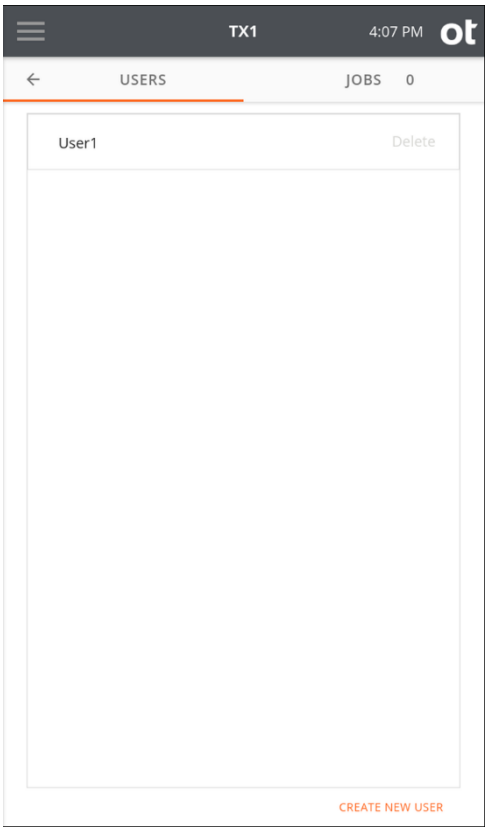

Note: Only users with administrator rights (including default User1) can access this **User Management** area of TX1. If a user is logged in who does not have administrator rights enabled, the Users button on the side navigation menu will be grayed out and unusable.

Tapping on any user in the list will show a **User Management** screen which contains all the available options for each user. The **User Management** screen for default User1 is shown below.

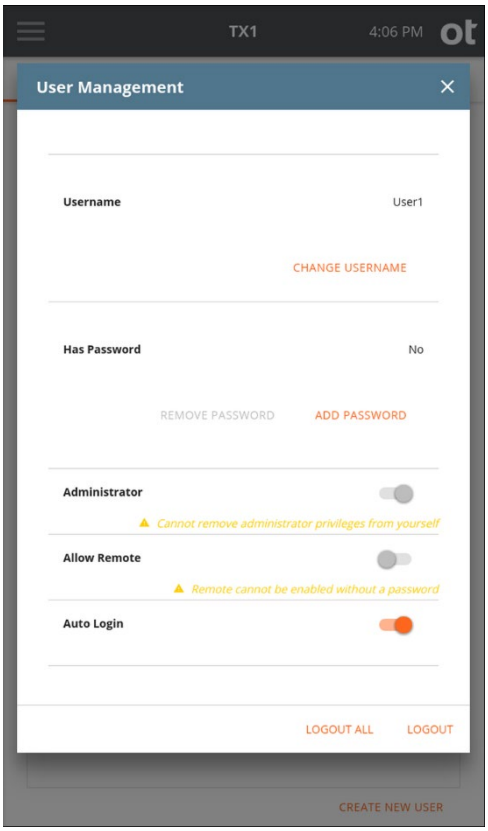

On this screen, an administrator can change any user's username and password, as well as set administrator rights, allow remote access, and enable auto login. Note that remote access cannot be enabled for any users that have no password setup, and auto login cannot be enabled for any that do have a password setup. Also, turning off administrator rights for a currently logged-in administrator is disallowed.

As you use the **User Management** screen to set up new and edit existing users, you may occasionally be prompted to reenter your password. This is a security feature that automatically disallows administrator level changes after 30 seconds has passed without any user interaction. Promptly going through all your user management activities and leaving this screen is recommended to avoid having to reenter your password. At the bottom of each **User Management** screen, either one or two buttons are displayed. All **User Management** screens will show the **Logout All** option. This will logout all instances of that particular user (local and/or remote logins). When the administrator is logged in and viewing their own **User Management** screen, (as shown above for default User1), a **Logout** button is also visible. This button logs out only that instance (local or remote) of that user.

To create a new user, simply tap the **Create New User** button at the bottom right of the **Users** list screen and enter the username and initial password for the new user and tap submit. A **User Management** screen will appear for the new user allowing complete user configuration (as shown in the screenshot above for User1).

Note that the currently logged-in local user is always shown at the bottom of the side navigation bar. Also, the username associated with the logged-in user is shown in the forensic log. For systems that do not have multiple users set up, the default User1 will be shown in the side navigation bar and in the forensic log.

A sample **Users** list screen with multiple users setup is shown below for reference. Should you have further questions while setting up users on your TX1 units, contact Customer Support.

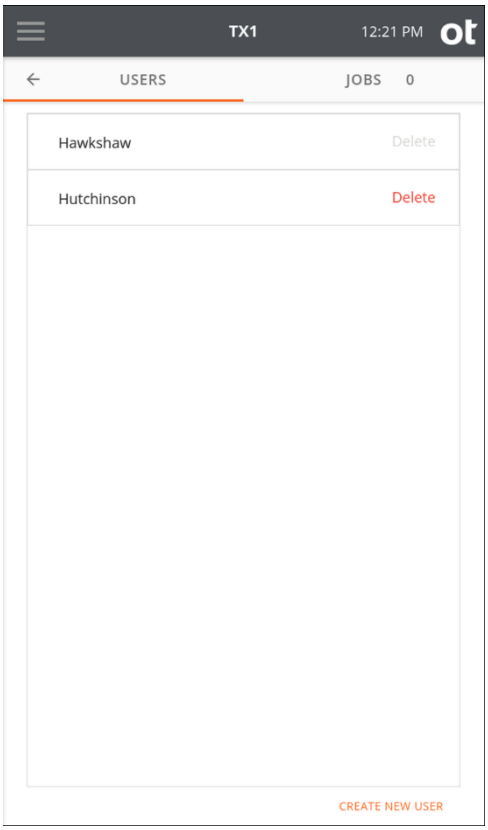

3.2.5 Locking the system

It may be desirable to lock your TX1 system while unattended to ensure no settings are changed or that your active/queued jobs are not altered in any way. To lock your system, simply tap on the Lock System item on the side navigation menu. A screen will appear that allows for entry of a six-digit personal identification number (PIN), as shown below.

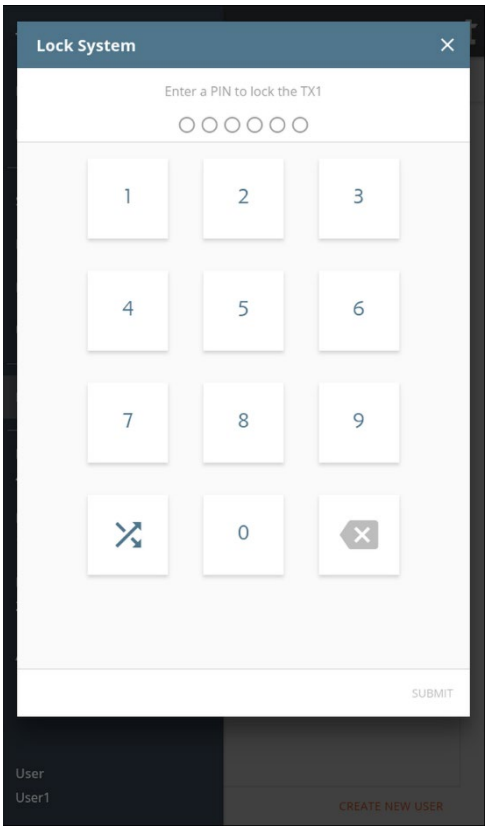

After entering the desired PIN, tap the **Submit** button in the lower right corner of the screen to begin the locking process. The system will prompt you for a second entry of the same PIN to confirm the desired digits have been entered. After verifying that both PINs match, the system will be locked. A message will appear at the top of the lock screen stating the time at which the system was locked.

To unlock the system, simply enter the current PIN and tap the **Submit** button in the lower right corner of the screen.

Note: The button at the bottom left of the keypad allows for randomizing the layout of the digits on the keypad. This can be used to ensure that commonly used PINs do not create a distinct pattern on the screen.

This PIN locking mechanism is temporary in the sense that a power cycle of TX1 will remove the lock.

3.2.6 Updating TX1 firmware

TX1 firmware is stored on an SD card located in the back of the unit. When a TX1 firmware update becomes available on the Tableau website, you can update your TX1 firmware a number of ways. You can use the **Tableau Firmware Update** (TFU) utility to update the SD card on any Windows-based computer system. Alternatively, TFU (starting

with version 7.32) can be used to save the TX1 firmware package file (.tx1_pkg) to any location available to your host system, which can then be made available to your TX1 to update its own SD card (without having to remove it from the unit). Finally, you can update the firmware remotely, using the web interface.

Note: A firmware update cannot be started while a job is running. This is true when initiating an update from the local user interface (after selecting the firmware package file) or after a remote user has uploaded a firmware package file to TX1. It is recommended that all active users on a given TX1 (local and remote) collaborate to ensure no jobs are running or will be started before a firmware update is initiated.

To update your TX1 firmware, select one of these options:

- **SD card removal update method** – To remove the SD card, turn the TX1 power off, then firmly push the SD card inward until it clicks, and release. The SD card will pop out. Gently remove the card and store it safely until it can be updated with TFU and returned to the unit. When re-inserting the SD card, be sure to gently push it in (label side up with angled edge toward the unit) until it clicks, then release. The SD card locks in position when it is properly inserted.
- **Package file update method** – Note that the SD card does not need to be removed for this option. Use TFU (version 7.32 and later) to save the required TX1 programming file (.tx1_pkg file) to the desired location (for example, an attached USB drive or to a network share that will later be accessible to TX1). On the TX1 that is being updated (using either the local or remote UI), make sure that the memory device (USB drive or network share) that contains the .tx1_pkg file is mounted and accessible to TX1. Then tap the **Firmware Version** item (or **About**) on the side navigation menu, and then the **Browse for Firmware** button, as shown in the screenshot below.

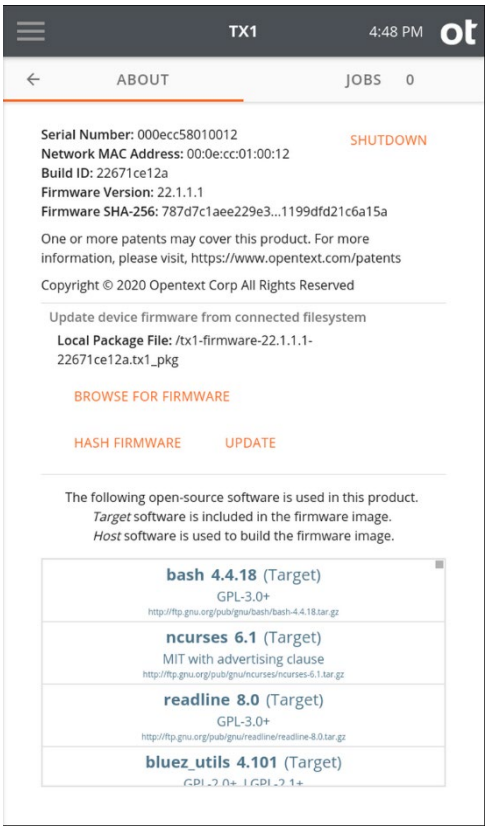

A browse screen appears that lists all mounted drives/shares. Tap on the desired drive/share and then use the **Browse** window to navigate to the folder that contains the TX1 firmware package file. Select the desired .tx1_pkg file, and then tap the Select button in the bottom right corner. The browse screen will automatically close, showing the name of the local package file that was selected. At that point, you may choose to hash the selected firmware package file or use it to update the unit's firmware. Note that tapping the **Update** button will automatically begin the update process, which takes over the TX1 unit allowing no user interactions until the unit successfully reboots itself after the firmware update.

Note: This package file update method performs an update of the SD card without reformatting it. Since the card was not reformatted, all previously stored user data (settings, logs, etc.) are retained. If a reformatted SD card is desired, use the TFU Prepare New TX1 SD Card feature to format and program your card.

- **Remote firmware update method** – Note that the SD card does NOT need to be removed for this option. On a remote browser window that is connected to the TX1 to be updated, open the side navigation menu, and select the **About** option. The remote **About** screen is shown below. Note that the **Upload device firmware** area is shown only in the remote user interface version of this **About** screen.

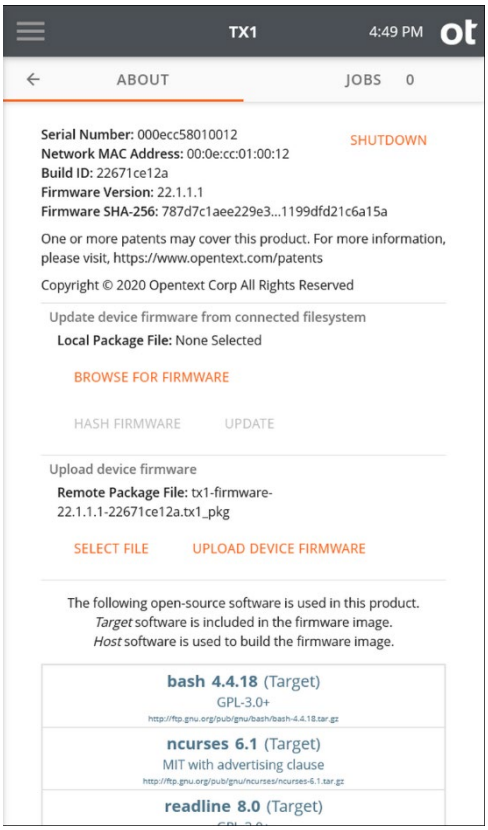

The first step in a remote firmware update is to select the new TX1 firmware package file on your local computer. To initiate this, tap the **SELECT FILE** button in the **Upload device firmware** area. This will launch a file browser window on your host system, allowing you to navigate to and select the desired firmware package file. The second step in this process is to upload the selected firmware package file to TX1. After checking the name of the selected package file in the **Upload device firmware** area, tap the **UPLOAD DEVICE FIRMWARE** button to initiate the firmware file upload. A progress bar will appear indicating that the file upload is in progress. Note that navigating away from this page when a firmware file upload is in progress will abort the firmware update process. Once the file has been fully received by TX1, it will automatically update its local firmware and reboot the unit.

Note: This remote firmware update method performs an update of the SD card without reformatting it. Since the card was not reformatted, all previously stored user data (settings, logs, etc.) are retained. If a reformatted SD card is desired, use the TFU **Prepare New TX1 SD Card** feature to format and program your card.

Note that, regardless of the firmware update method used, the hash of the currently loaded firmware package is calculated and displayed in the top portion of the About screen. This allows for at-a-glance verification that the proper firmware version is running and that it has not been altered.

3.3 Media utilities

Accessible from the Sources, USB Accessories, or Destinations buttons at the bottom of the Home screen (and all locations that provide drive lists), TX1 provides the following media utilities:

- Content breakdown
- Wipe (destination/accessory only)
- Format (destination/accessory only)
- Tableau encryption management (create/write on destination only; unlock on source/destination/accessory)
- Encryption unlock (source or destination; Opal, Bitlocker, APFS)
- HPA/DCO/AMA disable (ATA drives only)
- Blank check
- Browse filesystem
- SMART (ATA drives only)
- Export (iSCSI target)
- Eject

3.3.1 Content breakdown

The **Content Breakdown** utility is available for all mounted drives and offers a unique perspective on drive geometry that is quite intuitive and powerful. This screen provides a physical map style overview of the selected drive and all its partitions and filesystems. Tapping on each element (rectangular box) of the drive map provides basic information about that element and allows further exploration. Note that, on the right side of the screen, the selected element from the map is always shown at the top, but the parent element information is also shown below the selected element.

Drive and **Partition** elements include sector related information (sector size, sector count, and, for partitions, the starting sector) as well as a **View Hex** button which opens a window that shows the selected drive or partition hex data on a sector-by-sector basis. The screenshot below shows the drive **Content Breakdown** screen for a drive that has multiple partitions, with one of its partitions selected.

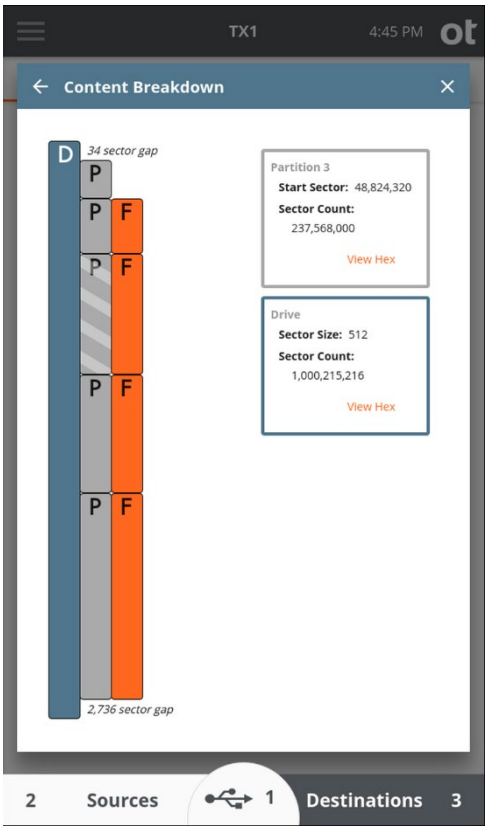

To view the raw hex data for a given drive/partition, tap the **View Hex** button within the information box. A sample hex view window is shown below. Each row shows 16 bytes of hex data, along with each byte's ASCII equivalent on the right side. The number with the plus sign to the left of each row represents the byte value offset from the start of the shown sector. The initial sector number (shown in the top blue header) is always the starting sector for the selected drive or partition. This starting sector number is also shown as the initial value in the rectangular box near the bottom of the screen (between the orange arrows). It is easy to navigate through the sectors within a given drive/partition and even into adjacent sectors by tapping the orange arrows. The single arrows will take you either one sector forward (right arrow) or one sector backward (left arrow). The double arrows will jump you to the next (or previous) boundary. That is to say, if you are in the middle of a given partition, the double-right arrow will jump to the end of that partition. Tapping the double-right arrow again in that scenario will take you to the beginning of the next partition or gap area (whichever happens to be there). These double-arrows make it easy to peruse an entire drive and see the very beginning and end of each drive element (partition or gap area) without having to go back to the map view to select a different element. Note that a specific sector number can also be entered in the box at the bottom of the screen. Simply tap the box, enter the desired sector number, and tap outside the entry field to go directly to that sector.

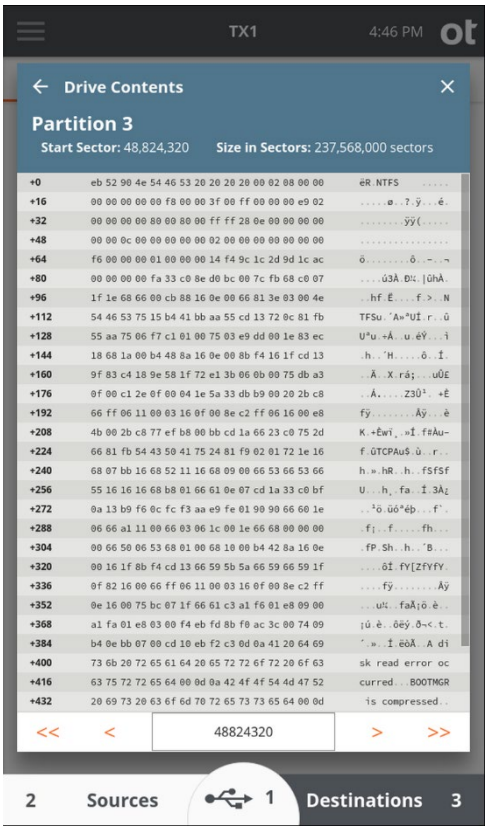

Filesystem elements show basic filesystem information (label, type, a free/used/total space) and a **Browse** button that opens TX1's standard browse modal, as shown in the screenshot below.

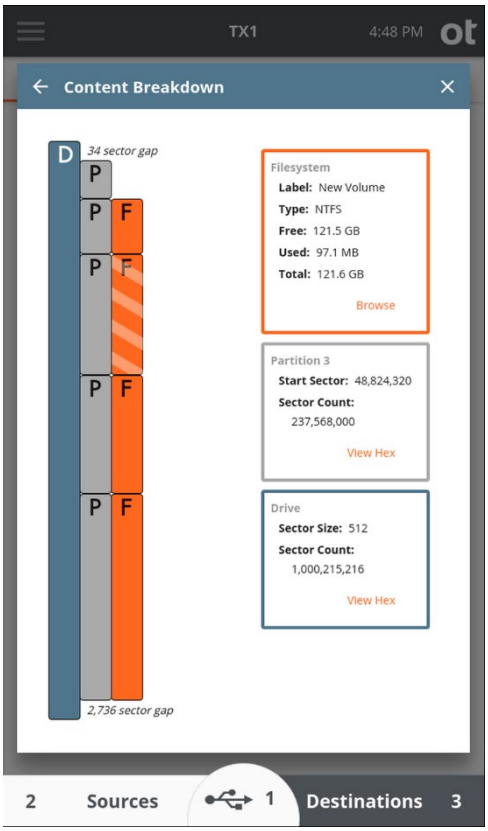

3.3.2 Wiping destination or accessory drives

The **Wipe** media utility provides three wipe types for destination and accessory drives. A summary of each wipe type is provided here, with more details in the table below.

- **Overwrite** – This method simply writes known pattern data to every accessible region of a drive. This can be done with a single pass or multiple passes. Verification is optional.
- **Secure Erase** – This method is only available for ATA based SSDs that support the command. TX1 need only issue the Secure Erase command to the drive, and the rest is done by the controller on the drive. Verification is not supported with this method since the state of the post-wipe data on the drive is drive manufacturer specific.
- **Sanitize** – This method is available for ATA and SCSI based media (both rotating and SSD) that support the command. Two wipe options are available for drives that support Sanitize – **Block Erase** and **Overwrite**. TX1 need only issue the Sanitize command to the drive, and the rest is done by the controller on the drive. Verification is optional. Note that an active Sanitize wipe may make a drive unresponsive for an extended period of time.

The exact differences between Secure Erase and Sanitize can be subtle, depending on the drive manufacturer's implementation. But, in general terms, Secure Erase is adequate for environments that are not concerned with removing any evidence of previous data in the physical memory chips. Secure Erase will guarantee that a typical host system read will return only wiped data, but someone with advanced capabilities to do chip-off memory structure analysis could theoretically discern previous data bit states. Sanitize is meant to cover situations that demand more secure data removal where advanced data retrieval techniques are of concern, with the downside of it taking much longer to complete.

Note: Secure Erase and Sanitize command requirements do not guarantee the final state of the data on wiped drives, which can result in wipe job failures that are out of TX1's control. From OpenText empirical testing over a large sample size of drives from different manufacturers, Secure Erase will reliably wipe drives in a very short period of time, but with a higher likelihood of a non-deterministic data state when complete, which makes reliable verification impossible. Sanitize has proven to be more reliable in clearing all data to zeros, which enables support of post-wipe verification. If you experience Sanitize wipe verification failures, please contact OpenText Customer Support to report the specific make and model of the drive, and the Tableau team will investigate.

The screenshot below shows the wipe screen for an SSD that supports Secure Erase.

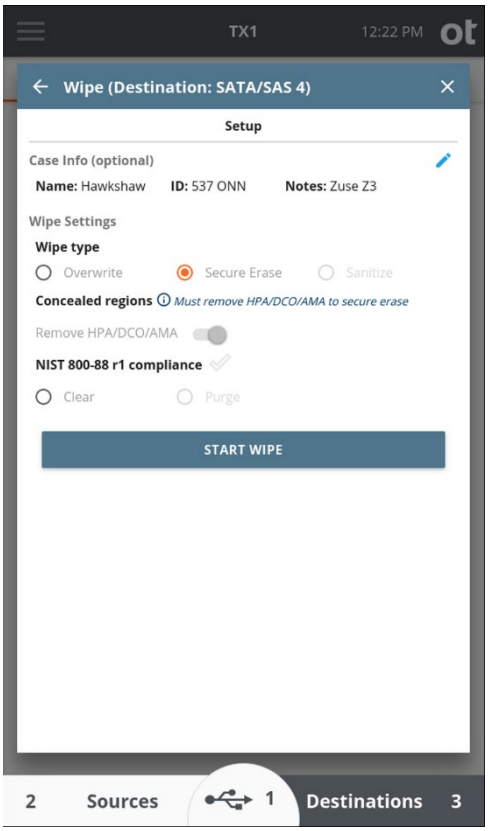

The screenshot below shows the wipe screen for an SSD that supports Sanitize.

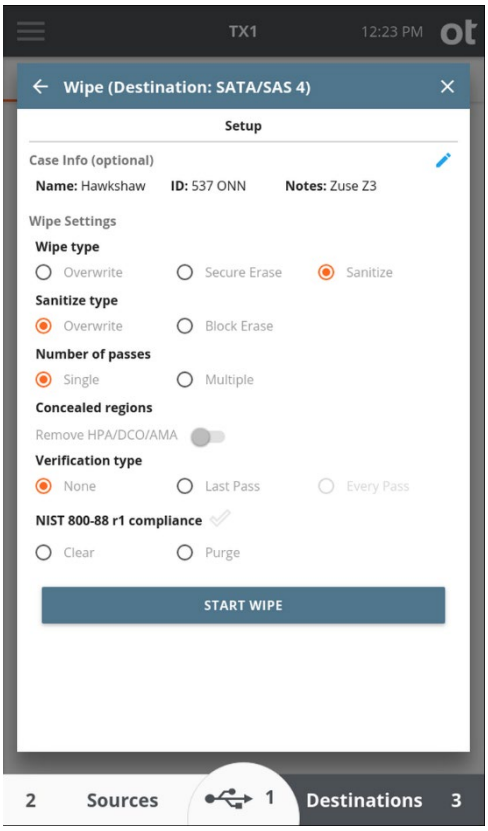

Note: Wiping drives results in sustained writing of the media, which can create abnormally high thermal operating conditions inside the drive. OpenText highly recommends using the TX1-S1 drive bay (which has active cooling) or an external drive cooler or fan when wiping media on TX1 to help prevent thermal damage to drives.

The following table provides **Wipe** option details:

Option	Description
Overwrite - Single Pass	TX1 will write a constant pattern to the destination or accessory drive in a single pass. If a custom wipe pattern is set, it will be written to the drive. Otherwise, zeros will be written. Verification is optional. Note that when an HPA/DCO/AMA configuration is present on a drive, a toggle may be set when configuring the wipe job to remove such configuration prior to starting the wipe, which will ensure the entire accessible drive space is overwritten.

Option	Description
Overwrite - Multiple Pass	TX1 performs three full write passes to the destination or accessory drive. The first pass writes zeros (0x0000) and the second pass writes ones (0xFFFF). When a custom data pattern is specified, it will be written only on the third pass. Otherwise, the third pass writes a randomly selected constant value between 0x0001 and 0xFFFE. Verification is optional. If enabled, it can be configured to verify after each wipe pass or after only the last pass. Note that when an HPA/DCO/AMA configuration is present on a drive, a toggle may be set to remove such configuration prior to starting the wipe, which will ensure the entire accessible drive space is overwritten.
Secure Erase (SSD only)	The ATA Secure Erase command instructs the drive to reset all available blocks to the erase state. How the erase state is implemented on the drive is not mandated by the ATA specification, which means the final data state on drives is manufacturer dependent (and not necessarily all zeros). Due to the indeterminate nature of the post-wipe data state, TX1 does not offer verification for Secure Erase wipes. Due to known issues with inconsistent and unreliable Secure Erase support on rotating drives (HDDs), TX1 only supports this feature on SSDs. Note that Secure Erase will erase all accessible drive space, but it will not necessarily erase over-provisioned space or other space reserved by the drive's internal controller. TX1 will force removal of any detected HPA/DCO/AMA configurations prior to starting a Secure Erase wipe, except for USB connected media. TX1 cannot remove HPA/DCO/AMA configurations for USB connected media, which means Secure Erase is not supported in that situation. See the note at the bottom of this table regarding handling of HPA/DCO/AMA configurations for USB connected media.
Sanitize - Block Erase (SSD only)	The ATA and SCSI Sanitize – Block Erase commands instruct the drive to erase all flash memory blocks. This is typically done electrically, not through writing of data to the drive. While the state of post-wipe data is not mandated by the ATA/SCSI specifications, Sanitize – Block Erase typically leaves a drive in a cleared (all zeros) state, which allows for post-wipe verification. Note that Sanitize – Block Erase will erase all user accessible drive space as well as over-provisioned space

Option	Description
	and any other space reserved by the drive's internal controller. TX1 will force removal of any detected HPA/DCO/AMA configurations prior to starting a Sanitize – Block Erase wipe, except for USB connected media. TX1 cannot remove HPA/DCO/AMA configurations for USB connected media, which means Sanitize – Block Erase is not supported in that situation. See the note at the bottom of this table regarding handling of HPA/DCO/AMA configurations for USB connected media.
Sanitize – Overwrite	The ATA and SCSI Sanitize – Overwrite command instructs the drive to overwrite all drive data in both storage and on-drive cache with zeros. This feature is typically implemented on HDDs but is available on some SSDs. Note that, for SSDs that support Sanitize – Overwrite, in addition to all user-accessible drive space, over-provisioned space and other space reserved by the drive's internal controller will also be erased. TX1 will force removal of any detected HPA/DCO/AMA configurations prior to starting a Sanitize – Overwrite wipe, except for USB connected media. TX1 cannot remove HPA/DCO/AMA configurations for USB connected media, which means Sanitize - Overwrite is not supported in that situation. See the note at the bottom of this table regarding handling of HPA/DCO/AMA configurations for USB connected media.

Note: Due to complexities associated with power cycling USB connected drives, HPA/DCO/AMA removal is not supported for such media on TX1 at this time, which precludes the use of Secure Erase and Sanitize wipe types. The presence of such configurations is detected and reported in the user interface and forensic logs for USB connected drives (as with all other drive interface types), but they cannot be removed directly by TX1 when connected to a USB port.

3.3.2.1 NIST 800-88r1 Compliance

In their document SP 800-88 r1: Guidelines for Media Sanitization, the National Institute of Standards and Technology (NIST) describes two data sanitization methods – **Clear** and **Purge**. TX1 indicates compliance with these data sanitization methods when the appropriate wipe settings are selected – the compliance checkmark turns from a dull gray outline to bold green, and the appropriate compliance radio button becomes selected.

Conversely, selecting either the **Clear** or **Purge** radio button, when enabled, will automatically select compliant wipe settings. TX1 can purge drives that support

conformant Sanitize commands. The **Purge** option is disabled if the Sanitize command is not supported by the drive.

The screenshots below show examples of NIST **Clear** and NIST **Purge** compliant drive wipe setups.

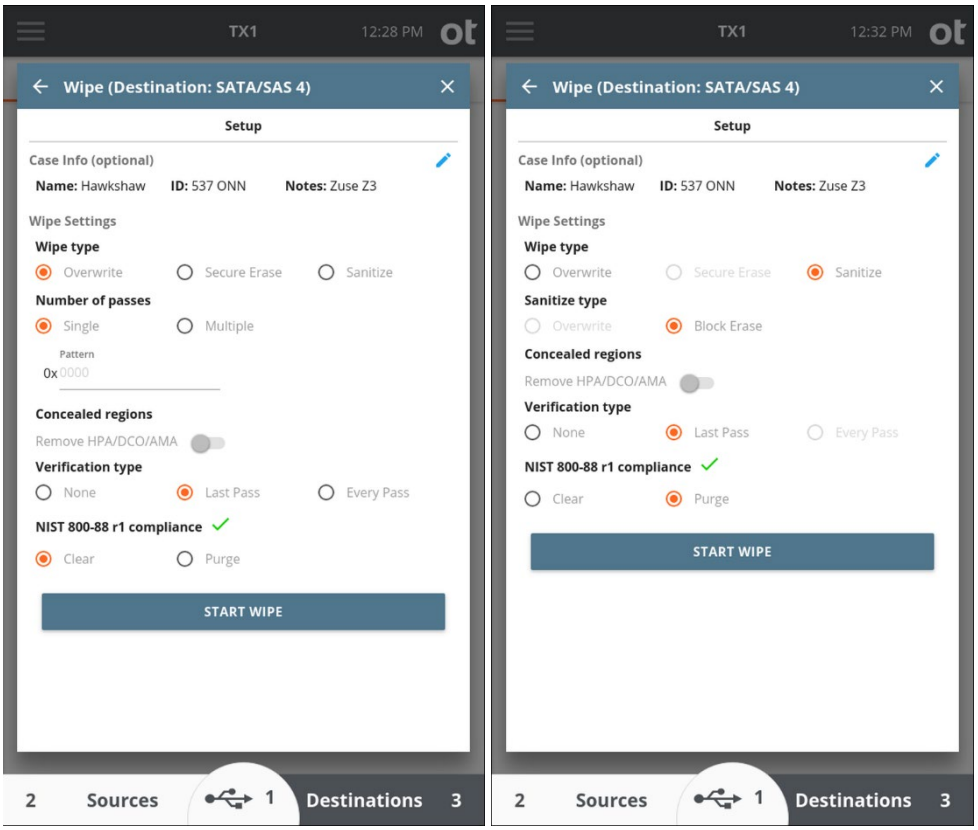

The following flowchart depicts how TX1 determines wipe setting conformity to NIST 800-88 r1. Note that this flowchart reflects handling of drives with no HPA/DCO/AMA settings.

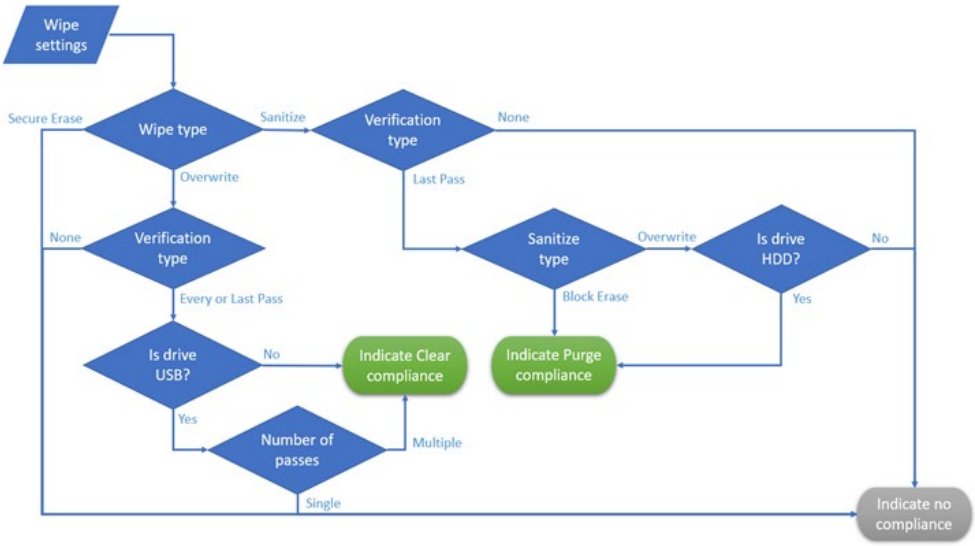

3.3.3 Formatting destination and accessory drives

To perform an image duplication to or save logs to a drive, you must format the destination or accessory drive with a file system that is recognizable by TX1. TX1 supports formatting destination drives in the following file system formats: **exFAT**, **NTFS**, **EXT4**, **FAT32**, or **HFS+**.

exFAT is recommended for best compatibility when accessing drives with all modern operation systems. **EXT4** is recommended for use with Linux forensic tools. HFS+ is recommended for use with MacOS forensic tools.

Note: When FAT is selected as the filesystem type for a destination drive format, TX1 will format the drive as FAT32. However, job logs (including the format log) and all user interface elements will simply show this as FAT. That is because TX1 supports reading from all FAT formats (12, 16, and 32), and simply identifying them all as FAT is considered acceptable and accurate for filesystem identification purposes.

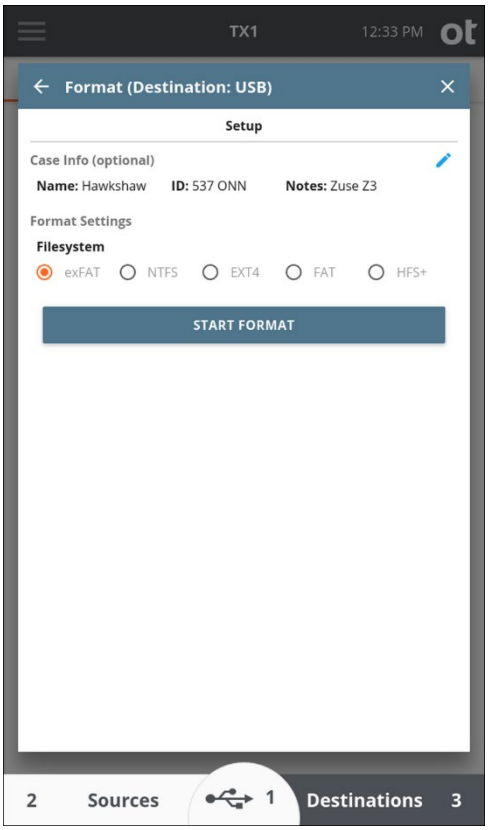

Note: TX1 cannot format a destination drive with an APFS filesystem, though it can mount a previously formatted APFS volume on any connected drive (source, destination, or accessory port).

3.3.4 Tableau encryption management

TX1 can encrypt destination and accessory drives using a password-based XTS-AES whole disk encryption. This Tableau-based encryption is compatible with the Tableau TD2u Forensic Duplicator and the open source VeraCrypt utility.

Note: The encryption process overwrites the destination drive, so remember to encrypt the destination drive before you copy data to it.

A Tableau-encrypted source drive can be unlocked with the password to allow browsing or imaging of the contents of the encrypted container.

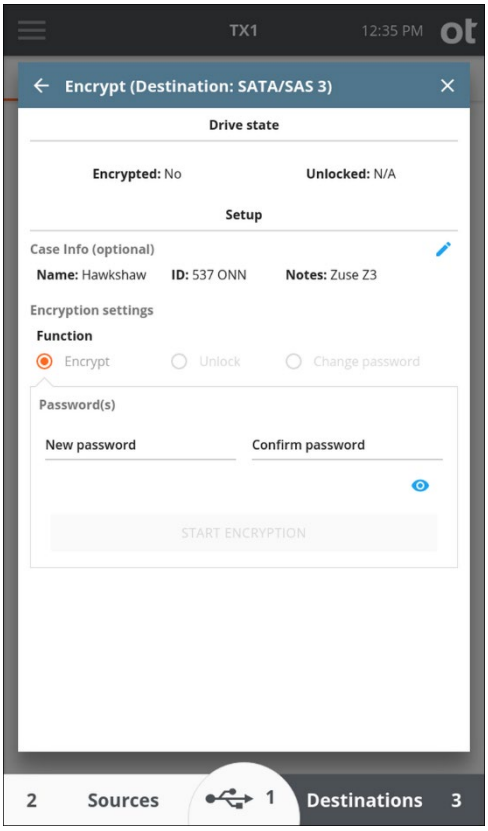

A Tableau-encrypted destination or accessory drive can be unlocked with the password to allow browsing or imaging to the encrypted container.

The encryption password can be changed only on destination and accessory drives, as it requires a write modification to the encrypted header.

OpenText is not able to recover lost passwords for TX1 encrypted media, so take appropriate steps to ensure you never lose your password.

To remove encryption from a drive, connect the drive to TX1 as a destination or accessory and wipe the drive.

3.3.5 Encryption unlock

TX1 can unlock drives/volumes that have been encrypted with APFS, BitLocker, and Opal. APFS encryption support is limited to source drives, but BitLocker and Opal encrypted media can be unlocked regardless of which port they are connected to. Note that this section is specific to non-Tableau encryption, mainly because none of these encryption types can be formatted directly on TX1. See the “Tableau encryption management” media utility on page 47 for details on usage of that method of encryption.

Whether dealing with APFS, BitLocker, or Opal encryption on a given drive/volume, the same **Encryption Unlock** media utility is used to unlock it. A pulldown field at the top of the **Encryption Unlock** screen lists all the detected encrypted types on a given drive, whether they be at the whole disk or volume level. Simply select the encrypted entity you want to unlock, enter the password (or BitLocker recovery key), and then tap the **Unlock** button to begin the encryption unlocking process. If successful, the progress bar will turn green. Close the **Encryption Unlock** screen to access the other TX1 functions with the now unlocked drive/volume. Once unlocked, each drive/volume can be used with any supported operations including browsing, imaging (physical or logical), and any applicable media utilities.

While unlocking APFS, BitLocker, and Opal encryption is simple and done using the same **Encryption Unlock** media utility, there are some notable differences in how TX1 handles these types of encryption that warrant special consideration, as covered in the sub-sections below.

3.3.5.1 Opal encryption

Opal Self Encrypting Drives (SEDs) that have had their encryption enabled in a Linux environment can be unlocked by TX1, as described in the beginning of Encryption unlock above. The presence of Opal encryption is noted in any area of the user interface that shows information about the attached drive, including drive tiles (which show in numerous locations), the Drive Details screen, and the Content Breakdown screen.

A locked Opal drive exposes no useful forensic information to TX1. The only options available for such media are ejection and unlocking. An unlocked Opal drive will appear as an unencrypted drive to the system and be usable for all supported forensic functions.

Note: The Opal standard does not specify an algorithm for generating a lock key from a plain text password. TX1 uses the Linux SEDUTIL function to report information about Opal drives and unlock them. This function uses an Opal-specific key generation algorithm as defined by the Trusted Computing Group. Other systems exist for enabling encryption on Opal drives (for example, BitLocker) which may employ a key derivation algorithm other than what the SEDUTIL function uses. Attempting to use a known password for such drives using TX1 will result in failed unlock attempts. Please contact OpenText Customer Support if you suspect you have run into such a situation.

3.3.5.2 BitLocker encryption

Drives and partitions that are encrypted with Microsoft BitLocker can be unlocked by TX1, as described in the beginning ofEncryption unlock above. The presence of BitLocker encryption is noted in any area of the user interface that shows information about the attached drive and/or partitions on the drive. This includes drive tiles (shown in the **Source** and **Destination** drive lists, among other locations), partition tiles (which show whenever a filesystem is being selected for an operation), the **Drive Details** screen, and the **Content Breakdown** screen.

Note: It is possible for BitLocker drives to have been originally encrypted and secured in a manner that TX1 will not be able to unlock/unencrypt. In particular, Smart Card and

Trusted Platform Module (TPM) methods secure a BitLocker encrypted drive with hardware-based interactions that are not supported by TX1.

Unlike an Opal SED, a BitLocker drive can be physically imaged (e01, ex01, dd, dmg) or cloned in its encrypted state. Such evidence can then be used with forensic investigation tools such as EnCase Forensic to unencrypt and analyze the evidence.

Once unlocked, the drive/partition can be used for any supported operations including browsing, logical imaging, and any applicable media utilities. Note that, while TX1 cannot format media as BitLocker, any previously formatted BitLocker drives/partitions can be unlocked and used as a destination for file-based operations such as writing image files and exporting logs.

Note: BitLocker encryption can be disabled, which is also known as **Clear Key** mode. While the data at rest remains encrypted in this mode, a password or recovery key is not required to unlock the encryption. On TX1, the method to unlock a disabled BitLocker drive/partition is the same as described above except that the **Password/Recovery Key** field will be disabled. TX1 will retrieve the **Clear Key** from the BitLocker metadata and use it to unlock the encrypted drive/partition.

When a drive/partition is BitLocker encrypted, it is assigned a Recovery ID number. TX1 will display the assigned BitLocker **Recovery ID** on the **Encryption Unlock** screen, which can help to identify specific drives that require a specific password/recovery key.

3.3.5.3 APFS encryption

Drive volumes that are encrypted with APFS can be unlocked by TX1, as described in the beginning of Encryption unlock above. The presence of an encrypted APFS volume is noted in any area of the user interface that shows information about the attached drive, including drive tiles (which show in numerous locations), the **Drive Details** screen, and the **Content Breakdown** screen.

Once unlocked, the APFS volume can be used for any supported operations including browsing, imaging (physical or logical), and any applicable media utilities.

It is important to note that there are distinct and critical differences in how TX1 handles the various encryption methods that can be unlocked on TX1 – APFS, BitLocker, Opal, and Tableau encryption. The table below summarizes how each of these encryption types will appear or be used in the identified TX1 operations in both their locked and unlocked states.

	APFS		Bitlocker		Opal		Tableau	
Operation	**Locked**	**Unlocked**	**Locked**	**Unlocked**	**Locked**	**Unlocked**	**Locked**	**Unlocked**
Logical Image (Source)	n/a (no file-systems to image from)	Selected files/ folders will be imaged	n/a (no file-systems to image from)	Selected files/ folders will be imaged	n/a (no reads possible)	Selected files/ folders will be imaged	n/a (no file-systems to image from)	Selected files/ folders will be imaged

	APFS		Bitlocker		Opal		Tableau	
Operation	**Locked**	**Unlocked**	**Locked**	**Unlocked**	**Locked**	**Unlocked**	**Locked**	**Unlocked**
Physical Image/ Clone (Source)	Full drive will be imaged/ cloned; encrypted state	Full drive will be imaged/ cloned; encrypted state	Full drive will be imaged/ cloned; encrypted state	Full drive will be imaged/ cloned; encrypted state	n/a (no reads possible)	Full drive will be imaged/ cloned; unencrypt-ed state	Full drive will be imaged/ cloned; encrypted state	Only the unlocked encryption container contents will be imaged/ cloned; unencrypt-ed state
Wipe (Dest)	Clears drive starting at sector/ block 0	Clears drive starting at sector/ block 0	Clears drive starting at sector/ block 0	Clears drive starting at sector/ block 0	n/a (no writes possible)	Clears drive starting at sector/ block 0	Clears drive starting at sector/ block 0	Clears only the unlocked encryption container leaving encryption intact
Format (Dest)	n/a (no APFS support on destina-tion)	n/a (no APFS support on destination)	Will overwrite existing format-ting, including BitLocker	Will overwrite existing formatting, including BitLocker	n/a (no writes possible)	Formats drive starting at sector/ block 0	Not allowed	Formats unlocked encryption container only leaving encryption intact
Blank check (Source/ Dest)	Checks full drive starting at sector/ block 0	Checks full drive starting at sector/ block 0	Checks full drive starting at sector/ block 0	Checks full drive starting at sector/ block 0	n/a (no reads possible)	Checks full drive starting at sector/ block 0	Not allowed	Checks contents of unlocked encryption container only

3.3.6 Disabling drive capacity limiting configurations

In the past, the most common method of intentionally limiting the reported capacity of a drive was by using the ATA HPA (host protected area) or DCO (device configuration overlay) feature sets. Starting with the ACS-3 (ATA/ATAPI Command Set 3) specification update, the concept of Addressable Maximum Address (AMA) was introduced. Newer drives may support this method of limiting the reported drive capacity. TX1 supports all these methods with automated detection, identification, and notification that will make dealing with them seamless and easy. From a forensic point of view, it is valuable to know if HPA, DCO, or AMA are in use. With that knowledge, the forensic practitioner can make an informed decision about whether or not to acquire data in the hidden regions of the drive.

Note that these methods (HPA/DCO and AMA) are mutually exclusive. A drive that supports HPA/DCO will not support AMA, and a drive that supports AMA will not support HPA/DCO. Also, while HPA and DCO are related features for a given drive, HPA has a unique attribute (volatile, or temporary, removal) that distinguishes it from DCO and AMA. For that reason, this section will cover volatile HPA removal as a separate topic before addressing non-volatile (permanent) removal of HPA/DCO or AMA.

3.3.6.1 Volatile HPA removal

HPA can be disabled without making a permanent modification to the drive. This is known as volatile, or temporary, removal of the HPA configuration. When a drive that has had its HPA removed in this manner is removed from TX1 (or is otherwise powered down) and then re-powered, it will always come back in its original state (with the original HPA configured and enabled). Since this is a temporary drive configuration change only (not a change to the data stored on the drive), TX1 automatically disables HPA on any drive connected to one of its source ports. Since DCO and AMA settings can only be disabled on a permanent basis, TX1 does not automatically disable them on connected source drives.

In the case of an automatic, volatile HPA removal from a connected source drive, the TX1 user interface makes it obvious what has occurred, as shown in the following screenshots.

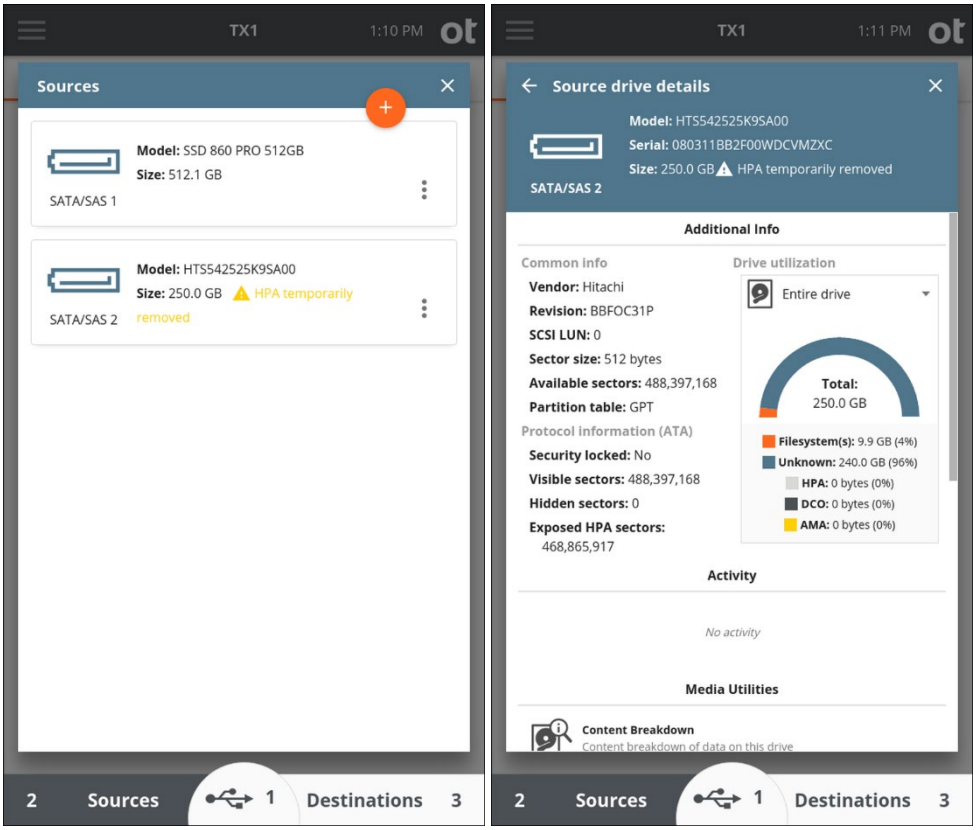

Referring to the drive details screenshot above, the fact that the HPA has been removed is reflected in the following areas on this screen:

- The **Size** field in the header reflects the full capacity of the drive (with HPA removed), along with a warning to draw attention to the HPA removal event.
- The drive utilization data on the right side accurately reflects the existence of 0 bytes of HPA (since the HPA was removed).

- The protocol information area shows the fact that no sectors are currently hidden, as well as how many were exposed when the HPA was removed.

TX1 never makes automatic changes to any drive capacity limiting configurations on destination drives. TX1 was designed to give the forensic practitioner complete control over the destination drive. If you choose to restrict the destination drive capacity using HPA, DCO, or AMA TX1 will not override that decision.

3.3.6.2 Non-volatile HPA/DCO/AMA removal

The HPA/DCO/AMA Disable media utility permanently disables the HPA, DCO, or AMA configurations on the source drive. Note that for HPA/DCO, you cannot remove a DCO-protected region on a drive without also removing any HPA-protected region, as defined by the ATA specification.

If a drive has a DCO or ACA configured, a red warning message is displayed on the drive tile indicating DCO or AMA is limiting the drive size (as shown in the DCO example below). Permanently disabling a DCO (and any HPA on that drive) or AMA is done from the media utilities portion of the drive details screen for a given drive. Drive details can be viewed through the **Sources** and **Destinations** areas of the main **Home** page or from the **Select a Source** or **Select a Destination** areas during duplication job setup.

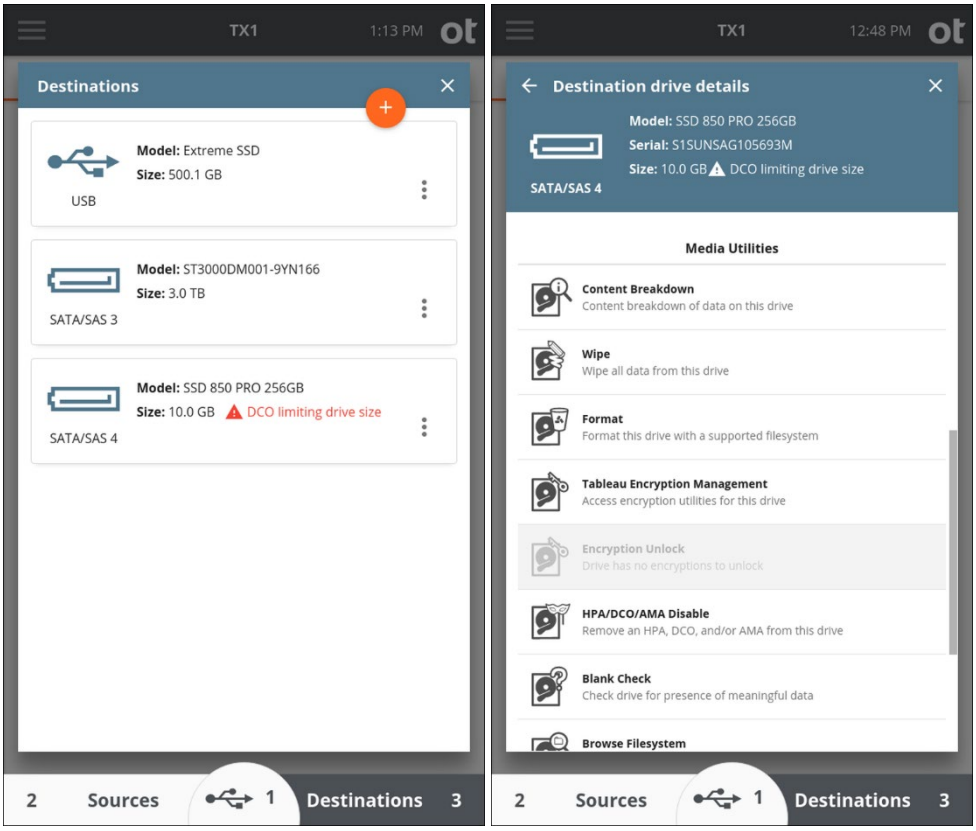

IDE drives with a DCO require special considerations with TX1. This is due to the fact that IDE drives are connected via the PCIe interface using a Tableau IDE Adapter (TDA7-5). DCO setting changes require power-cycling the drive which, for directly connected SATA drives, is done automatically by TX1. PCIe, as implemented on TX1, does not allow for hot drive removal (including power-cycling).

To disable a DCO on an IDE drive:

1. Confirm that TX1 is powered off.
2. Connect the IDE drive to a Tableau IDE Adapter (TDA7-5) and connect the adapter to the TX1 PCIe port. Note that IDE drive power does not come directly from the TDA7-5 adapter. A separate power cable is required to connect the IDE drive power interface to one of the source SATA/SAS drive power connections.
3. Connect any other drives required for the eventual duplication job and power on TX1.

   Note: TX1 will not allow any other jobs to be started or queued when an IDE drive HPA/DCO/AMA disable job is in progress. This includes disallowing the **Shelve DCO** feature for duplication jobs. This is because a system reset will be required to redetect the drive after the DCO is disabled, which would terminate any other jobs in progress.
4. After TX1 has booted to the **Home** screen and the IDE drive is detected, navigate to the **Sources** drive list and tap on the IDE drive tile to open its drive details screen.
5. Select the HPA/DCO/AMA **Disable** media utility, optionally enter case information, and then tap the **Remove HPA/DCO/AMA** button to start the operation. The system reset will occur immediately after TX1 issues the removal commands to the drive.

The log will reflect successful completion of the DCO disable operation after the command has been issued to the drive, before the system reset has occurred. TX1 has no way of knowing if the command actually completed at the drive level. The DCO state should be manually verified after the reboot is complete and before subsequent jobs are started.

TX1 also provides the ability to “shelve” a DCO or AMA, which means disabling a source drive DCO or AMA for the purposes of evidence duplication and then putting the same DCO/AMA back after the job is complete. See “Duplicating” on page 64 for more details on shelving a DCO.

3.3.7 Blank checking

The **Blank Check** utility checks a drive for the presence of meaningful data.

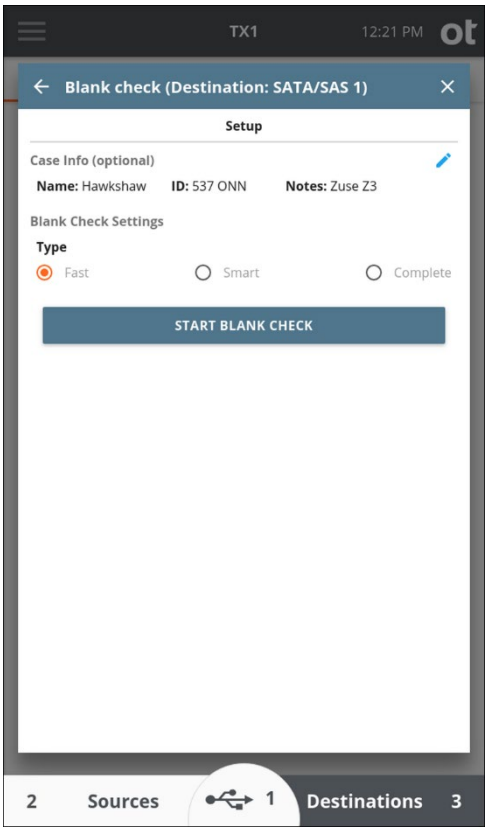

The following table provides **Blank Check** option details:

Option	Description
Fast	Quickly checks to determine if the drive appears to be blank by reading in and checking the sectors in the Master Boot Record, the Primary GPT, and the Secondary GPT.
Smart	Performs the Fast check, then reads in up to 75% of the available sectors randomly to determine whether they are blank. The blank check will stop as soon as a non-blank data pattern is detected.
Complete	Reads in up to 100% of the available sectors to check if the drive is blank. The blank check will stop as soon as a non-blank data pattern is detected.

A sector is considered blank if it contains only the same repeated 2-byte pattern. Any non-repeating pattern is considered to be non-blank. However, each individual sector

may contain different repeating patterns. If any sector is found to not be blank, the drive is not considered blank, and the blank check will stop.

Note: The **Fast** and **Smart** blank check options do not perform exhaustive checks of the entire drive. It is possible for a drive to appear to be blank according to the Fast or Smart check while still storing forensically relevant information.

3.3.8 Browse filesystem

The **Browse** function provides an easy way to view the contents of a recognized filesystem on any mounted drive, whether it is connected locally or via the network interface (iSCSI or CIFS). Simply tap the **Browse** button on the Home screen and select the desired drive/filesystem. The **Browse** operation is also accessible from the **Media Utilities** list in the drive details screen and from the filesystem details box within the **Content Breakdown** media utility. A sample **Browse** screen is shown below.

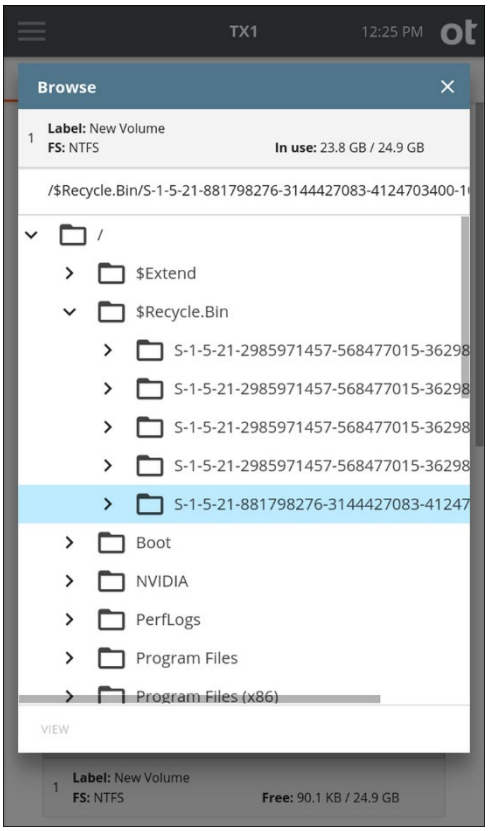

The **Browse** window provides basic information about the selected filesystem (label, type, and used space) in the first row. Note that, when multiple filesystems are present on a drive, tapping the top information row will allow for selection of a different filesystem on the same drive, without needing to back out of the **Browse** window to select the other filesystem. The second row of the **Browse** window shows the complete path name for

the currently selected folder/file. Below that is the main browse window, which shows the complete filesystem tree including all folders/files contained on the drive or share.

In the browser portion of the window, you can scroll up and down the list of folders/files and tap individual folders to drill down to the desired level to expose the names of individual files located on the drive. Note that the size of each file is shown at the end of the filename. Many users will find this utility helpful when attempting to triage a large evidence set and determine the priority by which each drive should be imaged, or when checking the contents of a destination drive to free up space by deleting unneeded directories and files.

Note: Certain text and image files can now be viewed directly on TX1. See “Viewing text and image files” on page 135 for more information on this feature.

3.3.9 SMART data

This media utility is available for ATA drives that support SMART data reporting. Selecting this feature will display available SMART data as reported by the drive.

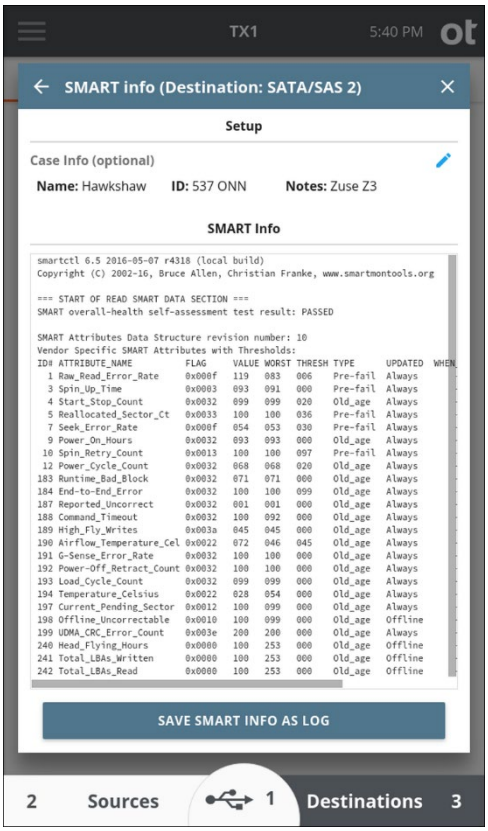

This information can be annotated with case info and saved as a log.

3.3.10 Export

This media utility allows the user to securely export any source, accessory, or destination drive as an iSCSI target. This makes the drive available to a remote user on any IP-based network (LAN/WAN/internet) via the Ethernet connection on the rear of TX1.

During the export process, TX1 will assign a unique IQN (iSCSI Qualified Name) to each exported drive. This IQN is used by the remote initiator to gain access to the exported media. The following is an example of an IQN produced by TX1 from exportation of a locally connected SATA drive: `iqn.2015-10.com.guid:sn-000ecc5801000c.sata.1`.

The breakdown of this number is as follows:

Number Portion	Description
`iqn.2015-10.com.guid:`	Base identification number/domain for all Tableau products (as provided by the iSCSI naming authority).
`sn-000ecc5801000c.`	The serial number of the specific TX1 in use, which serves as a unique identifier within the domain.
`sata.1`	TX1 specific suffix. The default is the protocol of the drive being exported with an incrementing number at the end. However, the user can edit this portion of the IQN via the optional Suffix entry field on the iSCSI Export setup screen (as shown below).

If successful, the final exported IQN is provided along with a listing of any initiator limits (IP address and/or IQN values). The IQN of the exported drive should then be available for selection via the **Discovery** function of an initiator on that same network.

Any exported drives can be un-exported by navigating to the iSCSI Export media utility for the drive and tapping on the **Remove Export** button in the lower right portion of the **iSCSI Export** screen.

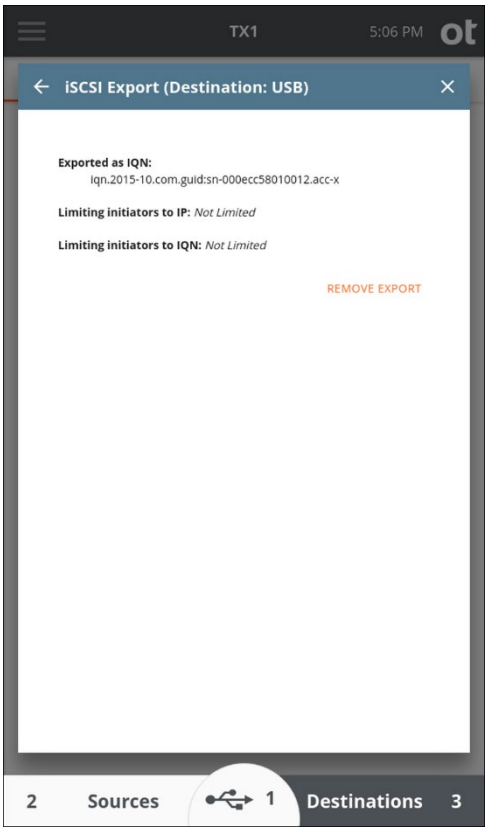

3.3.11 Eject

This media utility is provided to allow for safe ejection of attached drives. Ejecting a drive removes it from the system software in a safe manner and is recommended before unplugging any attached media from a powered TX1 and before powering down TX1 with drives attached. For destination and accessory drives in particular (since they are read/write), failure to eject a drive prior to removal from the system could corrupt the drive filesystem, which could result in loss of previously captured evidence/data. Note that ejection of media being used in a job will not be allowed until the job is complete.

In addition to quiescing the drive for system removal, pressing the **Eject** button will issue an ATA spin down command to drives that may support it. Spinning down rotating hard disk drives is recommended to minimize the chance of platter damage upon physical removal of the drive from the system. Note that not all drives support this command, and some may take longer to eject from the system due to lack of spin down command support. But this is considered a minor inconvenience compared to the benefit of minimizing the chance of drive damage.

3.4 Connecting drives

The following procedures provide the necessary steps for safely connecting drives to TX1.

TX1 operates as a standalone device or with the TX1-S1 drive bay. To connect the drive bay, ensure TX1 is powered off, align TX1 into place on top of the drive bay, and slide it back to lock it into place.

3.4.1 Source drives

Connect one or more drives to the TX1 source (left) side interfaces: SATA/SAS (x2), USB 3.0, PCIe, or FireWire.

Note: The PCIe interface on TX1 does not support hot-swapping. PCIe adapters with drives installed must be attached to TX1's PCIe port before powering up the unit. Also, removing a PCIe adapter from TX1 or removing a drive from a PCIe adapter before powering down TX1 will cause system instability and could possibly damage the drive, adapter, or TX1. TX1 must be powered down before removing the drive from any Tableau PCIe adapter or removing a PCIe drive/adapter from the TX1 PCIe port.

A Tableau PCIe adapter (sold separately) is required to acquire PCIe drives. The Tableau IDE-PCIe adapter (TDA7-5, sold separately) is required to acquire IDE drives. For detailed information on all the available PCIe based adapters for TX1, see the Adapters section of the Tableau Hardware products webpage (https://security.opentext.com/tableau/hardware?types=Adapters).

TX1 can acquire certain Apple computers that support Target Disk Mode via the USB 3.0 or FireWire source side connections. This can be done via three different types of Apple computer interfaces: USB-C, FireWire, and Thunderbolt 2. Commercially available adapters are required to convert these interfaces to USB 3.0 (for USB-C devices) or FireWire (for Thunderbolt 2 devices) for connecting to TX1. See the Adapters section of the Tableau Hardware products webpage (https://security.opentext.com/tableau/hardware?types=Adapters) or contact Customer Support for more information on the required adapters.

TX1 provides an Ethernet connection which enables network-based acquisition of iSCSI physical media (clones and physical or logical images) and mounted CIFS shares. See “Duplication over a network” on page 82 for more information.

Source drives are listed in the user interface in the order of TX1's physical port layout: iSCSI/CIFS, USB, FireWire, SATA/SAS 1, SATA/SAS 2, PCIe.

3.4.2 Destination drives

Connect one or more drives to the TX1 destination (right) side: SATA/SAS (x2) or USB 3.0. If using the Drive Bay, insert one or more 2.5” or 3.5” SATA/SAS drives.

Note: The SATA/SAS destination DC OUT ports directly on TX1 are enabled even when a TX1-S1 Drive Bay is connected, allowing for easy connection of up to four SATA/SAS destination drives. Please refer to the specifications section of this user guide for information regarding maximum power configurations.

TX1 also provides an Ethernet connection, which enables use of network-based iSCSI physical media and CIFS shares as destinations to store clone or image file data. See “Duplication over a network” on page 82 for more information.

Destination drives are listed in the user interface in the order of TX1's physical port layout: iSCSI/CIFS, USB, SATA/SAS 1, SATA/SAS 2.

3.4.3 Accessory drives

Connect up to two USB drives to the Accessory USB 3.0 ports on the front of TX1. These drives are mostly used for saving stored log files or updating TX1 firmware.

Note: The USB accessory ports on TX1 are NOT write-protected! Evidence media should never be connected to these ports.

Accessory drives are listed in the user interface in the order of TX1's physical port layout, with the leftmost USB drive always on top and the rightmost on the bottom.

3.4.4 Drive detection

After booting, TX1 begins detecting connected drives sequentially. Source and destination drive counts are displayed in the and **Destinations** buttons at the bottom of the TX1 **Home** screen, indicating the number of source and destination drives recognized. If any accessory USB drives are connected and detected, the center USB accessory portion of that bottom row will appear with the detected accessory drive count as well (as shown in the example below).

Tap the **Sources**, **USB Accessory**, or **Destinations** button to view more detail about the connected drives and to access media utilities.

The bottom of the TX1 screen with the source, USB accessory, and destination drive counts highlighted, is shown below:

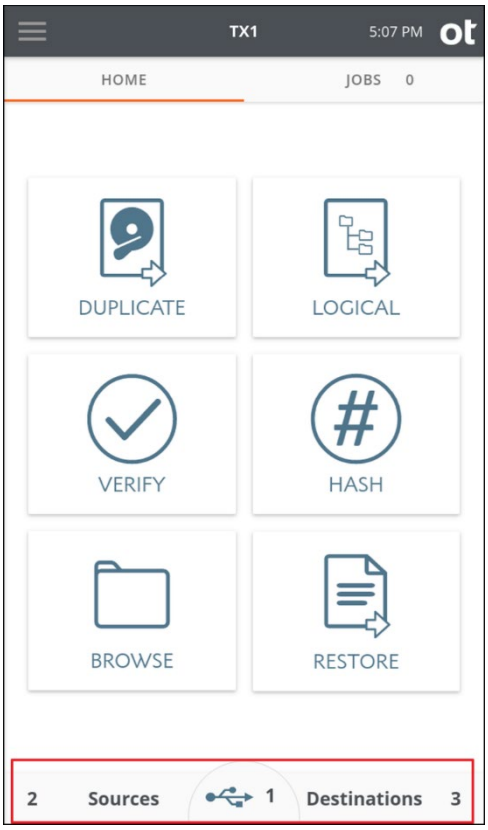

TX1 can detect USB drives that expose a CDFS volume. This is a common configuration for proprietary self-encrypting drives. The small CDFS volume typically contains an application that can be run on a host computer system, which allows for entering credentials that will unlock the drive. TX1 cannot run these proprietary applications, as they are typically made for x86-based Windows systems and thus TX1 cannot unlock these types of self-encrypting drives. That also means it cannot access the data volume of the drive (even in encrypted form) and thus cannot create an image of the drive. However, TX1 will detect these drives and report their type.

4 Using TX1

This chapter covers detailed procedures and information for using TX1.

4.1 Navigating TX1 features and options

The outline below maps TX1 navigation and feature structure.

4.1.1 Home screen

- Duplicate
- Logical
- Verify
- Hash
- Browse
- Restore
- Sources, Accessory Drives, and Destinations
    - View connected drive detail
    - Access media utilities

4.1.2 Jobs screen

- Job summary list and job details/status

4.1.3 Side navigation menu

- Home shortcut
- Logs
- Settings (system, network, and operation defaults)
- User Management
- Lock system
- About
- User

4.2 Preconditions checking

Before starting duplications and other jobs, TX1 automatically checks for preconditions. Some preconditions produce warnings, and you can choose to continue or cancel after viewing each one. Some preconditions are fatal and require that the duplication process be aborted.

4.3 Duplicating

TX1 will duplicate one source drive to up to four destination drives simultaneously per job. Up to two jobs can be active at any given time. For a given source/job, the destination drives can be a mix of directly connected drives and network shares. The destinations can be a mix of cloned and imaged copies.

Note: This section is focused on whole disk duplication operations. Please refer to “Logical imaging” on page 105 for details on that alternative acquisition method.

4.3.1 Cloning

A clone, also known as a disk-to-disk duplication, makes an exact copy of the source drive to the destination drive(s).

If a destination drive is not blank, TX1 displays a yellow warning to indicate that a clone will overwrite the contents of the destination drive. This reduces the risk of overwriting valuable data.

There is no need to format the destination media, as the clone will apply the file system of the source media (if one exists) to the destination media automatically. It is, however, a best practice to wipe destination media before duplicating to it as this can help to identify potentially defective media and bad sectors, and it can reduce the risk of cross-contaminating a duplication with stale data.

Note that, at the beginning of a clone job, TX1 prepares the destination drive by wiping sectors 0, 1, and end-of-drive minus 1. This ensures there is no stale partition table data on the drive, which reduces the possibility of drive detection issues at the end of the job.

Note: Because partition table information is relative to the sector size of the source drive, cloning to a destination drive with a different sector size is not allowed. TX1 will detect this sector size mismatch issue and warn the user. This condition will need to be rectified before the clone job can be started.

4.3.2 Imaging

An image, also known as disk-to-file duplication, copies the source drive to a series of files (sometimes called segments or chunks) on the destination drive. TX1 supports EnCase file formats ex01 and e01 and raw file formats dd and dmg. Compression is supported, and enabled by default, with ex01 and e01 file formats. File sizes from 2 GB per segment to Unlimited are supported. Smaller segments create more directory entries and Unlimited creates one large file segment.

Note: Not all image file size options are available in all situations. Due to filesystem addressing limitations, FAT32 formatted destinations have a maximum file size of 2 GB.

When imaging, the destination media must first be formatted with a recognized filesystem. Format destination drives by selecting the **Format** media utility in the drive details screen. The drive details screen can be accessed through the **Destinations** button on the **Home** screen or through the **Select Destinations** screen during setup of a duplication job.

If the destination drive is smaller than the source, a dd or dmg image will not fit on the destination drive. However, if using ex01 or e01, the source drive may fit on a smaller drive because these formats can compress the data before writing to the destination drive. There is no guarantee that the data will be compressed enough to fit on a smaller destination drive, especially if in cases where the data is mostly incompressible such as encrypted data.

Note: Use extreme caution when attempting to copy a source drive to a same size or smaller destination drive. Image file formatting adds overhead and, when coupled with incompressible data (such as encrypted data), a larger destination drive may be needed.

4.3.3 Performing a duplication

To perform a duplication:

1. Follow the steps listed in “Connecting drives” on page 61 to connect the source drive and all relevant destination drives.
2. From the **Home** screen, tap the **Duplicate** icon. The **Duplicate** job setup screen is displayed.

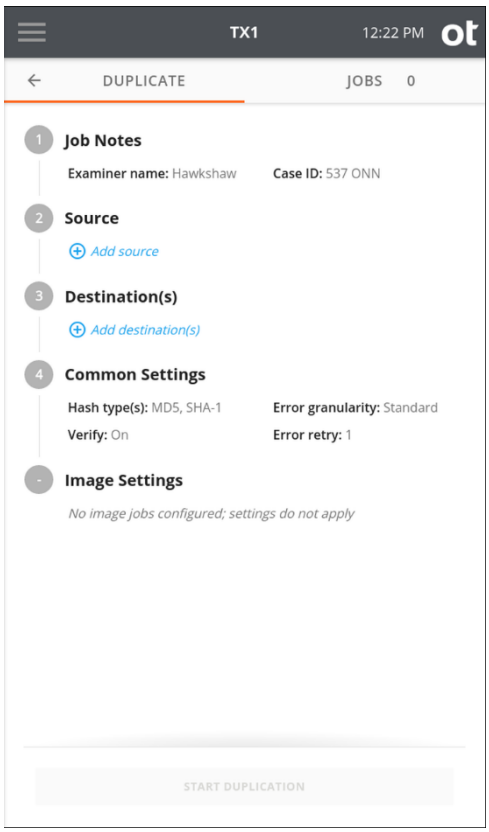

The job setup screen is organized in a natural workflow from top to bottom, but most steps and settings can be accessed in any order. The default values display for each step and setting. Tap on the step number or heading to expand the section and view or change the settings.

If only one source and one destination are connected, they are automatically selected. If you are satisfied with the default settings and the selected Source and **Destination(s)** drives, press the **Start Duplication** button at the bottom of the screen to begin the job.

3. To modify or enter job notes, tap the **1** or **Job Notes** heading to expand the section. Tap a text box to modify or enter **Name**, **Case ID**, or **Notes** values and the virtual keyboard is displayed on the bottom half of the screen. If desired, you can also attach a USB keyboard to one of the front Accessory USB ports to make data entry easier.

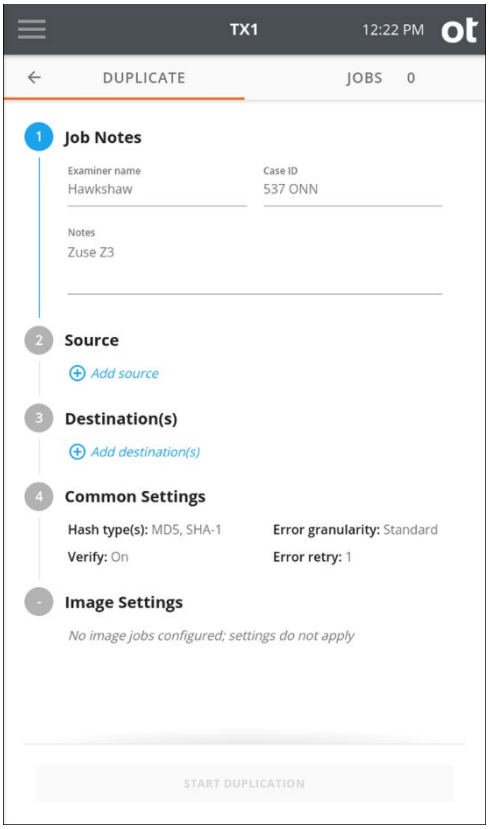

4. To change or add a source drive, tap the **2** or **Source heading**. From the source list modal that is displayed, select a drive from the list. A green check confirms your selection. Close the modal by tapping the X in the upper right corner or by tapping outside of the modal. Note that the source selection modal will automatically close once a drive has been selected. If a different source is desired, simply go back into the **Select a Source** screen by tapping on the **2** or **Source** heading from the **Duplication** stepper.

Notes:

- To help users identify which source drives have already been acquired, a green checkmark is shown in the bottom right portion of the drive tile for any drives that have been used in a previous, successful duplication job (clone or image). For currently active duplication jobs, a hollow checkmark will appear which will turn green once the duplication job has successfully completed.

- Within a screen displaying a list of drives, you can tap the options icon located in the right side of the drive tile to see more drive detail and access any available media utilities.

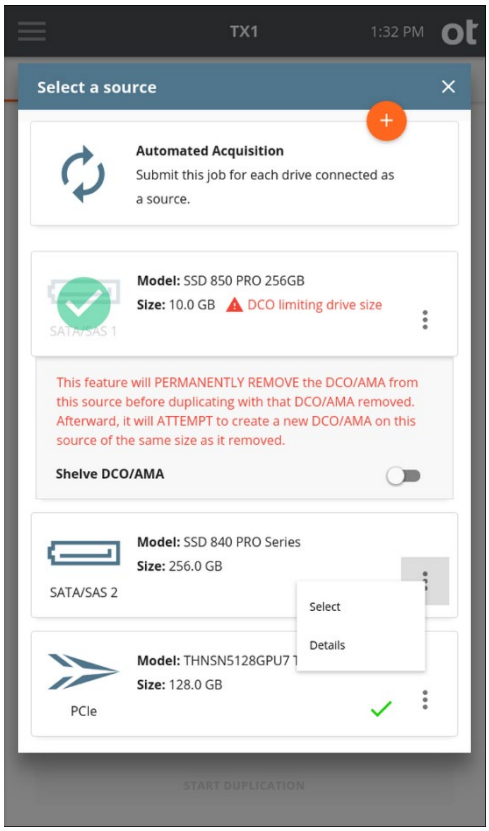

TX1 also allows “shelving” of a DCO or AMA for source drives. When enabled for a given source drive, the **Shelve DCO/AMA** feature will, after the duplication job is started, disable the DCO/AMA, complete the acquisition of the entire source drive, and then attempt to reapply the original DCO/AMA setting back to the source drive. This provides a convenient way to ensure that all source evidence is acquired and that the drive is returned to its original state at the end of the duplication job. See the source selection screenshot above for an example of a selected source drive with a DCO. Note that this **Shelve DCO/AMA** feature is only available for SATA drives that support DCO/AMA. See “Disabling drive capacity limiting configurations” on page 51 for more details regarding DCO/AMAs and how TX1 supports them.

Note: You cannot browse a source drive if the Shelve DCO/AMA feature is being used on an active job with that drive.

5. To change or add the destination drive(s) tap the 3 or Destination(s) heading. From the destination list modal that is displayed, select one or more drives from the list.

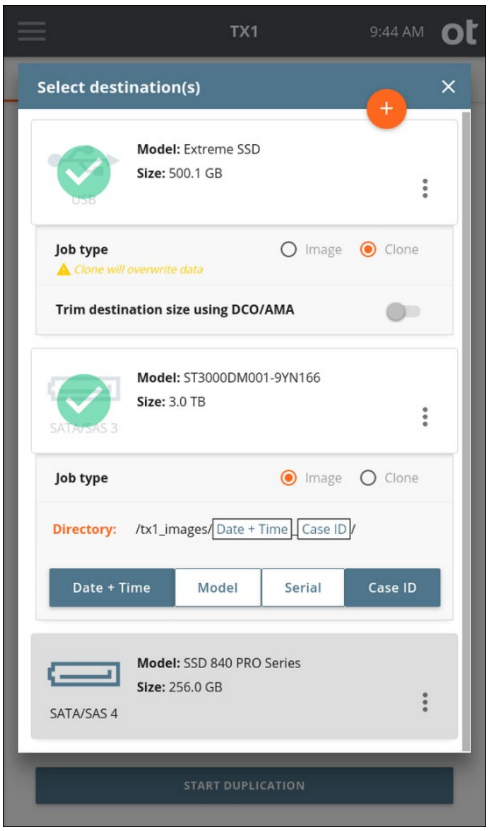

For each selected destination drive, a **Job type** panel expands below the Drive tile.

Select the **Clone** radio button to clone to the destination or select the **Image** radio button to image to the destination.

For clones, the option is displayed to trim the destination drive to be the same size as the source using a DCO or AMA.

Note: The trim feature is only available for SATA destinations that support DCO or AMA.

For images, the default destination directory path is displayed. To change the destination base directory, tap the orange **Directory** label to enter a **Browse** modal where you can select a different destination base directory, create and select a new directory, or delete a directory. Tap one or more of the four buttons (**Date and Time**, **Model**, **Serial**, or **Case ID**) under the directory path to add variables as names for a destination sub-directory. Each variable can only be selected once. Underscores are printed as separators between multiple variable names.

The **Image** option in the **Job type** panel will be disabled if the destination drive does not have a recognized filesystem. If an image destination is desired, format the destination drive by selecting the details option from the additional options menu (three vertical dots at the right side of the drive tile) or from the **Destinations** button on the **Home** screen.

To make a network share visible in the **Source** or **Destination** selection lists in a job setup screen, first add and mount the share from the **Sources** or **Destinations** buttons on the **Main** screen.

6. To change the common settings, tap the **4** or **Common Settings** heading.

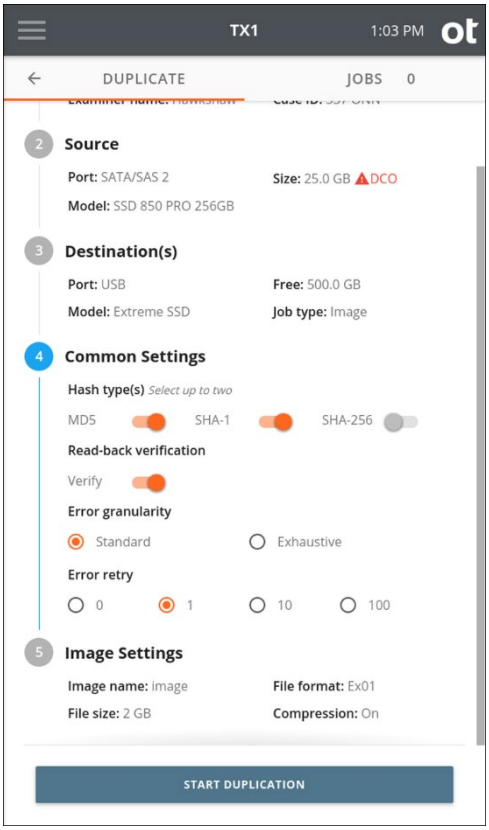

Select up to two hash types of MD5, SHA-1, and SHA-256. Some hash types will be disabled for certain image file formats, as described in the table below.

Duplication/ Image Type	Hash Type Options	Notes
DD/DMG/Clone	MD5, SHA-1, SHA-256	Can pick any two hash types. Limited to two to minimize chances of inadvertent performance degradation.

Duplication/ Image Type	Hash Type Options	Notes
E01	MD5 and SHA-1	E01 format does not currently support SHA-256. MD5 and SHA-1 forced on together since E01 format cannot support SHA-1 alone, and it was decided to not allow MD5 alone to simplify setting configurations.
EX01	MD5 and/or SHA-1	EX01 format does not currently support SHA-256. Can select any combination of MD5 and SHA-1 (including none if verification is not required) since EX01 format can handle all of them.

Error Granularity defines the granularity of failed reads from the source drive. The default setting of **Standard** attempts to recover data down to a 32 kB resolution. The **Exhaustive** setting attempts to recover data down to a single sector.

Note: The default error granularity setting is **Standard**, which will result in a minimum chunk of 32kB of source data (64 sectors for a 512B sector drive) that will get skipped and filled with zeros upon completion of the attempted reads (assuming no reads were successful). If this condition is encountered, consider changing the error granularity setting to be **Exhaustive**, which will result in repeated read attempts of the error region with decreasing sector sizes. This will maximize the amount of recoverable data and minimize the sectors that get skipped and filled with zeros.

Error retry defines the number of times TX1 will attempt to read sectors with errors before skipping the sector (and using a fill value of zeros). Be careful when selecting a value of 10 or 100 as it will drastically increase the duplication time when imaging source drives with hard errors.

7. If at least one destination Job type is image, then Step 5, Image Settings, is displayed as the last step. To change the image settings, tap the **5** or Image **Settings** heading.

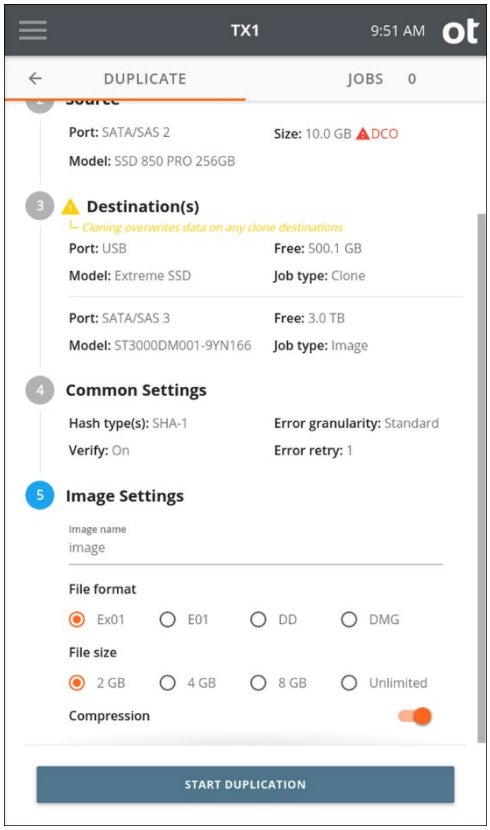

Image name defines the base filename for image segments. The default value is **image**. Tap the **Image name** text area to change the filename, then select a **File format** and **File size**.

Compression is enabled by default for ex01 and e01 file formats but can be disabled if desired. Compression is not supported for dd and dmg file formats.

Note: DMG files that were created from non-512 byte sector source media will not open on Apple devices by default. A special command line workaround is required when attempting to open such DMG files on an Apple device. TX1 will provide a warning in the **Image Settings** area when this condition is detected and DMG is selected as the file format.

8. Once you are satisfied with your settings and drive selection, tap the **Start Duplication** button. A **Job Status** screen is automatically displayed, as shown below.

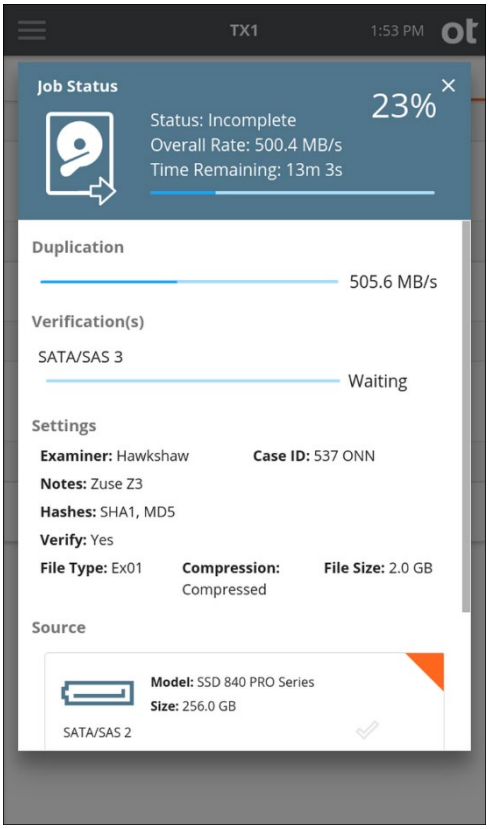

4.3.3.1 Files created during disk-to-file duplication

When performing an image, TX1 creates files (sometimes called segments or chunks) on the destination drive that contain the data copied from the drive.

Segments are written to the destination drive according to the following convention:

```
(image base directory)/
[directory name]/
 [filename].E01
 [filename].E02
 .
 .
 .
 [filename].E99
 [filename].log
 [filename].packed log
```

`[image base directory]` is defined in **Setting Defaults** or when selecting a destination drive. The default is `/tx1_images/`.

`[directory name]` is the image sub-directory name auto-generated for each acquisition and is defined in Setting Defaults or when selecting a destination drive during

duplication job setup. The default setting is Date and Time. See “Performing a duplication” on page 66 for more details.

`[filename]` is the base image filename and is defined in “Imaging” on page 65 during duplication job setup.

`[filename]` .001 (or .E01 or .Ex01) is the first segment or portion of the data copied from the source drive. All other segments have standard segment names (for example, `[filename].002, [filename].003`, and so on).

TX1 generates a .LOG file for each image job. It also creates a `.tx1_packed_log` file, which can be used to do a standalone verification of the original image or to Restore an image file to the original drive format.

4.3.4 Using automated acquisition

TX1 can be set up to run in Automated Acquisition mode, which will start a duplication job (clone or physical image) when any source drive is detected by the system. Custom job parameters can be set to the user’s liking when enabling this mode, and those settings will be used for all subsequent jobs within that automated run. This is a convenient way to quickly start jobs when you have many evidence drives to acquire as quickly as possible. This mode can also be used in a kiosk environment, where TX1 is always on waiting for a drive to be plugged in to start a duplication job. When used in conjunction with the ability to store forensic image files on a network-based iSCSI drive or CIFS share, this is a convenient way to quickly acquire as much evidence as possible with little user training, and make that evidence available to everyone in your network environment.

The steps below are focused on Automated Acquisition mode setup. Please refer to the more general Performing a duplication job setup workflow steps in the previous section for basic duplication job setup information.

Note: Automated acquisition is available for duplication (clone or image) jobs only. Logical imaging jobs typically involve custom analysis, triage, and targeted file set acquisitions, which are not aligned well with the goals of automated acquisition.

To set up Automated Acquisition mode:

1. From the **Home** screen, tap the **Duplicate** icon. The **Duplicate** job setup screen is displayed.

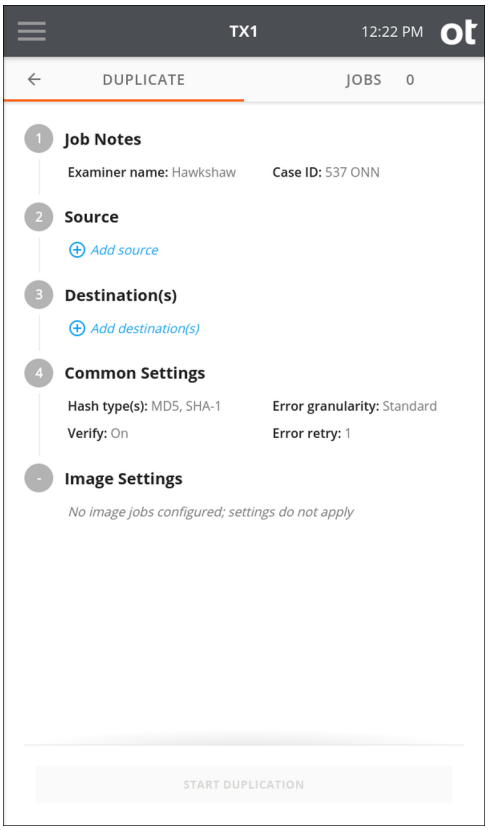

2. To modify or enter job notes, tap the **1** or **Job Notes** heading to expand the section. Tap a text box to modify or enter **Name**, **Case ID**, or **Notes** values and the virtual keyboard is displayed on the bottom half of the screen. If desired, you can also attach a USB keyboard to one of the front Accessory USB ports to make data entry easier. Note that all information entered in the **Job Notes** section will be used for each of the ensuing automated jobs.

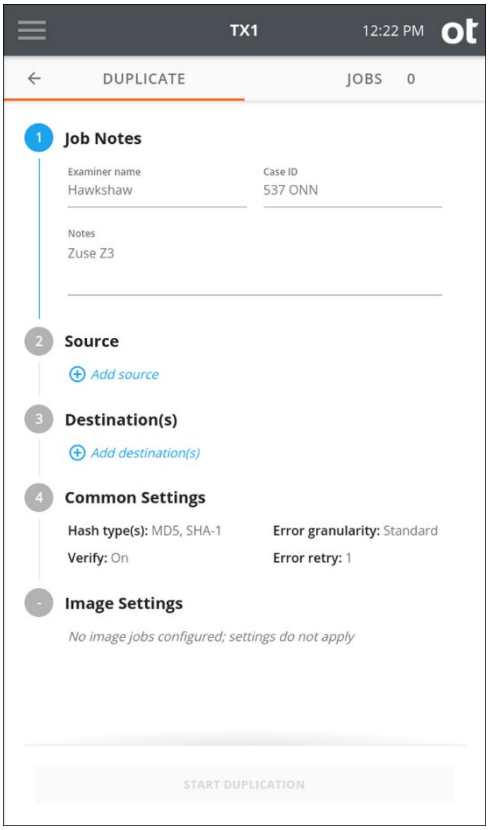

3. To set TX1 into Automated Acquisition mode, tap the **2** or **Source** heading. The top source drive tile is used for automated acquisition setup and will be present whether other source drives are connected or not. Tap the **Automated Acquisition** setup tile to start the automated job configuration process. Three options will be shown that need to be set before proceeding, as follows (and as shown in the screenshot below):

   - **Acquire currently connected** – If this option is enabled, any source drives currently connected to and detected by the system will have jobs automatically started when the Automated Acquisition mode is fully enabled.
   - **Acquire newly connected** – If this option is enabled, any source drives that are detected by the system for the duration of the automated acquisition run will have jobs automatically started as soon as drive detection is complete.
   - **Shelve AMA/DCO as needed** – If this option is enabled, any drives with detected AMA or DCO settings will have those settings disabled prior to the start of the job and then replaced after the job is complete. See “Disabling drive

capacity limiting configurations” on page 51 for more information regarding the Shelve AMA/DCO feature.

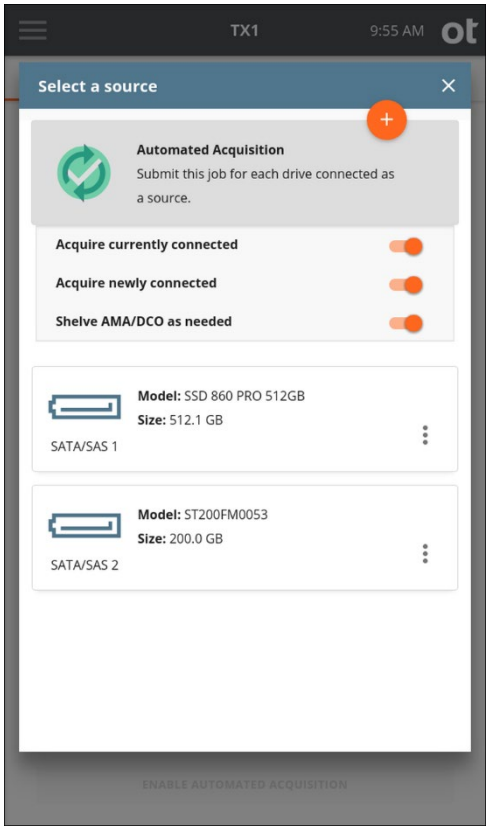

4. Your selection of Automated Acquisition mode will be confirmed by the automated acquisition icon in the left side of the drive tile turning green with a gray checkmark inside. Close the modal by tapping the **X** in the upper right corner or by tapping outside of the modal.

5. To change or add the destination drive(s) tap the **3** or **Destination(s)** heading. From the destination list modal that is displayed, select one or more drives from the list.

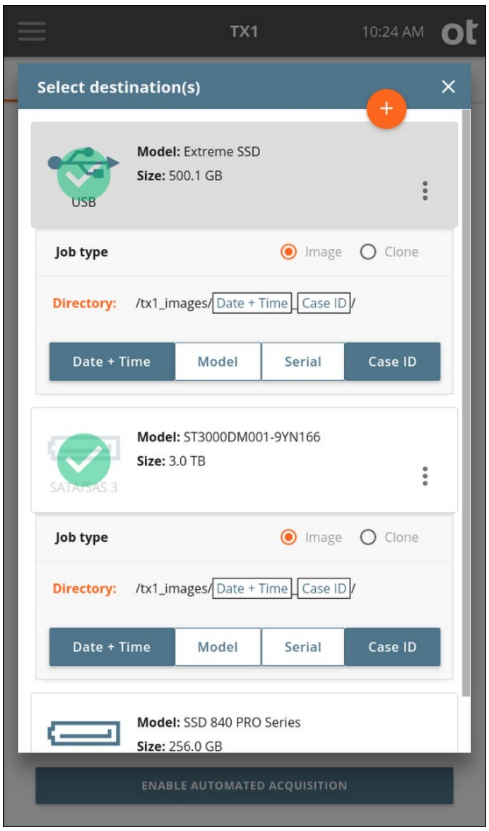

For each selected destination drive, a **Job type** panel expands below the **Drive** tile.

Note: Make sure to select **Image** as the job type for each destination drive. While it is possible to set a **Clone** job type here, that is not an option for automated acquisition jobs.

For image job types, the default destination directory path is displayed. To change the destination base directory, tap the orange **Directory** label to enter a **Browse** modal where you can select a different destination base directory, create and select a new directory, or delete a directory. Tap one or more of the four buttons (**Date and Time**, **Model**, **Serial**, or **Case ID**) under the directory path to add variables as names for a destination sub-directory. Each variable can only be selected once.

Underscores are printed as separators between multiple variable names.

The **Image** option in the **Job type** panel will be disabled if the destination drive does not have a recognized filesystem. If an image destination is desired, format the destination drive by selecting the details option from the additional options menu

(three vertical dots at the right side of the drive tile) or from the **Destinations** button on the **Home** screen.
To make a network share visible in the **Source** or **Destination** selection lists in a job setup screen, first add and mount the share from the **Sources** or **Destinations** buttons on the **Main** screen.

6. To change the common settings, tap the **4** or **Common Settings** heading.

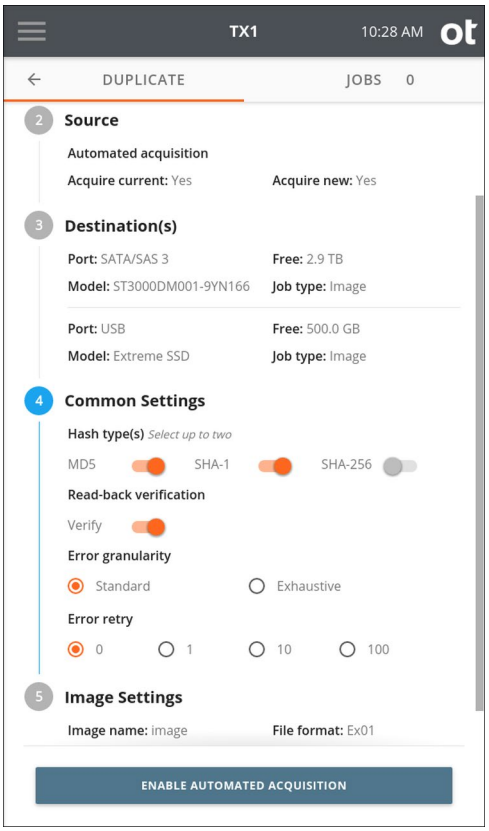

Select up to two hash types of **MD5**, **SHA-1**, and **SHA-256**. Some hash types will be disabled for certain image file formats such as ex01 and e01.

Error Granularity defines the granularity of failed reads from the source drive. The default setting of **Standard** attempts to recover data down to a 32 kB resolution. The **Exhaustive** setting attempts to recover data down to a single sector.

Note: The default error granularity setting is **Standard**, which will result in a minimum chunk of 32kB of source data (64 sectors for a 512B sector drive) that will get skipped and filled with zeros upon completion of the attempted reads (assuming no reads were successful). If this condition is encountered, consider changing the error granularity setting to be **Exhaustive**, which will result in repeated read attempts of the error region with decreasing sector sizes. This will maximize the amount of recoverable data and minimize the sectors that get skipped and filled with zeros.

Error retry defines the number of times TX1 will attempt to read sectors with errors before skipping the sector. Be careful when selecting a value of 10 or 100 as it will drastically increase the duplication time when imaging source drives with errors.

7. To change the image settings, tap the **5** or **Image Settings** heading.

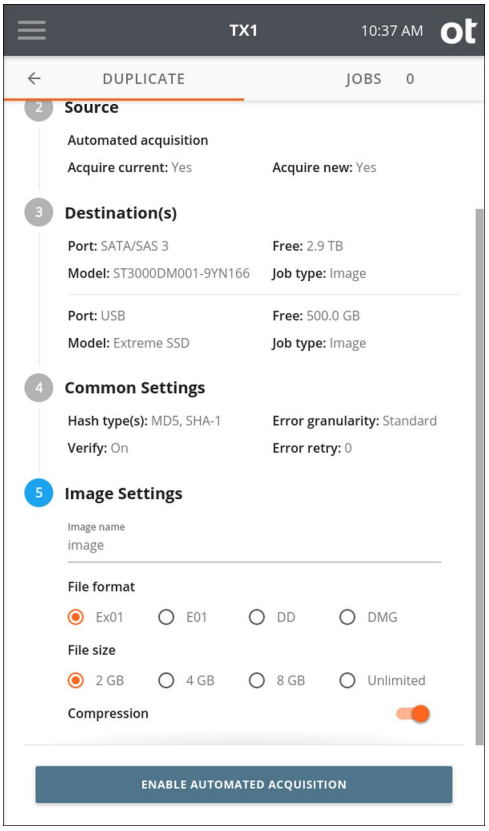

Image name defines the base filename for image segments. The default value is image. Tap the **Image name** text area to change the filename, then select a **File format** and **File size**.

Compression is enabled by default for `ex01` and `e01` file formats but can be disabled if desired. Compression is not supported for dd and dmg file formats.

8. Once you are satisfied with your settings and drive selections, tap the **ENABLE AUTOMATED ACQUISITION** button. The **Jobs** tab will be shown with information in the top **Automated Acquisition** area indicating that mode is active – the file type setting for all automated jobs and the number of automated jobs that have been started during the active automated acquisition job. Note that the automated acquisition indicator will be rotating in the Jobs tab header area next to the active/queued job count. This allows for at-a-glance awareness that automated acquisition mode is active, even when not looking at the Jobs tab.

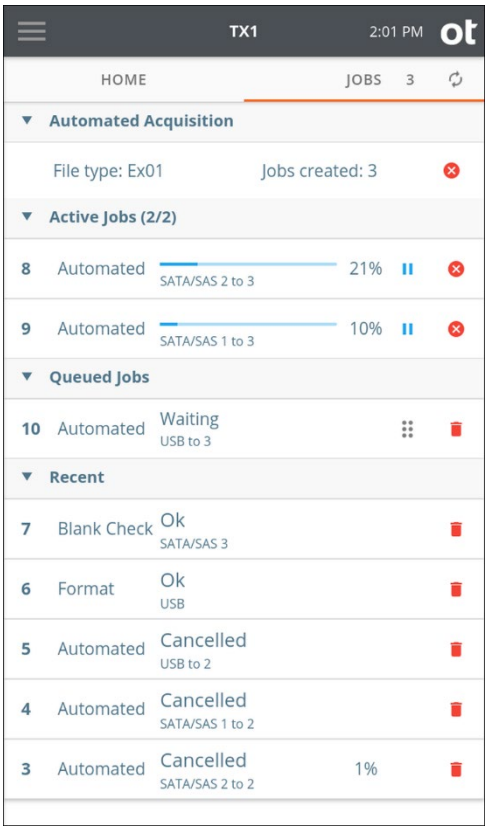

If any source drives were previously connected to the system and the **Acquire Currently Connected** option was set in step 3 above, then a job will be started or queued for each of the connected drives.

If no source drives were previously connected to the system or the **Acquire Currently Connected** option was disabled in step 3 above, then no automated jobs will be started until a new source drive is detected by the system.

It is important to note that, while in Automated Acquisition mode, new jobs can still be manually added, and they will run as soon as their required resources are available.

Any jobs that were automatically started by the system will show **(Auto)** next to the checkmark on the drive tile that indicates a drive was acquired (or is being acquire, in the case of a hollow checkmark). This helps you keep track of which jobs were manually or automatically initiated.

Note that automated acquisition jobs will never take you directly to the **Job Status** screen. These jobs are always shown in the **Jobs** tab as **Automated** (instead of **Duplication** or **Logical**), and their **Job Status** screen can be viewed at any time by tapping on the job row in the **Jobs** tab.

Automated Acquisition mode can be stopped by tapping the **Cancel** button on the right side of the **Automated Acquisition** job tile in the **Jobs** tab. Automated Acquisition job setup does not persist over a power cycle.

4.3.5 Duplication over a network

The 10–gigabit Ethernet interface of TX1 enables superior network imaging performance when combined with a properly configured 10–gigabit network infrastructure setup and destination storage server such as a Storage Area Network (SAN) or Network Attached Storage (NAS) or a Windows or Linux server configured with an iSCSI target or CIFS share. Duplication/Logical Imaging *from* network-based iSCSI targets and CIFS shares is also supported.

Note: TX1 must be connected to a network before attempting to mount and use iSCSI targets or CIFS shares.

4.3.5.1 Adding an iSCSI target

1. To add an iSCSI target as a source or destination drive, tap the **Sources** or **Destinations** button at the bottom of the **Home** screen. Then tap the orange plus button in the upper right corner of the drive list and tap **Mount iSCSI Target** to display the **iSCSI Discovery** screen.

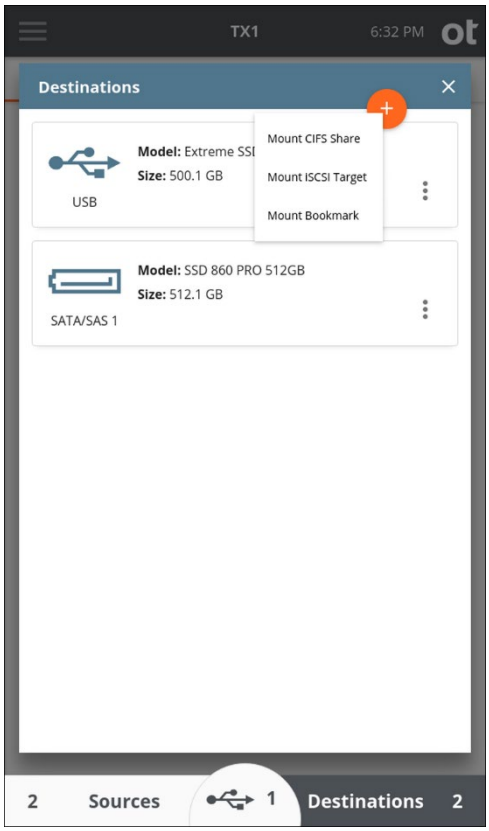

2. Enter the IP address of the iSCSI server by tapping on the **Address** field. If needed, change the default iSCSI **Port** from 3260 to the port used by the iSCSI server. If needed, enter a **Discovery Username** and **Discovery Password**.

   Note: Throughout the TX1 user interface, passwords can be shown as either plain text or hidden. Simply tap the eye icon to the right of the entry field to toggle the display mode.

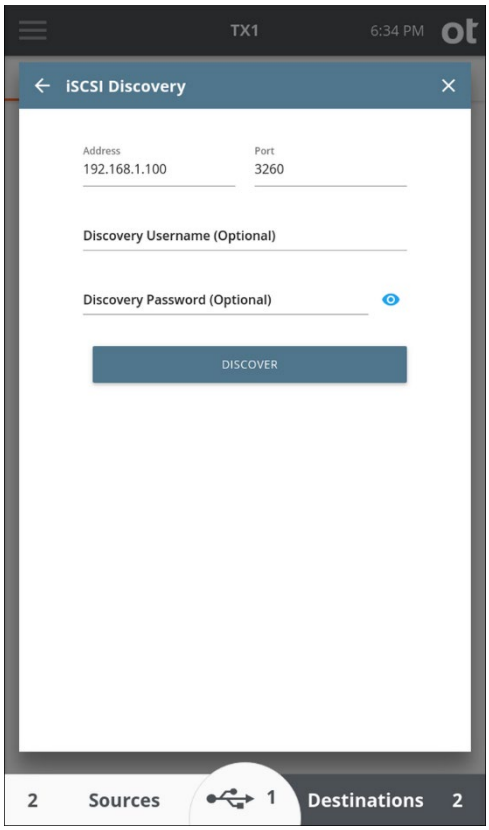

3. Tap the **Discover** button to discover available iSCSI targets. If the discovery is successful, a list of available iSCSI targets will appear below.

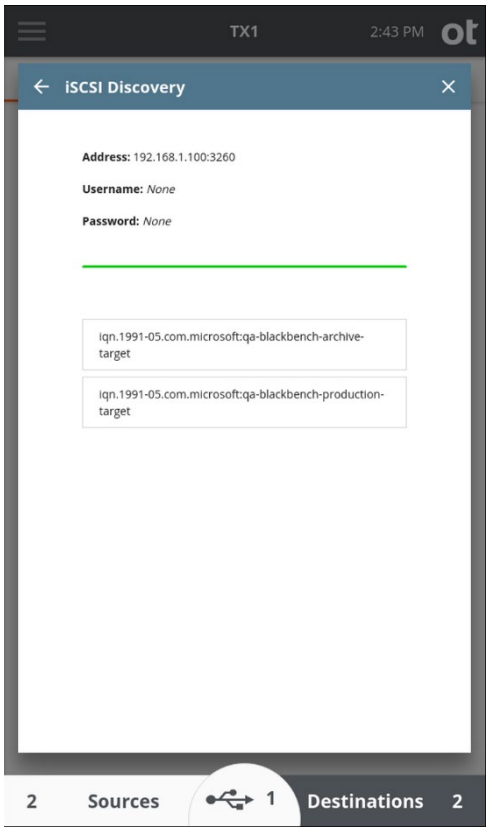

4. Tap an iSCSI target and the **iSCSI Login** screen is displayed. If needed, enter a login username, password, and a nickname (optional). Tap the **Login** button to login and mount the iSCSI target.

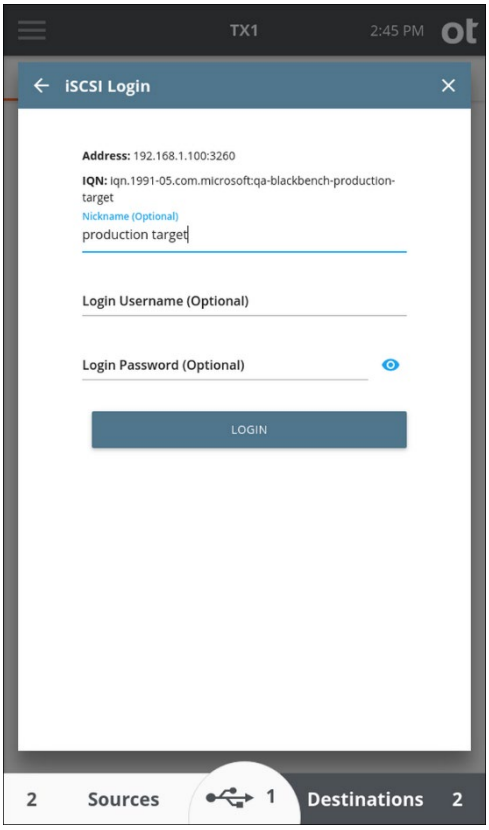

5. If the login is successful, you can optionally save the target as a Bookmark for convenient future access. To save a target as a bookmark tap the **Save As Bookmark** button under the iSCSI drive tile and enable or disable the desired **Username** and **Password** values to be saved. Then tap the **Save as Bookmark** button.

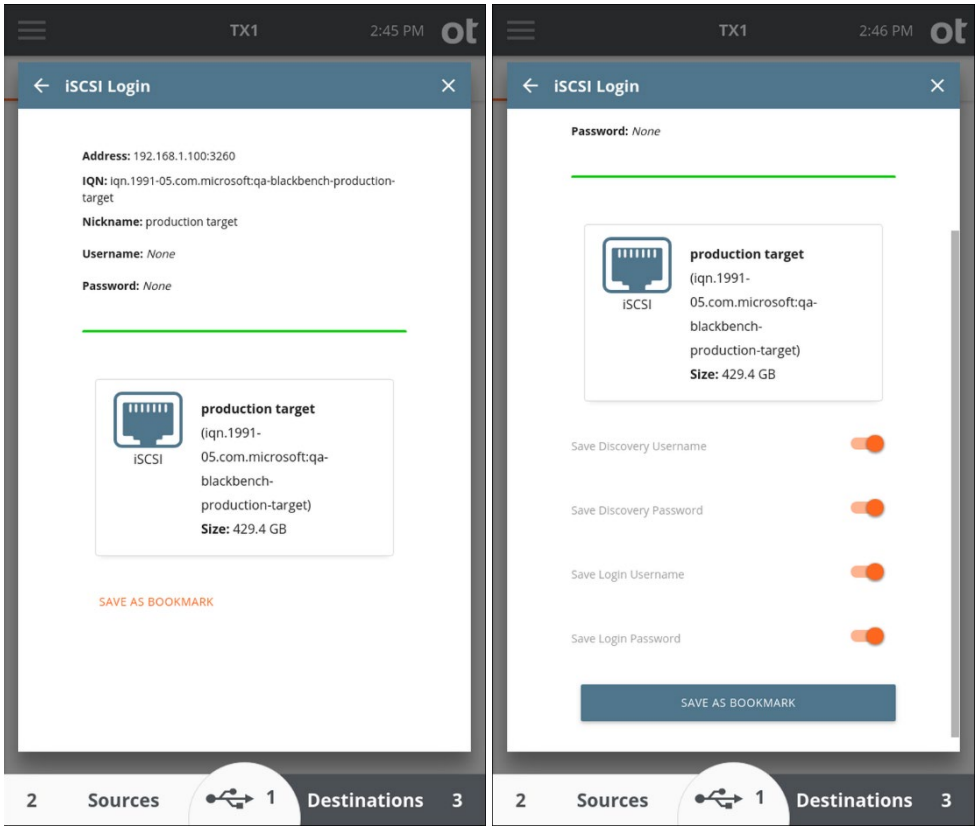

6. The target should now be listed in the Sources or Destinations drive list, depending on where you chose to mount it. The target can now be accessed like a normal drive for Duplication as a source (if mounted as a source), Duplication as a destination (if mounted as a destination), Hash (as a source), Verify (as a destination), and some media utilities.

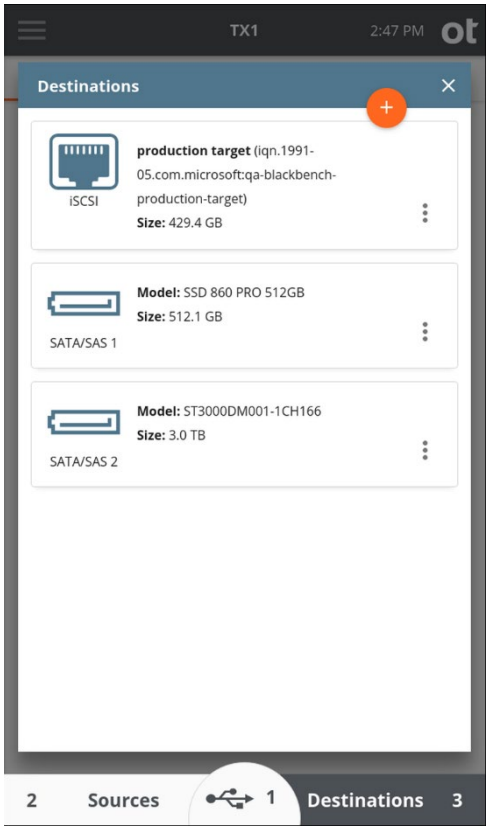

4.3.5.2 Adding a CIFS share

To add a CIFS share as a source or destination:

1. Tap the **Sources** or **Destinations** button at the bottom of the **Home** screen. Then tap the orange plus button in the upper right corner of the drive list and tap **Mount CIFS Share** to display the mounting screen.

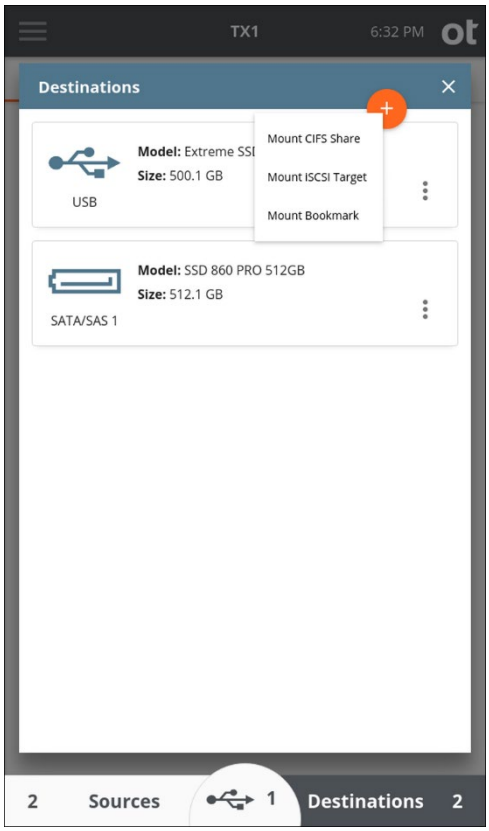

2. Enter the IP address of an available server and select **Next** or tap **List Servers** to select from a list of available servers on the network.

   Note: In Static IP setting cases or on networks with no domain name server (DNS), it is still possible to use a server's computer name to specify the share to mount.

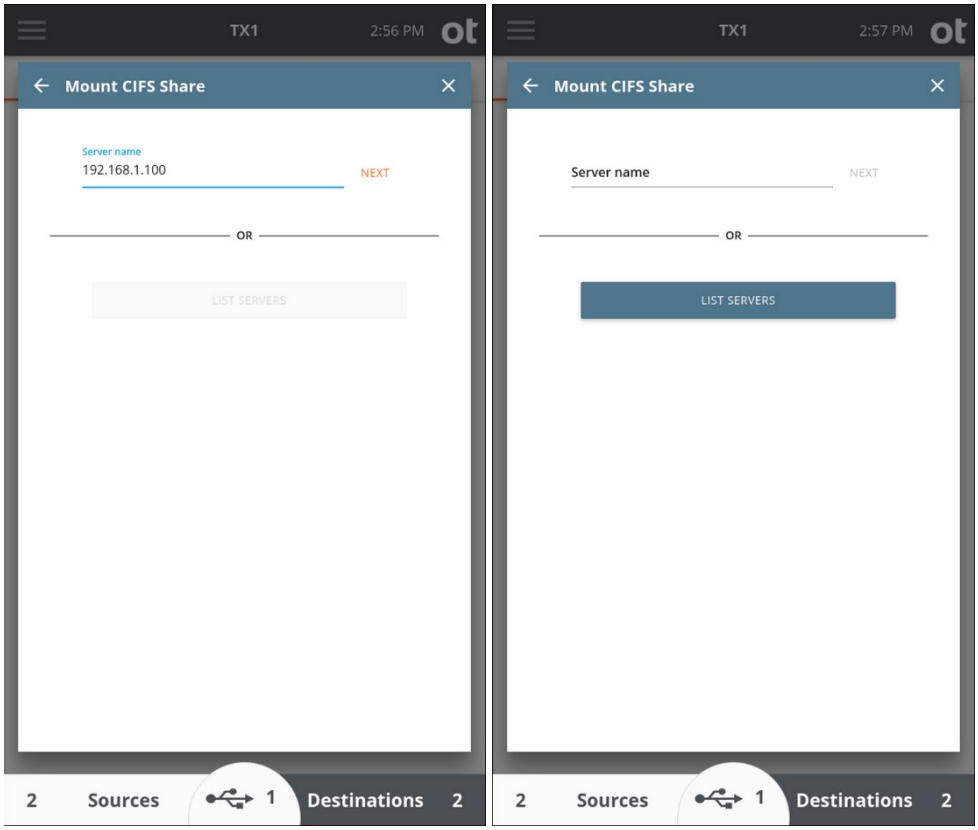

3. Enter a share name for the server listed in the status summary and select **Next** or tap **List Shares** to select from a list of available shares.

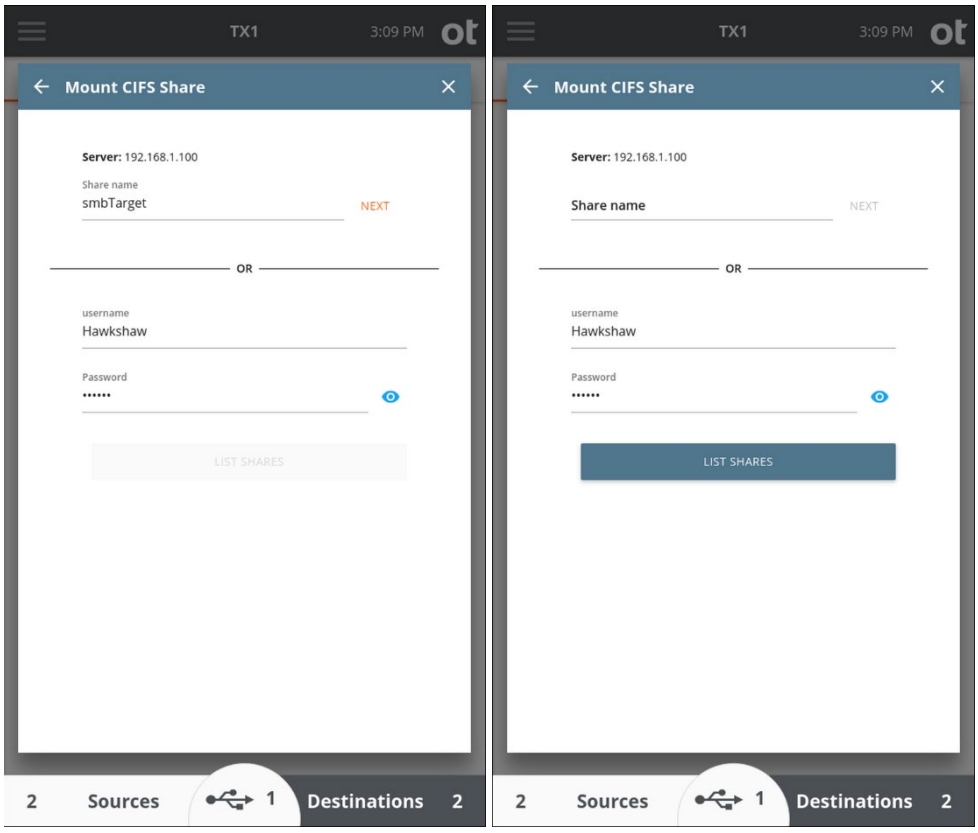

4. If you chose to use the **List Shares** feature, a list of available shares will be displayed with the currently connected shares identified by a grayed-out tile with a green check mark on the left. Tap the **Show Hidden Shares** slider to view default admin/hidden shares in the share list.

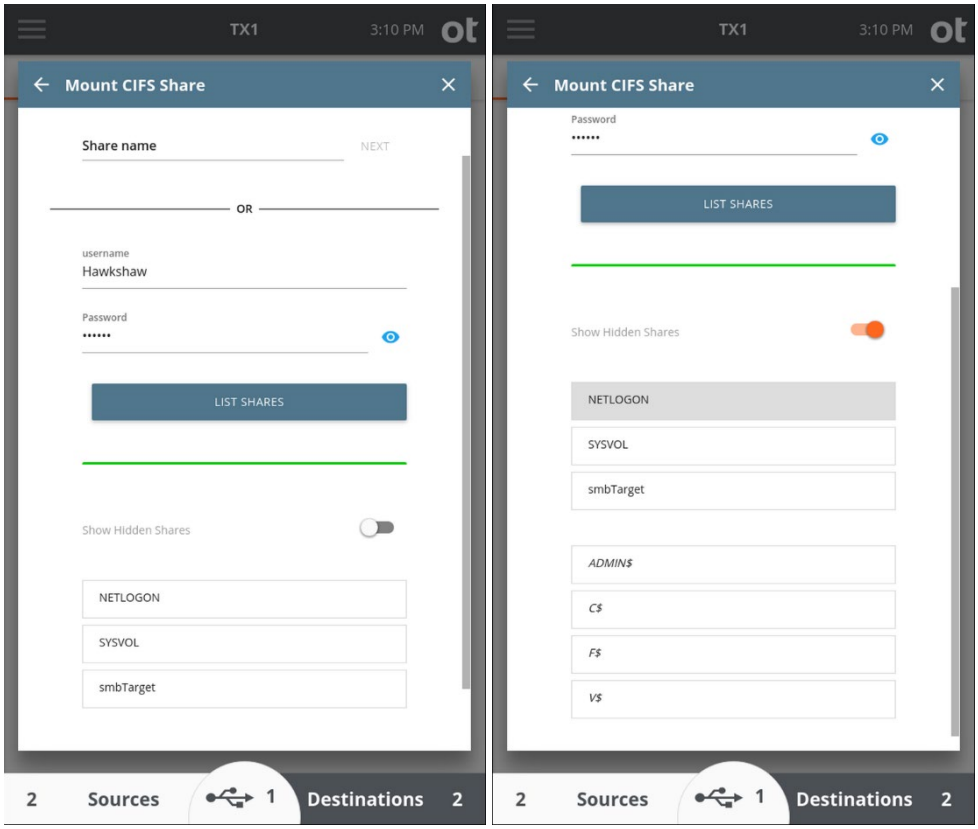

5. Enter a nickname for this CIFS share (optional) and enter a login username and password (if required). Choose the **SMB Version** and enable SMB 3.0 encryption (if desired), then tap the **Mount** button to login and mount the CIFS share.

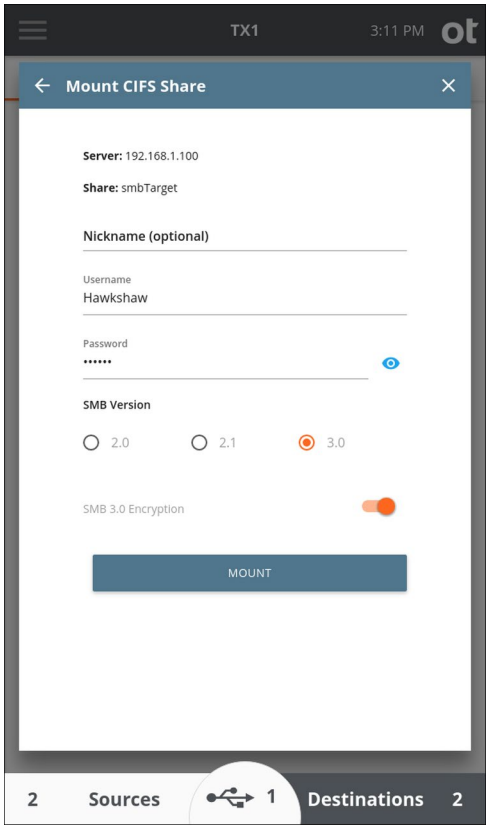

Note: Due to network security concerns, TX1 no longer supports SMB 1.0 as a mounting option for CIFS shares. Any bookmark previously stored that utilized SMB 1.0 will indicate that it is now invalid. The only course of action is to delete the existing bookmark and then remount and re-bookmark the share with a different SMB version selected.

6. The CIFS share should now be listed in the **Sources** or **Destinations** drive list, depending on where you chose to mount it. To save a share as a bookmark tap the **Save As Bookmark** button under the CIFS drive tile, enable or disable the desired Username and Password values to be saved, and then tap the Save as Bookmark button.

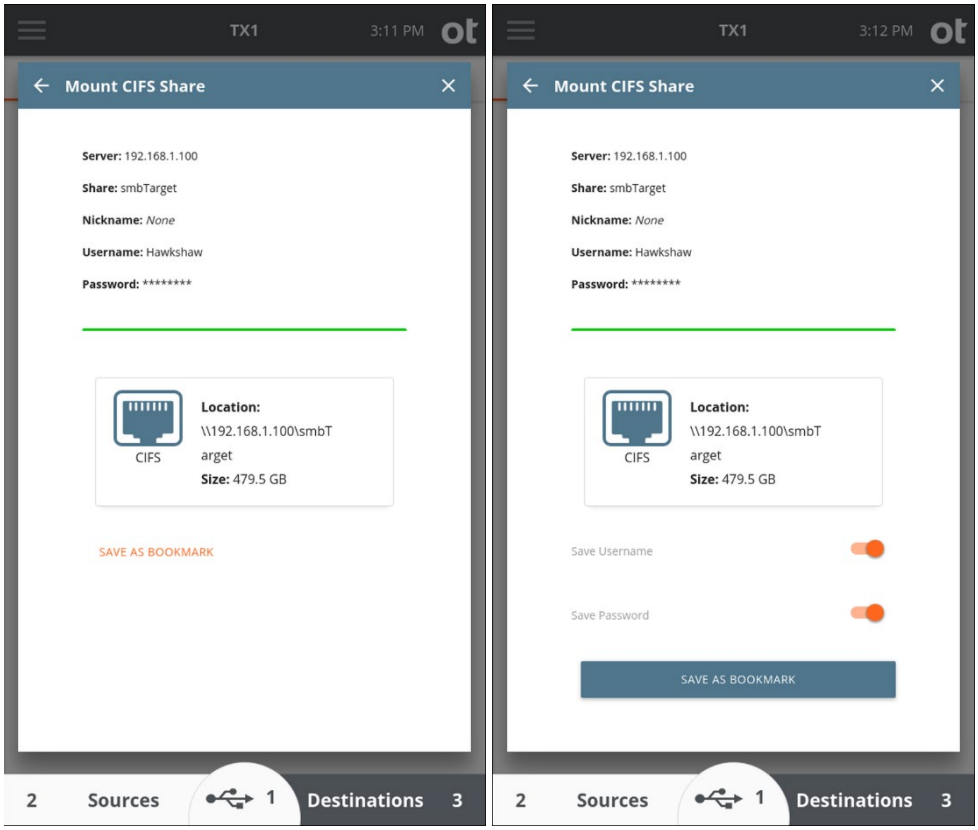

7. The bookmark is now saved (if selected). The share can now be accessed like any mounted filesystem for logical acquisition as a source (if mounted as a source), as a destination for physical and logical image files (if mounted as a destination), Verify (as a destination), Restore (as a source or destination), and some media utilities.

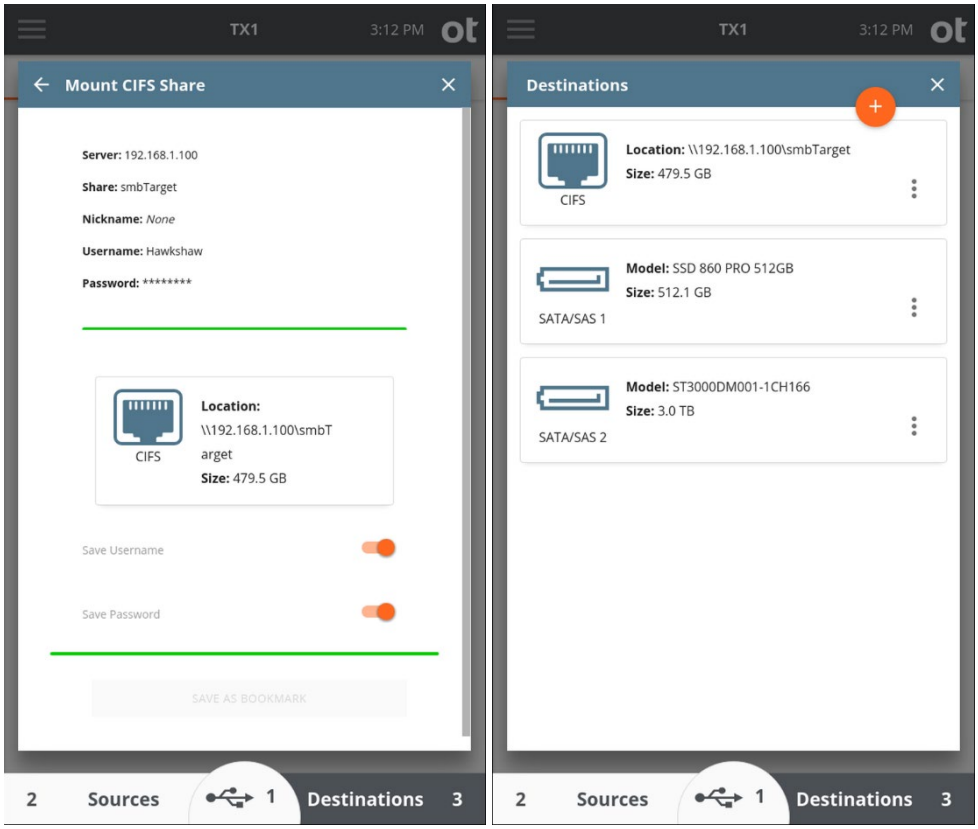

Note that a CIFS share takes the form of a filesystem (not a block device/drive) so you cannot perform a Clone Duplication to a CIFS share. The **Wipe**, **Blank Check** and **Format** options are also not available when a CIFS share is selected as a destination.

Note: The CIFS mounting steps above are shown for the case of a destination CIFS share. However, the same steps apply to mounting a CIFS share for use as the source of a logical imaging job, with the only difference being that the **Sources** drive list is selected from the bottom of the **Home** screen instead of **Destinations**.

4.3.6 Pausing and resuming a duplication job

In certain situations, significant amounts of imaging time can be saved by being able to pause and later resume a duplication job. TX1 has you covered, providing the means to pause and resume imaging jobs with the following output file formats: e01, ex01, dd, and dmg.

To pause a running duplication job, locate the desired job in the **Active Jobs** area of the **Jobs** tab, simply tap its **Pause** button ⏸, and confirm the desire to pause the job. The job will be moved to the **Recent** area with a status of **Paused**, as shown below.

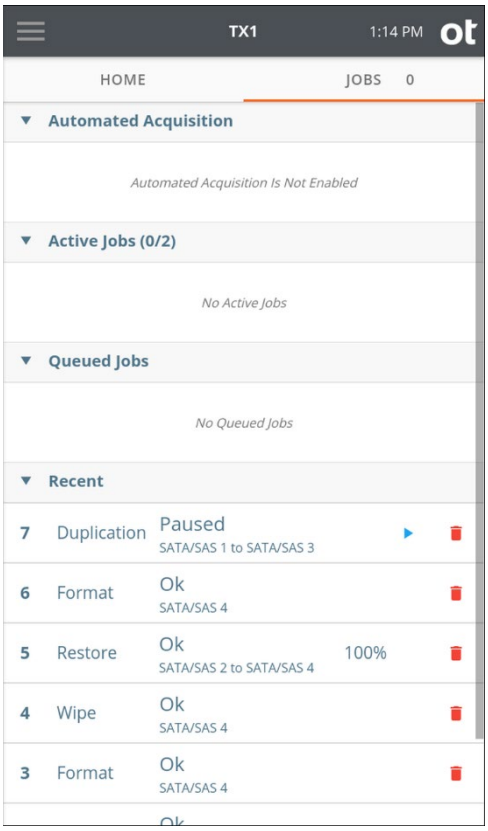

There are three distinct ways to resume a paused job:

- By tapping the **Play** button ▶ on the paused job in the **Recent** area of the **Jobs** tab.
- By tapping the **Resume Job** button in the header of the **Job Status** screen. (The Job Status screen can be viewed by tapping on the job in the Jobs tab.)
- By tapping the **Resume Job** button at the bottom of the forensic log for the paused job. All paused jobs display a paused status in the log list. Tapping the desired job log row displays the **Log Details** screen for that job, with the **Resume Job** button at the bottom.

Regardless of the method of resumption, a **Resume Duplication** screen will be displayed, as shown below. This screen allows for verification of the availability of the original job's source and destination drives before allowing resumption of the paused job. Note that, if Verification was not enabled in the original imaging job setup, the **Resume Duplication** screen offers a means of enabling it before resuming the job.

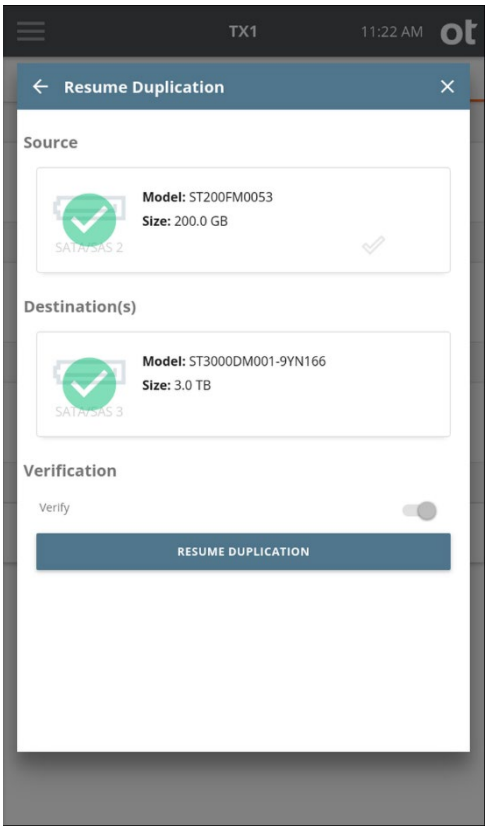

In addition to manually initiated pause and resume, TX1 supports power loss situations as well. For the supported job types (`e01, ex01, dd, dmg`), if power is unexpectedly lost during an imaging job, it can be resumed after power is restored and the system is booted up.

To resume a supported imaging job after a power loss event:

Note: As with manually paused jobs, the original source and destination drives are required by TX1 before the original job can be resumed after a power loss event. The steps/screenshots below assume the original drives are connected and available to the system after power was restored.

1. From the **Home** screen, navigate to the side navigation menu (available by tapping the menu icon at the top left of the **Home** screen).

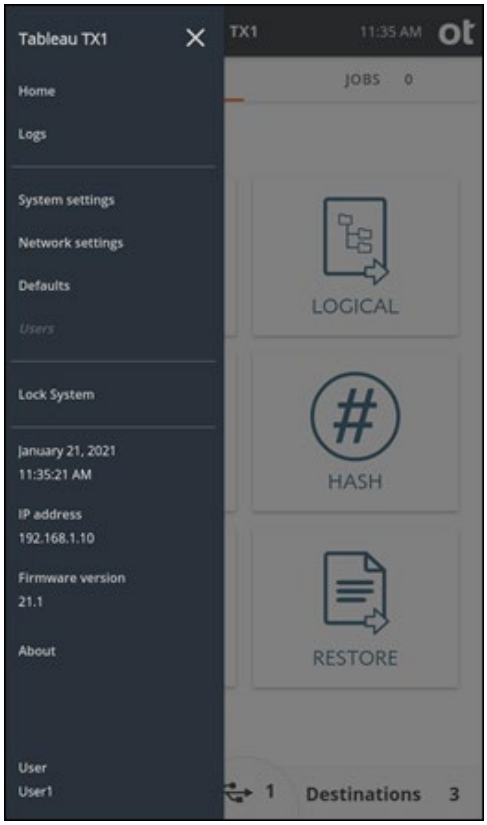

2. Tap the **Logs** menu item to see a list of all the stored job logs.

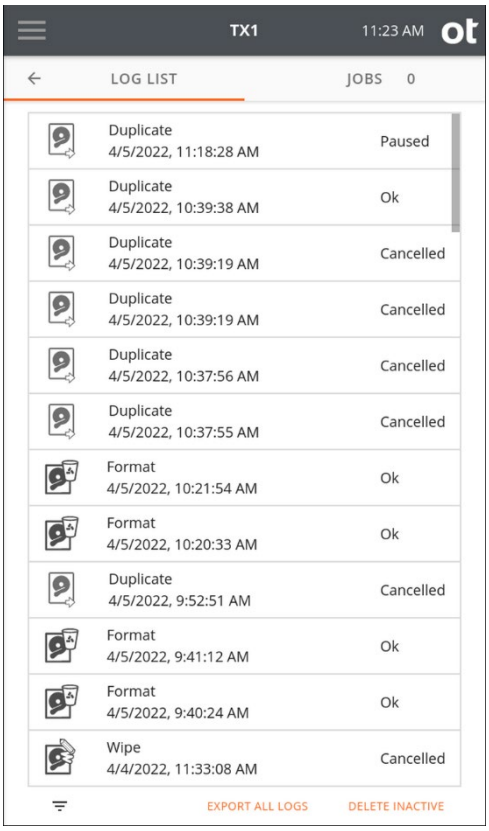

3. Find the desired paused job log (**Paused** status on the right, with the appropriate job start date and time shown) and tap on that log list entry to display the **Log Details** screen for that job log.

4. Review the log details to confirm this is the job that was running when power was lost that you intend to resume. Note that logs for completed jobs that experienced a power loss event will have a message at the top of the log indicating *** POSSIBLE POWER LOSS EVENT DETECTED ***. However, that message is added only after the power loss paused job has been resumed, so it will not be present when you initially view the paused log, prior to resuming that job.

   Verifying that the source and destination drives shown in the log represent the desired drives is recommended. When you are comfortable that this is the log for the job you want to resume, tap the **Resume Job** button at the bottom of the **Log Details** screen. The **Resume Duplication** screen for the paused job will appear, as shown below.

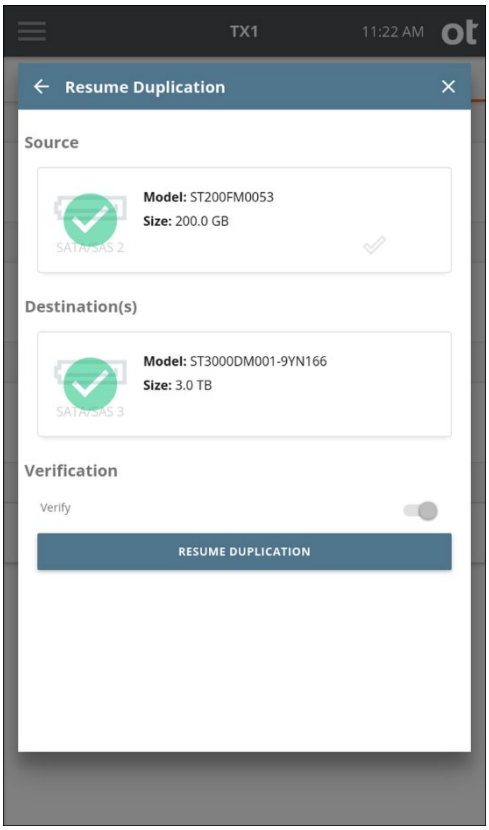

5. If the original drives are present and available to TX1, they will be shown with green checkmarks and the **Resume Duplication** button will be lit up. If your original job did not have Verification enabled, you can enable it on this screen if desired. Tap the **Resume Duplication** button to resume the job. The **Job Status** screen for the resumed job will be shown. From this point on, the job will carry on as if it had never been paused, with the exception of logged pause/resume events.

The forensic logs for paused and resumed jobs will provide some specific and unique information. The information differs slightly depending on the source of the pause event (manual or power loss). In the case of a manual pause event, a line will be added to the

log to indicate the date and time of the event. When unexpected power loss is the cause of the pause, there is no time for the system to log the pause time before shutting down, so that information is unavailable and thus not included in the log. In that case, a message is added to the log after the job is resumed to indicate that the missing pause information is likely due to a power loss event, and the job's elapsed time is not calculated, since it cannot be accurately determined. Each subsequent pause (if manually initiated) and resume event is logged, providing an accurate capture of how many pause/resume cycles occurred during the job. The following log sample shows a completed power loss paused/resumed job. Note that, had this been a manually paused/resumed job, the line with the possible power loss warning would be replaced by a **Paused** field, with the date and time of the pause event.

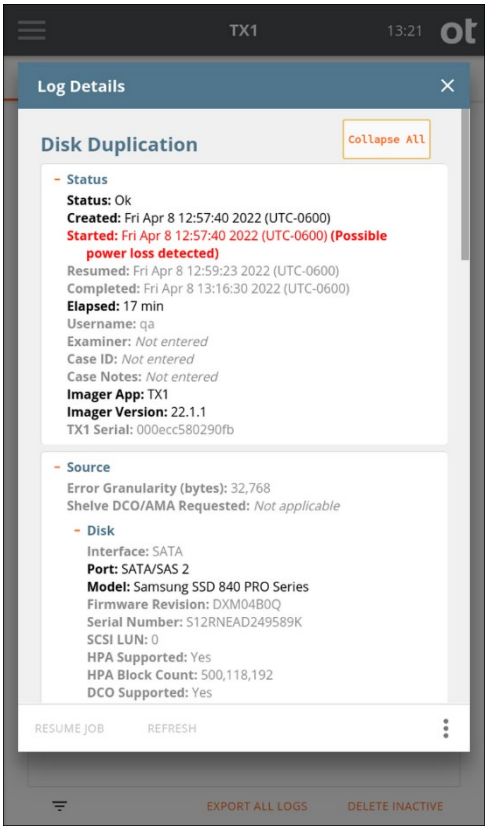

TX1 also allows the resumption of imaging jobs that failed for the following reasons:

- Source drive missing or disconnected
- Destination drive missing or disconnected
- Destination drive full

These types of failed jobs can be resumed through either the **Recent** jobs area in the **Jobs** tab or the job's log details screen, as described in the sections above. The best way to tell if a job is resumable or not is to check for the ability to resume it from the **Recent** jobs list or in the log details screen.

4.4 Hashing

Forensic practitioners may need to calculate the hash values, or fingerprints, for a source drive without making a copy of the drive. The Hash function can generate MD5, SHA-1, and SHA-256 hash values for a source drive. You can use up to two different hash algorithms in one operation.

To create a hash of a source drive:

1. From the **Home** screen, tap the **Hash** button.

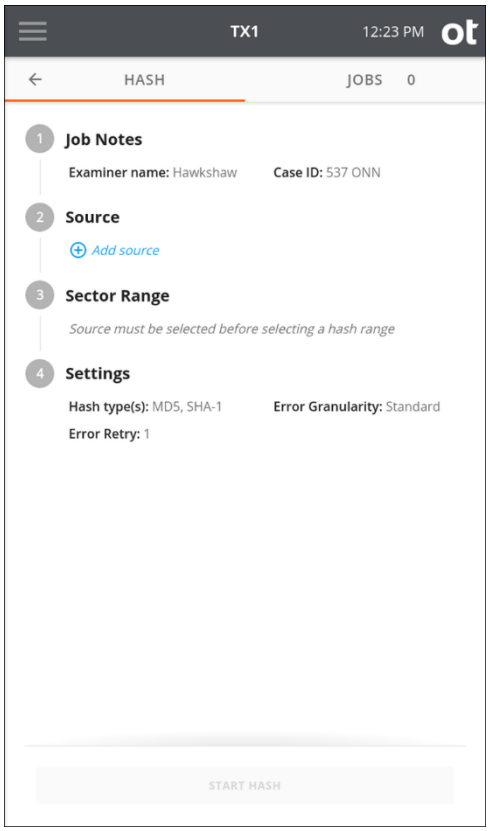

2. Enter **Job Notes** and select a **Source** drive.

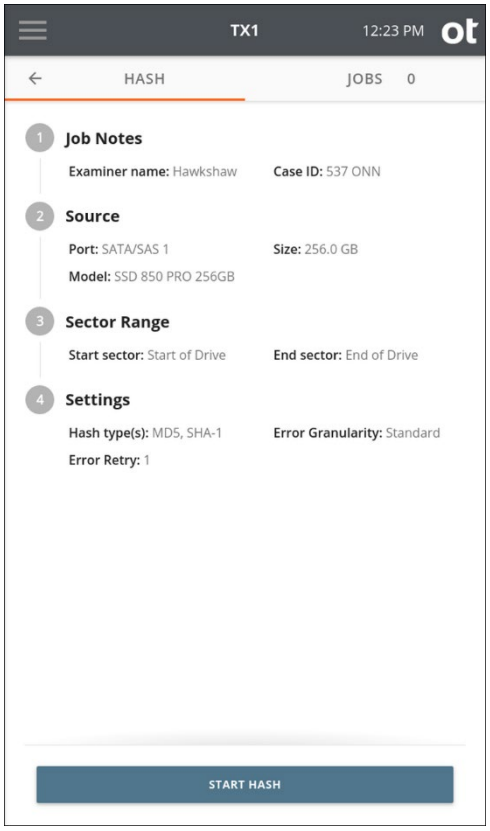

3. Select a **Sector Range** for the hash. The default settings will always provide a full drive hash, but certain situations (such as a failing source drive with bad sectors) could benefit from a partial drive hash. Partial drive hashes are defined by start and end sector values. These sector values can be set in two ways, as follows:

    1. Use the radio buttons to select your range. This will typically be partition-based if the drive has one or more partitions, or the entire drive if no partition table exists.

2. Manually enter the specific start and end sector numbers. Again, if you do not define a custom **Sector Range**, the entire drive will be hashed. After verifying the hash settings, tap the **Start Hash** button at the bottom of the screen.

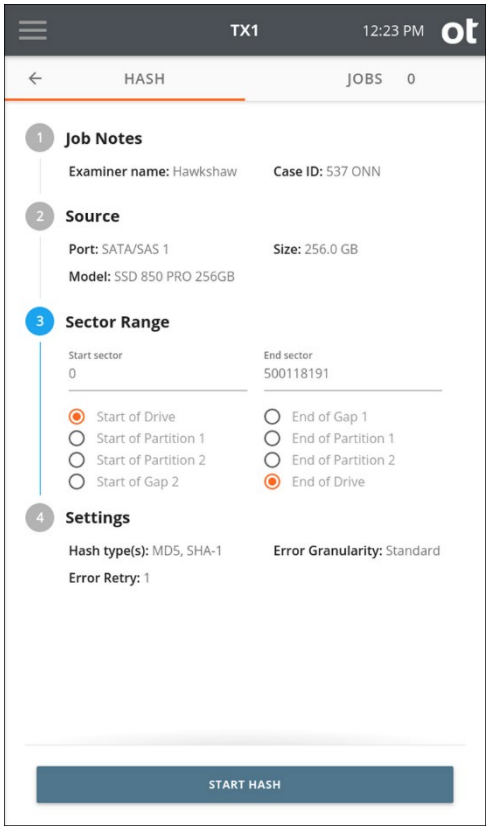

4. To cancel the hash operation, close the **Job Status** screen by tapping the **X** in the upper right corner, and then tap the **Cancel** button from the **Active Jobs** area at the top of the **Jobs** summary screen.

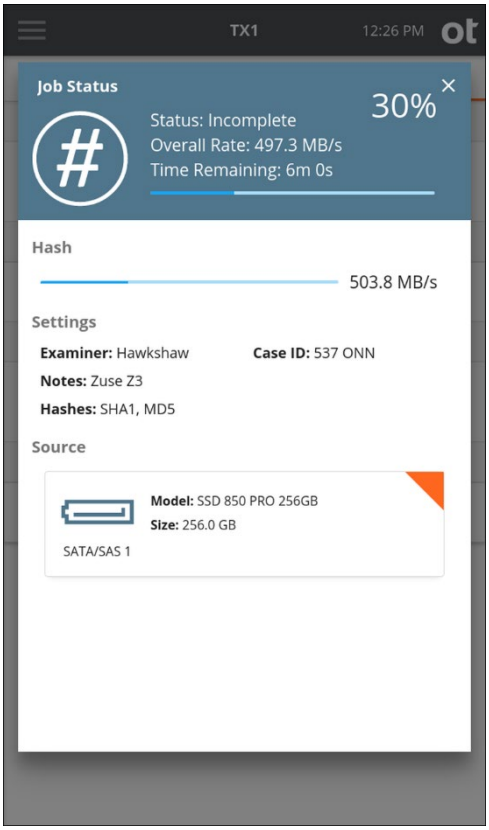

When the hash operation is finished, the results display on the screen. To access the forensic log of the hash job, simply tap the **View Log** button in the header of the completed **Job Status** screen. You can also view the log information by selecting **Logs** from the side navigation menu.

4.5 Logical imaging

TX1 provides a powerful logical imaging function that allows for file-based evidence acquisition from locally attached and network-based source filesystems. This tool saves valuable time by focusing on specific files of interest rather than acquiring the entire physical drive. TX1 logical imaging understands the structure of the recognized filesystem and acquires the desired source file data and/or metadata.

Logical imaging operations on TX1 can be targeted, limiting which files are acquired to only a subset of the source filesystem. TX1 allows both direct selection of contents to acquire, as well as rule-based searches to target specific files based on file type or other criteria. This allows for a more targeted and faster acquisition of only the files that are of forensic interest. Used in conjunction with physical disk imaging, logical imaging enables

rapid acquisition of source file data, providing TX1 users the ability to balance thoroughness with acquisition time and effort for the demands of a given case.

The logical imaging function can be configured to create logical evidence files (lx01 format) and/or metadata lists (comma separated value csv format). The industry standard lx01 logical evidence file format can be used with a variety of post-acquisition forensic analysis software tools, such as industry leading EnCase. The CSV metadata files can be configured to contain the metadata from only the files that were acquired in the lx01 file or from all the source files. Acquiring all source file/folder metadata provides a trail of what data was and was not acquired during the operation. This allows an investigator to quickly analyze a summary of the filesystem involved in the job.

Due to the wide range of variables in a logical image job, such as file data compressibility and acquisition dataset size, it is not possible to determine with certainty if the data from a source filesystem will fit on a destination filesystem. As a result, TX1 only warns the user that a destination may be too small when the used space of the source filesystem is larger than the available space on the destination, but the job can still be started. In a worst-case scenario (compression disabled and no source file down-selection), it is possible for the destination filesystem to become full, thus causing the job to fail.

Note: Use caution when attempting to logically image from a source filesystem to a smaller destination filesystem.

4.5.1 Performing a logical image acquisition

To perform a logical image acquisition:

1. Follow the steps listed in “Connecting drives” on page 61 to connect the source drive and all relevant destination drives.

2. From the **Home** screen, tap the **Logical** icon. The **Logical Image** job setup screen will be displayed, as shown below.

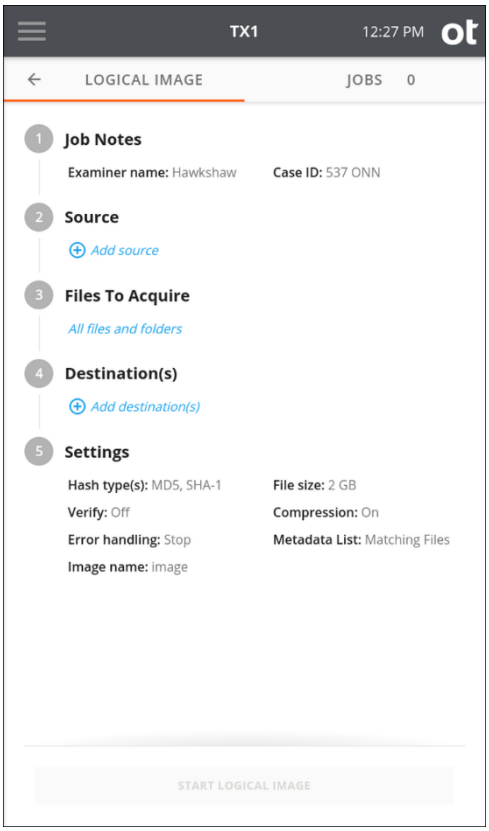

The job setup screen is organized in a natural workflow from top to bottom, but the steps and settings can be accessed in any order. The default values display for each step and setting. Tap the step number or heading to expand the section and view or change the settings.

If only one source drive with one recognized filesystem and one formatted destination drive are connected, they are automatically selected. If you are satisfied with the default settings and the selected source and destination filesystems, press the **Start Logical Image** button at the bottom of the screen to begin the job.

Note: Logical image job source auto-selection is distinct from duplication job source auto-selection in that a logical image job operates at the filesystem level, not the drive level. Therefore, a sole source drive will not be auto-selected for a logical image job if it contains more than one recognized filesystem.

3. To modify or enter job notes, tap the **1** or **Job Notes** heading to expand the section. Tap a text box to modify or enter **Examiner name**, **Case ID**, or notes values and the virtual keyboard is displayed on the bottom half of the screen. If desired, you can also attach a USB keyboard to one of the front Accessory USB ports.

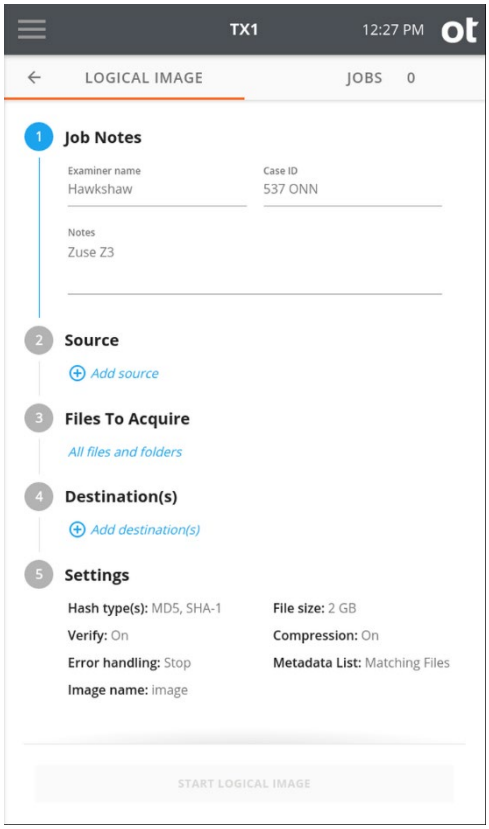

4. To select a filesystem from an available source drive, tap the 2 or Source heading. A list of attached drives will appear, with a filesystem summary tile shown under each drive for any filesystems recognized by TX1. Note that network shares can also be used as logical image job sources. To make a network share visible in the **Select a source** list, tap the orange plus button at the top right of the modal and follow the share mounting workflow. Source network shares can be mounted from the Sources button on the main screen as well.

   Once the source drive/share list is set, simply tap the desired filesystem tile to select it as the source for the logical image job. A green check confirms your selection and the source selection modal will auto-close. If a different source filesystem is desired (or to verify the details of the selected filesystem), simply go back into the **Select a source** screen by tapping the **2** or **Source** heading from the logical image job setup stepper. Close the modal by tapping the X in the upper right corner or by tapping outside of the modal.

Note: Within any screen displaying a list of drives, you can tap the options icon (three vertical dots) located on the right side of the drive tile to see more drive detail and access any available media utilities.

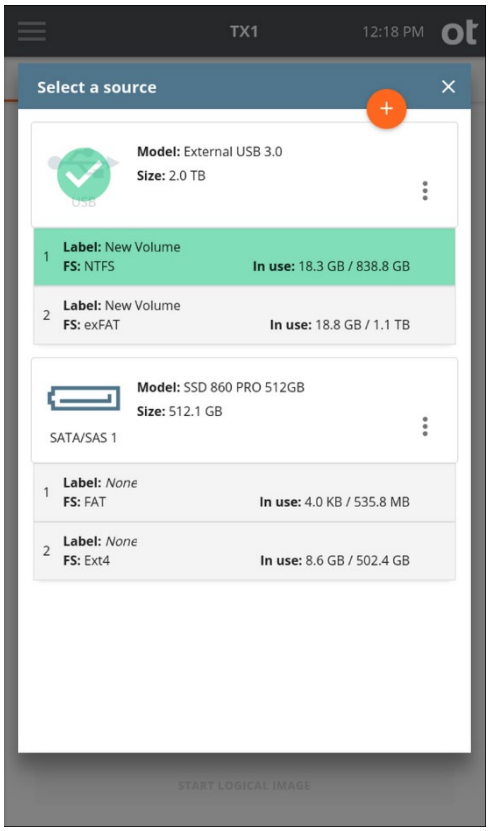

Note: Unlike a physical duplication job, the option of shelving a source drive DCO/AMA (removing it and then re-applying it at the end of the job) does not exist in logical imaging. The existence of a DCO or AMA will be obvious (per warnings in multiple locations), but the DCO/AMA will need to be permanently removed using the manual HPA/DCO/AMA Disable media utility before gaining access to all portions of the source media.

5. The next step is to determine which files and folders should be acquired. Start this process by tapping the **3** or **Files to Acquire** heading in the job setup stepper (resulting in the screen shown below). The default setting is to acquire all files and folders. Use the default setting if your job does not benefit from targeted down-selection of source files/folders and skip to the Destination(s) selection step below.

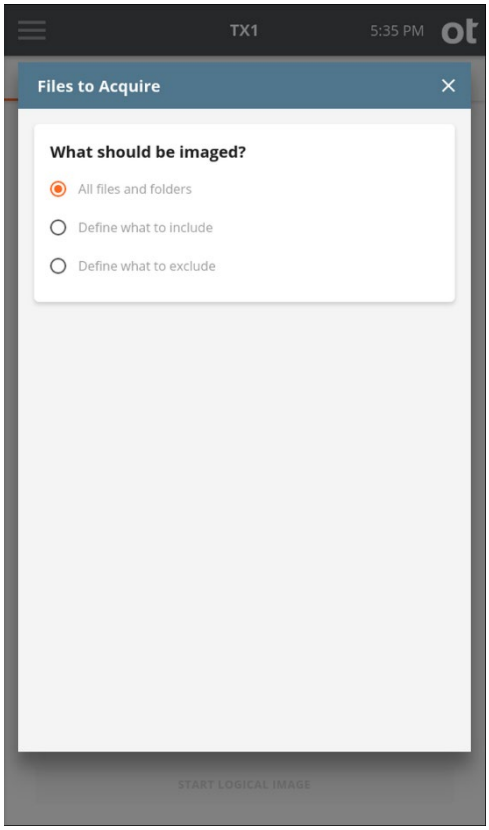

As shown in the screenshot above, the default setting of **All files and folders** is initially selected. If no source file or folder down-selection is desired, simply exit this step by closing the **Files to Acquire** modal window by tapping the X in the upper right corner or by tapping outside of the modal.

If down-selection of the source dataset is desired, there are two different starting points, as follows:

- **Define what to include** – This setting assumes that NO files and folders will be acquired, which requires defining which ones to INCLUDE.
- **Define what to exclude** – This setting assumes that ALL files and folders will be acquired, which requires defining which ones to EXCLUDE.

Note: The instructions in this section focus on the Basic mode of searching for files and folders to be acquired. This Basic mode search provides a powerful yet simple and straightforward way to focus on forensically valuable file-based evidence. See

"Advanced logical imaging setup" on page 125 for details on the optional Advanced mode search.

Regardless of whether you are including items in an empty dataset or excluding items from a full dataset, the same setup style is used to limit what is acquired, as covered in detail below.

6. Select files and folders to include/exclude.

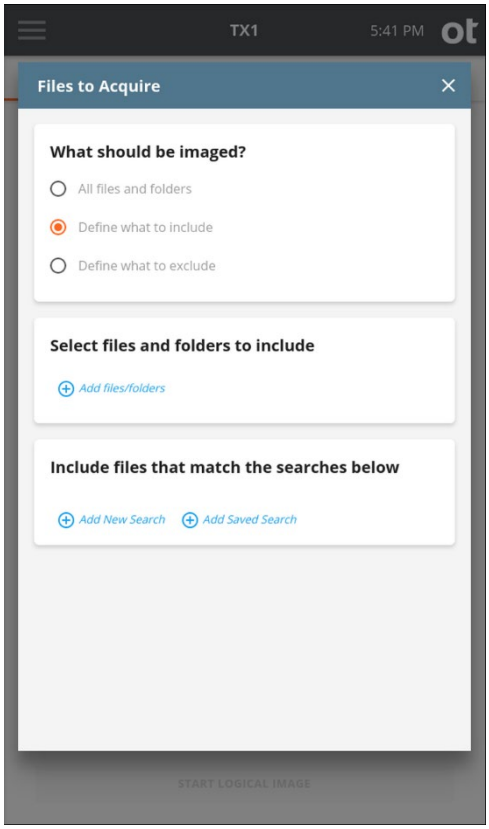

In the example shown above, **Define what to include** is selected, which means the initial acquisition dataset is empty. To manually select which files and folders to include, tap the blue circle/plus sign icon labeled **Add files/folders** in the **Select files and folders to include** section. This will launch a browse modal window which shows the entire contents of the selected source filesystem, as shown below.

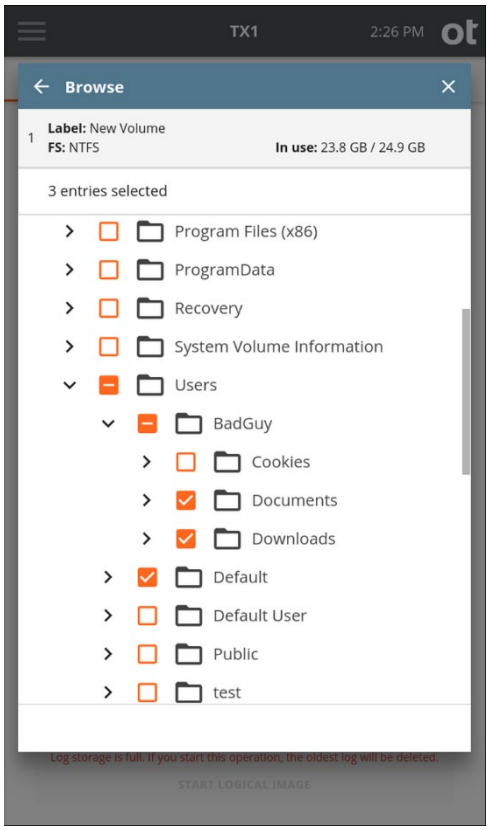

Individual files and/or folders can be manually selected simply by clicking on the orange box to the left of the desired item. In the example above, we have chosen to include `/Users/BadGuy/Documents`, `/Users/BadGuy/Downloads`, and `Users/Default`.

Note the following items related to manual file/folder selection in this browse modal:

- Selections are limited to 50 files/folders. Children files/folders of a selected parent folder do not count towards this limit.
- If a parent folder is selected, individual children cannot be deselected. First deselect the parent, then go in and select the desired children.

Notice that the total count of selected items is shown in the row between the filesystem information at the top and the directory tree at the bottom. This count reflects the manually selected items and does not include children of a selected parent folder.

For more details about the criteria by which you can include or exclude files for logical imaging, see “Include/exclude criteria” on page 117 and its subsections.

7. To select the destination drive(s) tap the **4** or **Destination(s)** heading. All available destination drives will be shown in a modal window, with a tile shown below each drive to show its recognized filesystem(s). Tap the desired destination filesystem(s)

(up to four) and the selected filesystem tile will show in green, with a checkbox added to its parent drive tile. Selecting a filesystem will also open a drawer under the filesystem tile, which shows the options for logical imaging destinations, as follows:

- **Job type:** This allows selection of the desired output files types. The main logical evidence file format for TX1 is lx01, which will contain all the data and metadata for every acquired file and folder. Additionally, TX1 offers the ability to generate a Metadata file in the common CSV (comma separated values) format, which contains all the available metadata for every acquired file and folder. Optionally, the metadata output file can be configured to contain all the metadata for all the files/folders on the source drive (whether the actual files/folders were acquired or not). That setting is made in the **Settings** area of the logical image job setup screen. The **Job type** setting options for the selected filesystem are: **Lx01**, **Lx01 + Metadata**, and **Metadata**. See “Source file metadata” on page 129 below for more information regarding captured metadata.
- **Directory:** This identifies in which directory the logical image job output files will be stored. The default directory will initially be shown, which can be changed by tapping on the orange **Directory** label to enter a **Browse** modal where you can select a different destination base directory, create and select a new directory, or delete a directory. Tap one or more of the four buttons (**Date and Time**, **Model**, **Serial**, or **Case ID**) under the directory path to add variables as names for a destination sub-directory. Each variable can only be selected once. Underscores are printed as separators between multiple variable names.

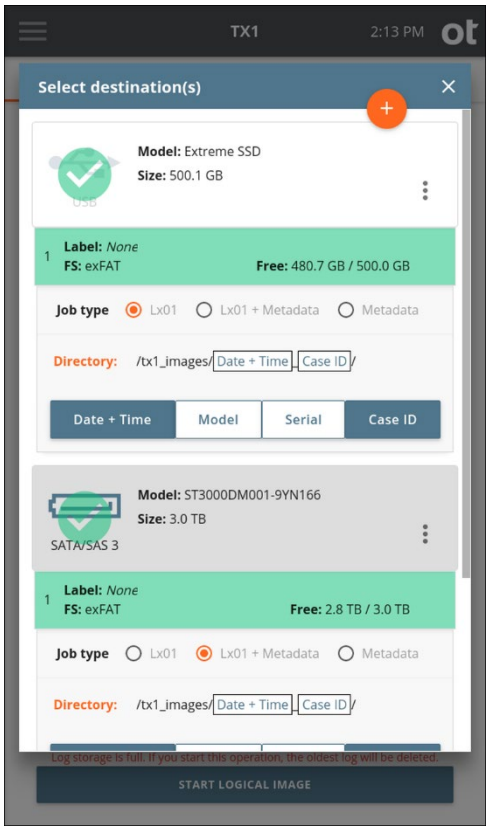

Destination drives with no recognized filesystems are grayed out, with a warning message stating no filesystem is available. Such a drive can be formatted through the media utilities available on the drive details screen, which can be accessed by tapping the additional options menu (three vertical dots) at the right side of the drive tile, or from the **Destinations** button on the **Home** screen.

Network shares can also be used as logical image job destinations. To make a network share visible in the **Select destinations** list, tap the orange plus button at the top right of the modal and follow the share mounting workflow. Destination network shares can be mounted from the **Destinations** button on the main screen as well.

8. To change the job settings, tap the **5** or **Settings** heading.

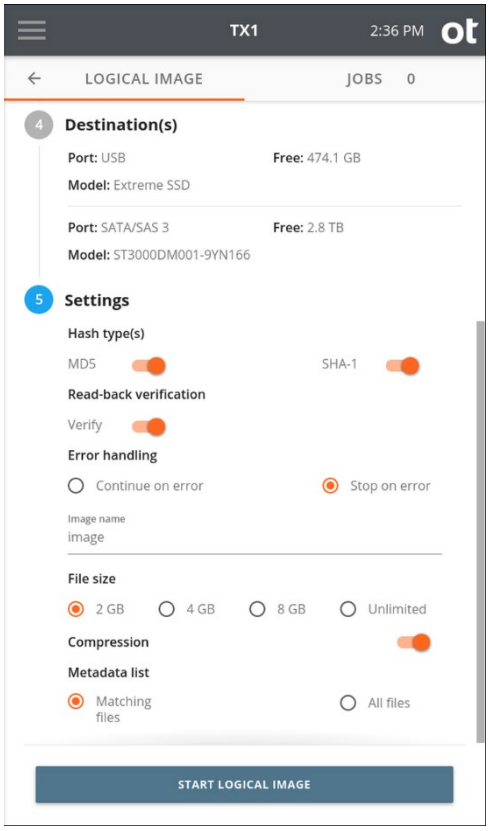

Select the desired **Hash type** for the logical image job - **MD5** and/or **SHA-1**. Note that hash values for source files will be calculated based on the chosen hash settings, even if no lx01 outputs are requested. In that case, the file data is still read to allow for hash calculation, and the file-based hash values are stored in the metadata output file.

Enable or disable **Read-back verification**. Read-back verification is not possible if no lx01 outputs have been configured or if no hashes have been selected.

Set the desired source read Error handling mechanism for the job – **Continue on error** or **Stop on error**. Note that read error handling in logical image jobs is distinct from that of physical image jobs in that retries are not allowed. This is because file read errors will typically not succeed after an initial failed read. Therefore, for logical image jobs, the only error handling setting decision to make is whether to continue with the job (skipping the unreadable file(s)) or to stop the job as soon as the first file read error is encountered.

Note: When a read of a given source file fails and the **Continue on error** setting is active, the unreadable file will not be acquired. In that case, the metadata output file will show an error condition in the entry for the unreadable file, and the lx01 will

indicate the error condition, which enables forensic analysis tools such as EnCase to indicate an error for the affected file(s).

The default output **Image name** is shown and can be changed by tapping the field and typing in the desired name.

The default image **File size** is shown and can be changed by tapping the desired size. Note that this setting is unavailable if no lx01 output has been requested.

Compression can be enabled/disabled. Note that this setting is unavailable if no lx01 output has been requested.

If a **Metadata list** was chosen as part of the destination settings, the following options for what to include in that metadata list will be shown at the bottom of the **Settings** section: **Matching files** or **All files**. If **Matching files** is selected, only the metadata for files that were acquired during the job will be included, whereas All files will obviously include a metadata entry for each file on the source, regardless of whether it was acquired or not. The latter setting may be useful in cases where time or other constraints only allowed partial source file gathering.

9. Once you are satisfied with all the logical image job settings, tap the **Start Logical Image** button.

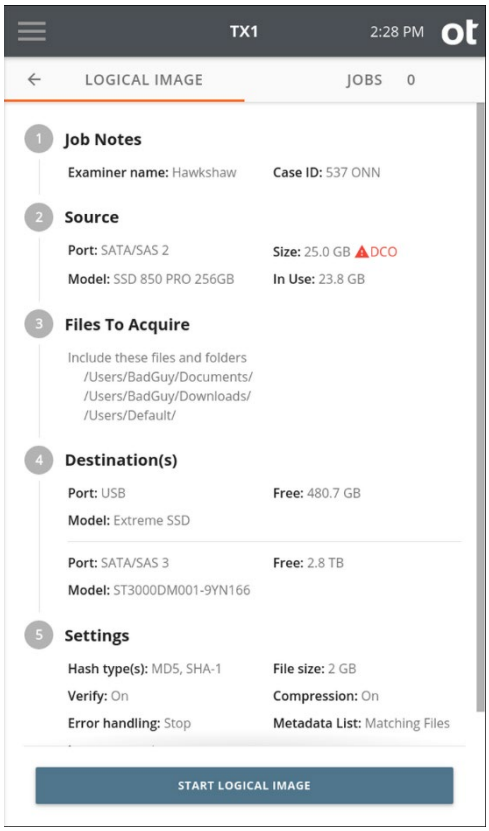

4.5.2 Include/exclude criteria

From the main **Files to Acquire** window, searches can be added that will allow for targeted acquisition of specific file/directory criteria. Select either **Add New Search** or **Add Saved Search** in the **Include files that match the searches below** box to specify a search or series of searches to apply to the logical imaging job.

While the screenshots and workflows below reflect the method of creating new searches during job setup, making use of the saved searches option is encouraged as an efficient time-saver. This is especially true if specific search types are commonly used for your logical imaging jobs. Before continuing with information about adding new searches, here are a few pointers related to the saved searches feature:

- New searches can be created/imported and saved searches can be edited from the **Defaults** setting area on the side navigation menu.
- During logical image job setup, new searches can be created, and previously saved searches can be used/edited.

- Searches within a given job setup can be a mix of new and saved searches. Simply tap the desired search entry method at the bottom of the last search criteria box (**Add New Search** or **Add Saved Search**) to add another search to the job.
- If a saved search is edited in the job setup area, those changes can optionally be saved back to the originally named search, saved as a newly named search (by editing the name field), or just used during the current job, without altering the original saved search that was added. To save the in-line changes, simply tap the **Save Changes** button at the bottom left of the search window. Please be aware that editing saved searches for a given job without saving the changes could lead to confusion about what the named, saved search means. For that reason, it is recommended to save any in-line changes back to the saved search that it started from or as a new saved search.
- Naming your searches provides a convenient way to identify the type of criteria used in the search, which is particularly useful for a list of saved searches. However, naming of searches is optional. If you do not enter a name for a given saved search, the default name of **Saved Search** will be used with a number at the end that increments for each newly saved, unnamed search.

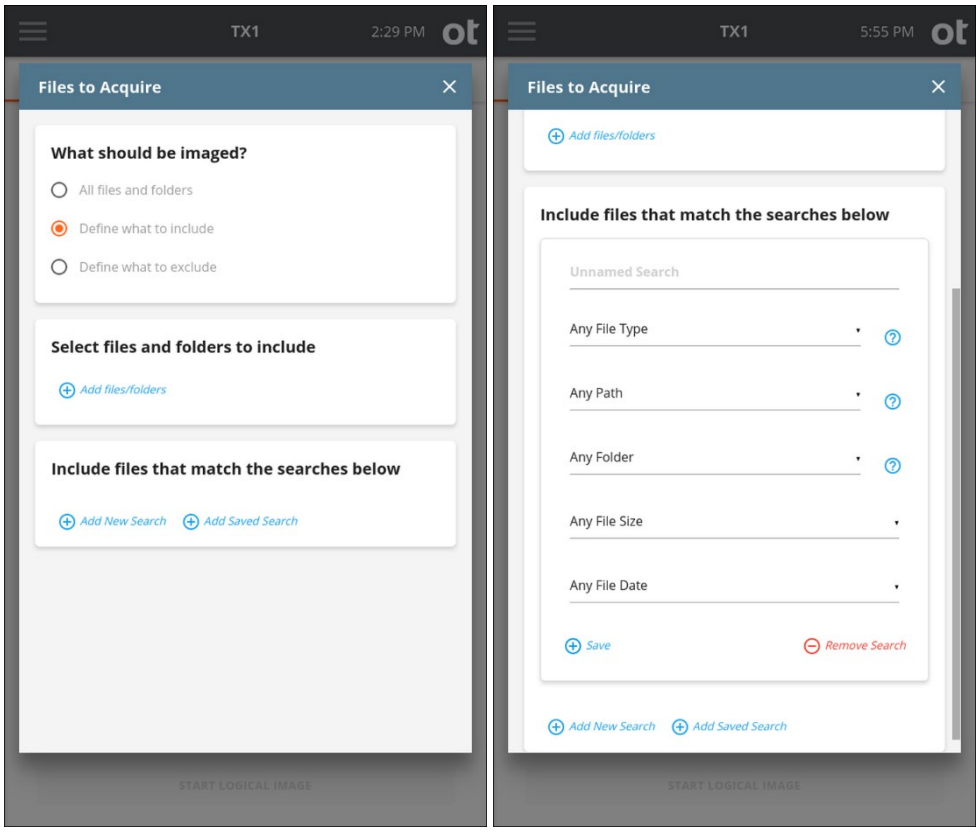

In the example on the right above, we have decided to add a new search using the **Include files that match the searches below** function. Again, since we started with

Define what to include in the top selection box, our initial acquisition dataset is empty. Any configured searches will potentially add to that empty acquisition dataset.

The file search options are covered in detail in the following sections. Before getting into those details, here is some general information regarding logical image searches:

- A search defines parameters for the kinds of files that are of interest to a forensic investigation. Adding searches will include/exclude every instance of that kind of file.
- Multiple searches can be configured for a given job (using the Add Search button at the bottom of the search setup window), but they are logically independent from each other. This means that if a search gets a hit on a file, it will be added to/removed from the acquisition dataset even if a subsequent search is configured to ignore that same kind of file.
- Within a single search box, all the criteria defined must match for a given file to be included in (or excluded from) the acquisition dataset.
- The first entry in each search setup box provides a name field for the search. The default name is **Unnamed Search**. Changing the name to something more specific may help when reviewing the summary of all searches in the logical image job setup screen or when viewing the forensic log associated with a logical image job.
- Each of the search parameter fields makes use of drop-down selection boxes to help guide the setup of each parameter.
- Text fields used for matching file names, file extensions, and file paths can use any Unicode characters TX1-supported filesystems might contain.

Note: Unicode is the standard for encoding visual text characters/symbols into digital values to allow computer systems to understand what characters/symbols are being referenced for proper display and processing. It was defined to allow digital encoding of special characters/symbols (including language specific accent glyphs, for example) that are not covered by the ASCII standard. TX1 uses UTF-8 (Unicode Transformation Format – 8 bit) to encode all possible characters/symbols used in all filesystems that TX1 supports and to use those encoded values when matching characters during logical image searches. For characters that are composed of multiple, distinct glyphs (for example the German umlaut over the letter A), TX1 uses the NFC (Normalization Function Composition) standard to normalize the various encoding methods into a common hex value for matching purposes. For more information, search for “UTF-8” and “Unicode equivalence” on the internet.

The various search parameter fields that TX1 supports are covered in the following sections.

4.5.2.1 File type

The **File type** search parameter restricts the search to apply only to files that match a list of file extensions. Each search can have any number of file type constraints. When multiple file type constraints are included in a given search, a match of any one of them will include/exclude a given file.

The following file type search parameters are available:

- Archives
- Databases
- Documents
- Emails
- Multimedia
- Pictures
- Custom

Each type other than **Custom** has a predefined list of extensions known to be associated with that type of file. The lists can be seen by tapping the blue help button (circle with question mark) to the right of the file type parameter field. All extensions associated with each file type are found in “File extensions” on page 137.

Using the **Custom** field allows for manual entry of any extension values to match against. The entry of custom extension values is done outside of the pull-down selection box.

The forensic log associated with the logical image job contains exactly what extensions were used for each search, along with which selections were used to create that list.

Note: Searching by file type does not use file signature analysis to determine what type a file is. Only the file name extension is used to determine a match. If file signature analysis is required for a given job, a physical image should be made (possibly in addition to a logical image) to ensure all source data is available for use with external forensic data analysis tools, such as EnCase Forensic.

4.5.2.2 Path

The **Path** search parameter restricts the search to apply only to files with a specific, user-defined string in either the filename or directory path. A field for entering the desired search string appears after selecting one of the options. The Path search parameter options are as follows:

- **Filename Contains** – restricts the search to only apply to files that contain the given string somewhere in the filename.
- **Path Contains** – restricts the search to only apply to files that contain the search string somewhere in the full path (directory or filename).

Wildcards can be used to search for only a portion of a file or folder in a path-based search. The available wildcards and examples for each are shown in the table below. Note that this wildcard search is based on the Linux glob search rules, which can be referenced online for additional information.

Note: While wildcard searching can be a powerful tool to quickly search for files of interest, it is critical that these wildcard rules are fully understood before making use of them in an actual case job. Misunderstanding exactly how a given wildcard rule will function could result in inaccurate search results and missed evidence.

Wildcard Character	Matching Rule	Examples
*	Matches any number of any characters (including none).	**Law*** Matches: Law, Laws, Lawyer Does not match: NoLaw, La, aw ***Law*** Matches: Law, NoLaw, Lawyer Does not match: La, aw
?	Matches any single character.	**?it** Matches: Bit, bit, Sit, sit Does not match: it, slit, bits
[abc]	Matches one instance of any of the characters between the brackets.	**[SB]it** Matches: Sit, Bit Does not match: sit, bit, it **Note:** The literal bracket characters (‘[‘ or ‘]’) can be matched as part of the path/file name if they are the first character inside the brackets, such as [[abc].
[!abc]	Matches one instance of any of the characters that are not what is between the ‘!’ and the ‘]’.	**[!S]it** Matches: sit, Bit, bit Does not match: Sit
-	Used within brackets, hyphens denote a range of characters to be matched.	**money[1-5]; equivalent to money[12345]** Matches: money1, money2, money5 Does not match: money, money6 **Note:** To match a hyphen itself within brackets, it must be the first or last item inside the brackets.

4.5.2.3 Folder

The **Folder** search parameter restricts the search to apply only to a specific folder or kinds of folders. The following folder search parameters are available:

- User Folders
- Operating System Folders
- Non-Operating System Folders
- Custom Folder

The **User Folders** and **Operating System Folders** items have a predefined list of folders known to be associated with them, which can be seen by tapping the help button to the right of the folder parameter field. Selecting either of those pre-defined values will limit the search to only those folders. Selecting the **Non-Operating System Folders** value will limit the search to any folders that are not in the predefined **Operating System Folders** list. All predefined paths associated with the **User Folders** and **Operating System Folders** are found in “Folders” on page 128.

The **Custom Folder** setting allows you to browse the source for a folder to match. After selecting **Custom Folder** from the drop-down list, simply tap the **Select Folder** button to launch a browse modal and select the desired folder.

Note: The full, absolute path name must be specified for **Custom Folder** searches. The addition of the full absolute path is automatic when adding a **Custom Folder** search during logical imaging job setup; however, when adding a new search in the **Default** settings area (via the side navigation menu), manual entry of the desired **Custom Folder** path is required. For manual entry, you must use the forward slash (/) to specify the root directory and spell out the complete path to the custom folder name (for example, /users/suspectname/pics). Entering only a folder name without the absolute path nomenclature will result in missed evidence. If you want to locate a folder name without regard for the absolute path on the drive, use the **Path Contains** search option and enter the desired folder name (for example, suspectname/pics).

4.5.2.4 File size

The **File Size** search parameter restricts the search to apply to only files in a certain range of sizes. The following file size search parameters are available:

- File Sizes >=
- File Sizes <=
- File Sizes in Range

File Sizes >= lets you specify a file size in bytes, KB, MB, or GB and match only files greater than or equal to the specified size.

File Sizes <= lets you specify a file size in bytes, KB, MB, or GB and match only files less than or equal to the specified size.

File Sizes in Range lets you specify two file sizes, and match only files with size in between the two specified sizes.

4.5.2.5 File date

The **File Date** search parameter restricts the search to apply to only files in timestamp ranges, as follows:

- File Dates >=
- File Dates <=
- File Dates in Range

File Dates >= lets you specify a date and only match files with one or more timestamps on or after the given date.

File Dates <= lets you specify a date and will only match files with one or more timestamps on or before the given date.

File Dates in Range lets you specify two dates and will only match files with one or more timestamps on or between the two given dates.

Dates are entered by typing them into the text boxes in YYY-MM-DD format. Dates will match the rule if any of the TX1 supported timestamps for that filesystem match the **File Date** setting.

Note: Some filesystem types, particularly FAT, do not support time zone independent timestamps. There is no way to determine the time zone used for FAT time stamps created by Windows. TX1 treats these timestamps as if created with a UTC+0000 timestamp.

4.5.3 About the logical imaging process

4.5.3.1 Logical image job status

Once a logical image job has been started, the **Job Status** screen will automatically be displayed, as shown below.

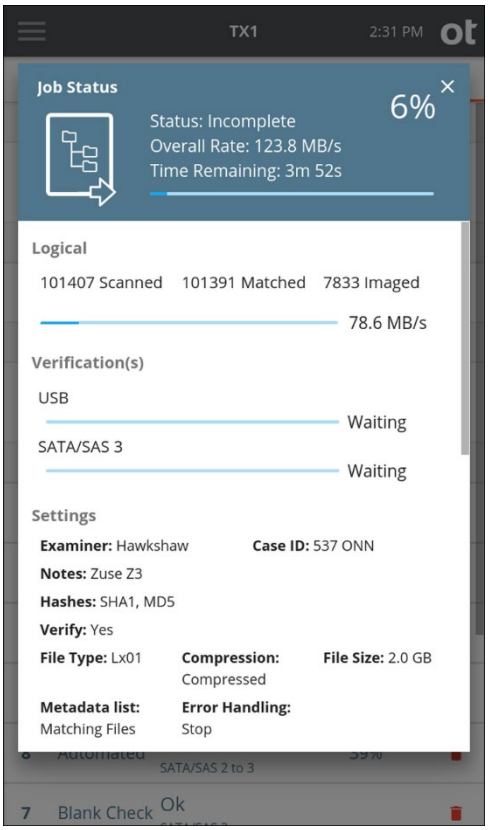

This status screen is similar to other TX1 job types, with the following notable differences:

- While the operation is scanning for more files to acquire, the progress bar is displayed as an indeterminate bar (throbbing/pulsing bar with no data rate displayed). TX1 does not know how many bytes it needs to acquire until the scan is complete.
- Three new status fields show the live status of the scanned files:
    - **Scanned** shows the number of files the job has checked so far to see if it should include them.
    - **Matched** shows the number of files that will be acquired by this job out of the number of files scanned.
    - **Imaged** shows the number of files that have been fully acquired out of the number of files matched.
- The **Settings** section includes the settings specific to logical imaging that have been configured for the active job.
- The **Included Files** section shows a text summary of the rules used by the active job to determine which files should be acquired.

Note: Filesystem read errors encountered during logical imaging jobs may result in unpredictable acquisition behavior. When they occur, such errors are indicated by a red warning message at the top of the Logical section of the status screen. You will also see

non-matching values in the **Matched** and **Imaged** fields on the **Job Status** screen at the end of the job and errors noted in the job's metadata file. If you suspect drive/filesystem read errors during a logical imaging job, we recommend that you clone or physically image the drive (`e01, ex01, dd, dmg`) instead of trying to do a logical image. In addition to a physical image (or if a physical image is not possible), try logically imaging the source in multiple, smaller jobs instead of trying to gather all files/folders in one job. If the errors happen to be in less forensically interesting areas of the filesystem, this could result in a more valuable file/folder acquisition set.

4.5.3.2 Files created during logical imaging

When performing a logical image on TX1, multiple different files may be output to each destination depending on the job configuration, as follows:

- `{image_name}.log` contains the forensic log of the logical imaging operation.
- `{image_name}.Lx01, {image_name}.Lx02, …` are the forensic evidence files for the operation. They contain all the data and metadata for each file and folder acquired.
- `{image_name}.csv` is a comma separated value store of all the metadata for every file and folder acquired. Optionally, this file also contains all the metadata for files and folders that were not acquired. This type of file can easily be imported into many common data processing applications such as Microsoft Excel. CSV file data contents and format information can be found at "Source file metadata" on page 139.
- `{image_name}.tx1_packed_log` contains a TX1 readable copy of the forensic log that can be used for later standalone verification of the lx01 file set.

All the above output files are generated when a given destination is configured to be the **Lx01 + Metadata** job type. No CSV metadata file is generated for the Lx01 job type. No `lx01` file set or `.tx1_packed_log` file is generated for the **Metadata** job type.

If all destinations are configured to be the **Metadata** job type, and no hashes are configured, the file data for each file will not be read at all. This allows an investigator to quickly create a record of all source file metadata.

4.5.3.3 Logical image verification

Verification of lx01 files differs from verification of physical imaging operations because, in an lx01 file, there is no overall hash. Each file's data stored in the lx01 has an associated hash that was calculated during the original acquisition. The logical imaging verification function reads back the file data from the lx01 on the destination, calculates a new hash value for each file, and compares that hash value to the originally stored hash value. A failure of any one file to match the original hash value will result in a verification/job failure.

4.5.3.4 Advanced logical imaging setup

The logical imaging setup of TX1 can be switched to **Advanced Logical Imaging** Setup mode. By enabling this setting in the **Default** settings screen, two additional search setup options are activated. If **Advanced Logical Imaging Setup** is enabled but none of its

additional features are used, TX1 behaves the same as **Basic** mode except for minor changes to log text.

Per Search Include/Exclude switches

In **Basic Logical Imaging Setup** mode, all searches either include files or all searches exclude files. With **Advanced Logical Imaging Setup**, this can optionally be modified on a per search basis. This allows setup of complicated logical image jobs that specify searches in precedence over other searches and over file/folder selection. Note that this mode can be convoluted and confusing from a search logic perspective, which is why the **Basic** search mode is used as the default.

Whether or not to include or exclude a given file/folder is determined by going down the list of searches first. If a given search matches the file/folder, that file/folder is included or excluded based upon the setting for that search. The earlier search always takes precedence in this way.

If no search matches a file/folder, then whether that particular file/folder is included or excluded is decided by the manual selection of files and folders.

- Define what to include – all files and folders not matching Include searches and not manually selected are not included.
- Define what to exclude – all files and folders not matching Exclude searches and not manually selected are included.

Note that new searches will use the top level include/exclude setting by default.

Invert switches for individual search parameters

In **Advanced Logical Imaging Setup** mode, each search parameter has an invert toggle. This lets the parameter match the opposite set of files from what it normally matches for that parameter.

For example, if you create a search that matches **Archives** and **File Sizes in Range** 1 MB to 100 MB, two switches will appear next to those two settings. If you enable the invert switch next to **Archives**, the rule now matches any file that is not an Archive. If you enable the switch next to the File Sizes range, the rule will now match any file outside the range 1MB to 100MB. You can switch one or both conditions to match the kind of files you are searching for.

In general, **Basic** search mode is recommended unless there is a compelling reason to switch to **Advanced Logical Imaging Setup** mode to target specific source evidence.

4.5.4 File extensions

During logical imaging, TX1 can search for file types with any of the following extensions.

File Type	File Extension
Archives	“7z”, “7zip” “zom”, “apk”, “xxe”, “uug”, “mim”, “tz”, “arj”, “zsm”, “zze”, “boo”, “bkp”, “bak”, “sav”, “bac”, “ful”, “bag”, “zso”, “bplist”, “bhx”, “mhk”, “bz”, “bz2”, “ckit”, “boz”, “ish”, “rar”, “r01”, “jar”, “cru”, “cif”, “gnu”, “gz”, “tgz”, “gzip”, “ha”, “hap”, “lzs”, “lha”, “arc”, “lzh”, “pak”, “hqx”, “image”, “sparseimage”, “marc”, “cab”, “b64”, “gho”, “rpm”, “rzip”, “rz”, “sea”, “sdn”, “stf”, “squashfs”, “sqz”, “sit”, “sitx”, “td0”, “ufa”, “bar”, “cpio”, “taz”, “z”, “tar”, “tgz”, “uu”, “uue”, “xcr”, “yz”, “yc”, “zdg”, “zip”, “zoo”
Databases	“asl”, “cdb”, “ntx”, “csv”, “dbf”, “idx”, “ind”, “dbk”, “bch”, “db2”, “db3”, “ndx”, “cvt”, “crp”, “mdx”, “$db”, “db$”, “fp”, “pjx”, “mda”, “mdt”, “mdn”, “mdb”, “mde”, “ldb”, “mar”, “mdw”, “mny”, “wdb”, “mlb”, “odb”, “sqlite”, “sdb”, “db”, “sqlite3”, “sqlite”, “kexi”, “shm”, “db-wal”, “sqlite-wal”, “cix”, “dba”
Documents	“cch”, “cfl”, “cht”, “ch3”, “aft”, “abc”, “xlc”, “gra”, “opx”, “adt”, “smf”, “pfc”, “att”, “bfx”, “brk”, “ezf”, “can”, “cci”, “ccitt”, “cpf”, “cfp”, “ef3”, “fcx”, “ftf”, “f96”, “dxn”, “gam”, “cg3”, “fax”, “tbf”, “jet”, “bk”, “kfx”, “awd”, “oaz”, “prd”, “tef”, “sci”, “tri”, “wpf”, “q”, “cvp”, “mif”, “zvd”, “key”, “sld”, “cpp”, “ch4”, “crp”, “cpr”, “pps”, “ppt”, “pot”, “pptx”, “odp”, “otp”, “sxi”, “sti”, “numbers”, “slk”, “lss”, “wks”, “123”, “gph”, “wk4”, “xlk”, “xls”, “xlsx”, “xlt”, “xlb”, “xlw”, “ods”, “ots”, “sxc”, “stc”, “wkq”, “wrk”, “rpc”, “rpn”, “ltr”, “fdf”, “pdf”, “sam”, “pages”, “asc”, “etx”, “bib”, “asp”, “aspx”, “ascx”, “bbs”, “big5”, “big”, “chi”, “cwk”, “cws”, “chm”, “doc”, “sbj”, “cag”, “xla”, “mdz”, “wiz”, “mcc”, “oft”, “msc”, “dat”, “ppt”, “xls”, “pot”, “pub”, “cbt”, “diz”, “efx”, “evy”, “xml”, “faq”, “hwp”, “aw”, “oth”, “htm”, “html”, “asp”, “aspx”, “log”, “wke”, “wk1”, “wks”, “fm3”, “wk3”, “fmt”, “ami”, “adx”, “ntf”, “id”, “ide”, “mht”, “mhtml”, “obt”, “cue”, “ybk”, “obd”, “cpe”, “cov”, “docx”, “doc”, “wbk”, “asd”, “dot”, “wps”, “new”, “odg”, “otg”, “odf”, “odt”, “ott”, “old”, “cas”, “cbk”, “sxd”, “std”, “sxm”, “sxw”, “stw”, “pub”, “cat”, “qdf”, “qsd”, “abd”, “1st”, “rpt”, “rtf”, “shtml”, “set”, “txt”, “c00”, “chp”, “vs?”, “htm”, “html”, “wri”, “wp5”, “bk!”, “wpd”, “tv1”, “tv2”, “tv3”, “tv4”, “tv5”, “tv6”, “tv7”, “tv8”, “tv9”, “bv1”, “bv2”, “bv3”, “bv4”, “bv5”, “bv6”, “bv7”, “bv8”, “bv9”, “wpt”, “blk”, “bk1”, “bk2”, “bk3”, “bk4”, “bk5”, “bk6”, “bk7”, “bk8”, “bk9”, “wp”, “xml”, “man”, “manifest”, “config”, “cfg”, “xsd”, “resx”, “msc”, “admx”, “slt”, “xsl”, “xrm-ms”, “mum”, “dtd”, “xsl”
Emails	“pfc”, “org”, “emlx”, “cca”, “eml”, “ccm”, “nsf”, “mbox”, “edb”, “msg”, “nws”, “pab”, “sch”, “sc2”, “scd”, “que”, “dbx”, “pst”, “wab”, “crd”

File Type	File Extension
Multimedia	“bnk”, “rol”, “amr”, “amf”, “aif”, “aiff”, “avr”, “cda”, “aifc”, “cdm”, “idf”, “aac”, “pcm”, “ra”, “ram”, “wav”, “wma”, “zad”, “asf”, “awm”, “awa”, “divx”, “vob”, “f4p”, “f4v”, “swf”, “dvr-ms”, “mp4”, “asr”, “3g2”, “wm”, “wmv”, “filmstrip”, “flc”, “m4r”, “m4p”, “qtm”, “ic1”, “ic2”, “ic3”, “snd”, “avi”, “voc”, “dvm”, “flv”, “lza”, “mmm”, “mp3”, “m3d”, “mpg”, “mpeg”, “mps”, “mpv”, “mpa”, “mp2”, “13”, “m1s”, “m1v”, “m1a”, “m2s”, “m2v”, “m2a”, “m4a”, “m4b”, “m4v”, “mp4”, “ani”, “avi”, “wav”, “rmi”, “idf”, “mtm”, “midi”, “mid”, “rmi”, “qtch”, “moov”, “movie”, “mov”, “mov”, “qt”, “rm”, “smjpeg”, “ibk”, “3gp”, “asf”, “wma”, “wmf”, “wmv”
Pictures	“3dmf”, “l64”, “n64”, “sbj”, “acb”, “ais”, “dxr”, “eps”, “ai”, “pdd”, “psb”, “psd”, “psd”, “pdd”, “ps”, “amff”, “sdw”, “art”, “ind”, “cil”, “b&w”, “dxf”, “dwg”, “3ds”, “flc”, “fli”, “cel”, “bif”, “b8”, “bit”, “bob”, “bgi”, “cvg”, “ccrf”, “cob”, “scn”, “cr2”, “crw”, “cvs”, “cam”, “dpx”, “clp”, “cps”, “rix”, “scd”, “sce”, “scr”, “scp”, “scg”, “scu”, “sci”, “sck”, “scq”, “scl”, “scf”, “scn”, “sco”, “scz”, “cpi”, “c64”, “ce1”, “ce2”, “gmf”, “csb”, “cpx”, “csi”, “jff”, “cmf”, “cpt”, “map”, “cmx”, “cmv”, “bmf”, “xar”, “abk”, “cdr”, “pat”, “cdt”, “cdx”, “ct”, “cur”, “dif”, “lbm”, “bbm”, “anm”, “gem”, “dvi”, “ce”, “cft”, “cut”, “dhp”, “pal”, “ed5”, “ed6”, “emf”, “qfx”, “fpx”, “fif”, “fmv”, “img”, “gif”, “gl”, “pdw”, “org”, “hpc”, “hpg”, “pcl”, “hgl”, “gca”, “iax”, “ica”, “sbp”, “ilbm”, “rlc”, “infini-d”, “igf”, “pix”, “iff”, “sgo”, “je”, “imj”, “jfif”, “jif”, “jpg”, “jpeg”, “jpe”, “jtf”, “cpr”, “kdc”, “dcs”, “xif”, “pcd”, “kiz”, “kqp”, “icns”, “pict”, “pct”, “mac”, “mnd”, “afw”, “af2”, “af3”, “ds4”, “dsr”, “qsf”, “dsx”, “dsf”, “dst”, “pp5”, “pp4”, “s3d”, “drw”, “sg”, “sep”, “sp”, “bmp”, “dib”, “mic”, “cag”, “msp”, “db”, “mgl”, “mpw”, “mpf”, “mri”, “mng”, “dcx”, “gal”, “mpt”, “nan”, “nif”, “nef”, “bga”, “psp”, “max”, “pcx”, “cdi”, “abm”, “pic”, “pxr”, “pbm”, “pnm”, “ppm”, “pgm”, “pgm”, “png”, “ppm”, “epi”, “j6i”, “pig”, “rle”, “sct”, “sid”, “pax”, “bw”, “shg”, “rgb”, “sgi”, “slb”, “sld”, “ssk”, “rix”, “arw”, “3d2”, “sod”, “ras”, “tga”, “p10”, “thn”, “tif”, “tiff”, “raw”, “cgm”, “cbd”, “dem”, “wvl”, “wi”, “ged”, “wbf”, “wmf”, “ico”, “spl”, “wpg”, “nff”, “bm”, “cbm”, “xpm”, “xwd”, “xbm”, “yuv”, “yuv3”, “zeiss”, “zgm”, “pcc”, “heif”, “heifs”, “heic”, “heics”, “avci”, “avcs”, “avif”, “avifs”

4.5.5 Folders

During logical imaging, TX1 can use the following predefined paths associated with User and Operating System folders.

Folder Type	Folder Name
User Folders	“/Users/”, “/Documents And Settings/”, “/WinNT/Profiles/”, “/var/users/”, “/users/”, “/u01/”, “/user/”, “/home/”, “/root/”, “/export/home/”

Folder Type	Folder Name
Operating System Folders	“/Windows/”, “/WinNT/”, “/System/”, “/Program Files/”, “/Program Files (x86)/”, “/ProgramData/”, “/Applications/”, “/bin/”, “/dev/”, “/etc/”, “/sbin/”, “/usr/”, “/boot/”, “/lib/”, “/proc/”, “/sys/”, “/unix/”

4.5.6 Source file metadata

Logical imaging with TX1 includes source file metadata in the csv output file.

Column	Content
Path	Contains the full, filesystem-relative path for this entry. Example: /users/charles/pictures.
Type	Either contains “Directory”, “Symlink”, or “File”, depending on what kind of entry this row represents.
Filesize	The file size, in bytes, of the entry. This field is empty for directories.
Creation Date	The IS0 8601 UTC date/time string for the creation date of this entry. This field is empty if the creation date is unavailable.
Accessed Date	The IS0 8601 UTC date/time string for the accessed date of this entry. This field is empty if the accessed date is unavailable.
Modified Date	The IS0 8601 UTC date/time string for the modified date of this entry. This field is empty if the modified date is unavailable.
Written Date	The IS0 8601 UTC date/time string for the written date of this entry. This field is empty if the written date is unavailable.
MD5 Hash	The MD5 Hash of the entry. This field is empty for directories. It is also empty if no MD5 hash was calculated, no MD5 hash was configured, or the entry did not match the rules for acquisition.
SHA1 Hash	The SHA1 Hash of the entry. This field is empty for directories. It is also empty if no SHA1 hash was calculated, no SHA1 hash was configured, or the entry did not match the rules for acquisition.

Column	Content
File Status	OK if there were no problems reading file data/metadata. ERRORS if there were errors reading file data and/or metadata. This field is empty for directories.
Matched Rules	“Y” if the file matched the acquisition’s rules for inclusion.

4.6 Verifying

The standalone **Verify** function verifies the integrity of an existing image file by reading back the data from the image file, calculating a hash value of that data, and then comparing that calculated hash value with the value of the original acquisition hash.

Note that, while the same **Verify** function can be used for standalone verification of physical and logical images, the underlying mechanism is different. This is because physical images contain whole disk acquisition hash values and logical images contain file-based acquisition hash values. No difference will be noticed during the verification job itself, but the source image type will make a difference in how the results are reported. For a physical image verification job, the drive level readback hash values will be reported in the forensic log. For a logical image verification job, a simple pass/fail indication will be reported in the forensic log, which indicates that all the file-based acquisition hashes matched the readback hash values. If any individual file in a logical image file fails to verify, the entire verification job will show as failed.

To verify an image file on a destination drive:

1. From the **Home** screen, tap the **Verify** button.

2. Enter **Job notes** and select a **Destination** drive.

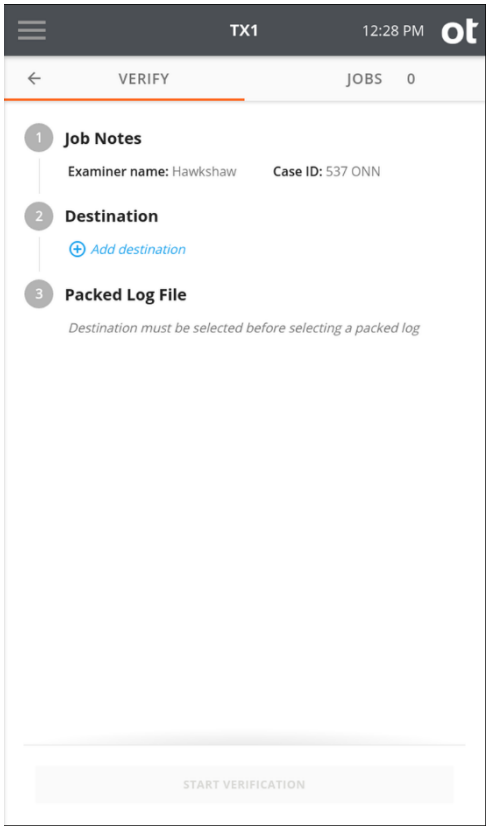

3. Select a **Packed log file**. Browse the destination and locate an existing TX1 packed log file.

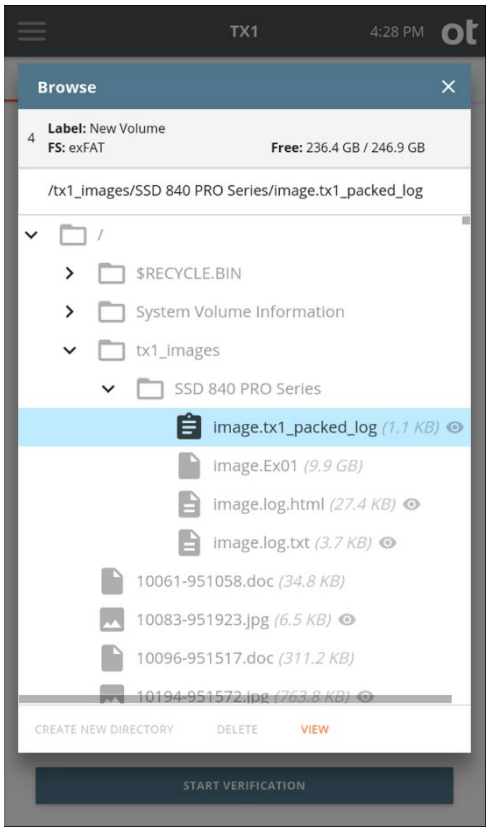

Note: The packed log files will always appear at the top of the file list in a given source folder when browsing. This provides easy access to these types of files in situations where there are many segment files.

4. Tap the **Start Verification** button at the bottom of the screen.

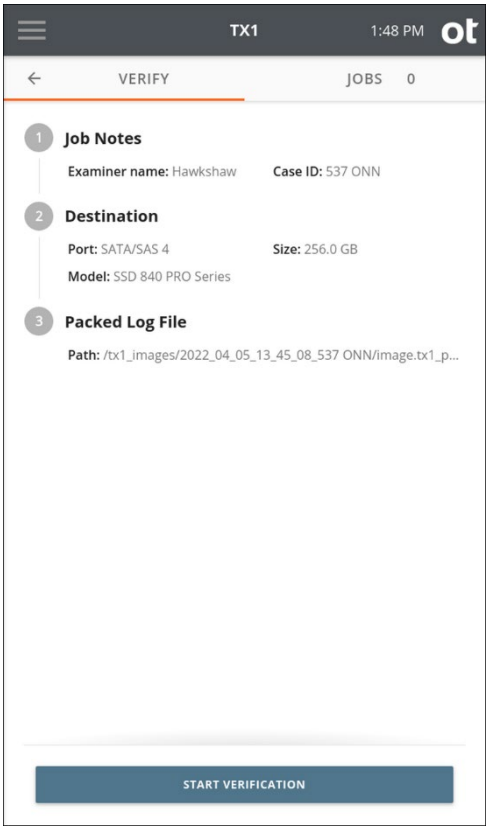

The verification process begins. A **Job Status** modal displays the verification status.

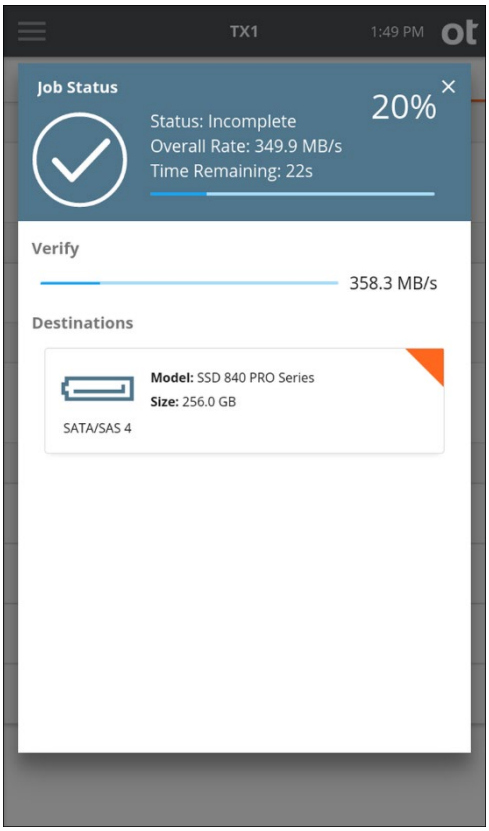

5. To cancel the **Verify** operation, tap the **Cancel** button from the **Jobs** summary screen.

   When the **Verify** operation is complete, the results are displayed on a final **Job Status** screen. The log for the completed job can be easily viewed by tapping on the **View Log** link on the right side of the top status bar or from the side navigation menu.

4.7 Browsing

The **Browse** function provides an easy way to view the contents of a recognized filesystem on any mounted drive, whether it is connected locally or via the network interface (iSCSI or CIFS). Simply tap the **Browse** button on the **Home** screen and select the desired drive/filesystem. The Browse operation is also accessible from the **Media Utilities** list in the drive details screen and from the filesystem details box within the **Content Breakdown** media utility. A sample browse screen is shown below.

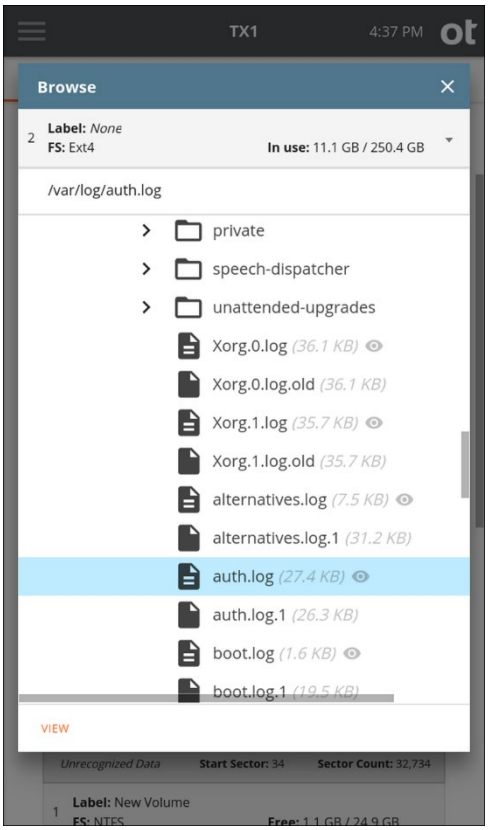

In the browser portion of the window, you can scroll up and down the list of folders/files and tap individual folders to drill down to the desired level to expose the names of individual files located on the drive. The size of each file is shown at the end of the filename. Many users will find this utility helpful when attempting to triage a large evidence set and determine the priority by which each drive should be imaged, or when checking the contents of a destination drive to free up space by deleting unneeded directories and files.

4.7.1 Viewing text and image files

TX1 provides the ability to view certain text and image file types. This is a valuable feature that will allow for on-the-fly analysis/triage of supported file types which could help home in on valuable evidence before deciding whether to acquire a given drive. Files that can be viewed directly on TX1 are indicated on the **Browse** screen with an eye icon at the end of the filename/size information. Additionally, viewable files show an alternate leading icon before the filename, as follows: folded paper icon with lines on it for text files; and a mountain range icon for image files, as can be seen in the sample filesystem browse view above. Selecting the desired viewable file will activate the **View** button at the bottom of the **Browse** screen. Simply tap the **View** button to see the text or image file directly on TX1. The screenshot below shows a sample image file displayed on TX1.

TX1-generated packed log files (extension “tx1_packed_log”) can also be viewed directly on TX1. These files are used as input for **Restore** and **Verify** jobs. Being able to view these files before starting one of those jobs can help ensure the desired file is selected.

Note: Viewing large text files (larger than 256 KB) results in undesirable screen update effects while scrolling through the file on TX1. For that reason, only the first 256 KB of data is shown when viewing text files directly on TX1. Files larger than 256 KB are truncated with a message at the end of the file in the viewing window that informs of the truncation. Download the file (using the remote web interface) or create an lx01 image for use in a forensic analysis tool such as EnCase Forensic to see the full contents of any large text files.

TX1 determines which text and image files are viewable by file extension only. The following file extensions are viewable on TX1:

- **Text type:** .bat, .c, .conf, .csv, .h, .htm, .html, .ini, .js, .json, .log, .nfo, .py, .readme, .sh, .text, .tsv, .txt, .xml
- **Image type:** .apng, .bmp, .gif, .ico, .cur, .jpg, .jpeg, .jfif, .pjpeg, .pjp, .png, .svg, .webp
- **TX1 packed log type:** “tx1_packed_log”

While it is possible to view the text version of a TX1 forensic log file via the **Browse** feature, for security reasons, the HTML log files will appear as raw HTML text. In order to

view these HTML logs in their styled format, please use the Logs menu item on the side navigation bar.

4.8 Restoring

The **Restore** function allows for recreation of the original drive format from a previously created TX1 image file. The uses for this feature are varied but include the ability to use a restored drive as a system boot disk and to simply create an archival copy of the evidence in its original format for future case reference.

The **Restore** function works with all physical duplication image file types (`e01, ex01, dd, dmg`) It does not support restoration from a logical image file set (`lx01`).

Note that, at the beginning of a Restore job, TX1 prepares the destination drive by wiping sectors 0, 1, and end-of-drive minus 1. This ensures there is no stale partition table data on the drive which reduces the possibility of drive detection issues at the end of the job.

Note: Because partition table information is relative to the sector size of the source drive, restoring to a destination drive with a different sector size is not allowed. TX1 will detect this sector size mismatch issue and warn the user. This condition will need to be rectified before the Restore job can be started.

To restore a drive from an image file:

1. Attach the desired source and destination media to TX1.

   Note: The **Restore** function requires a Tableau packed log file as input, so ensure that the desired packed log file is on a source drive (local or network) before starting this operation.

2. From the **Home** screen, tap the **Restore** button.

3. Enter **Job notes**, select a **Source** drive, and then select a **Packed log file** by browsing the source and selecting the appropriate TX1 packed log file.

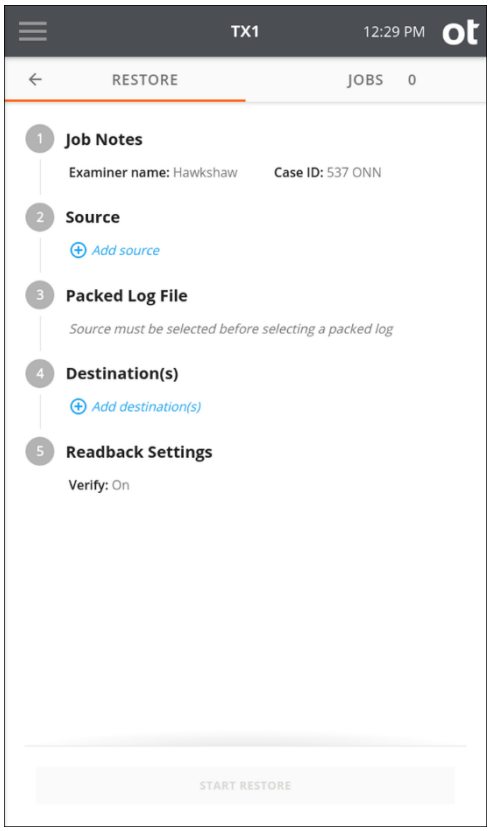

Note: The packed log files will always appear at the top of the file list in a given source folder when browsing. This provides easy access to these types of files in situations where there are many segment files.

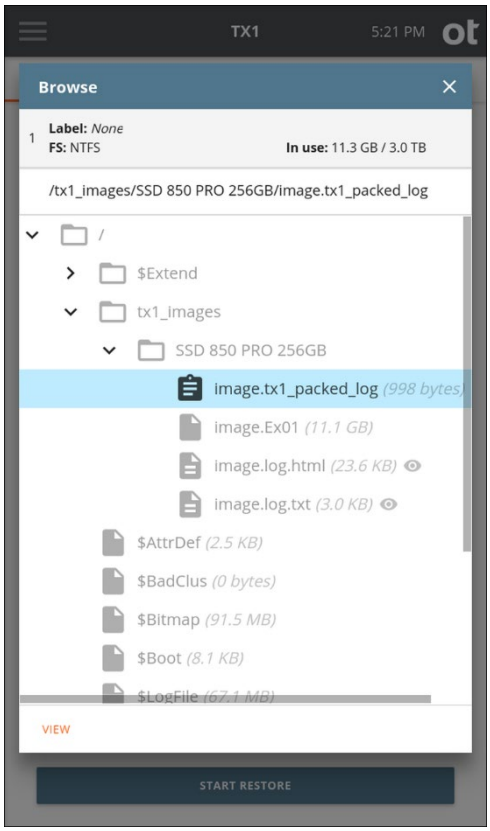

4. Select a **Destination** drive and, if desired, enable the **Trim** feature which will apply a DCO or AMA on the destination to make it appear exactly as large as the original source drive.

5. If desired, enable read-back verification. This will read the entire destination drive back after the **Restore** job is complete, calculate a read-back hash value, and compare that value with the original image file acquisition hash.

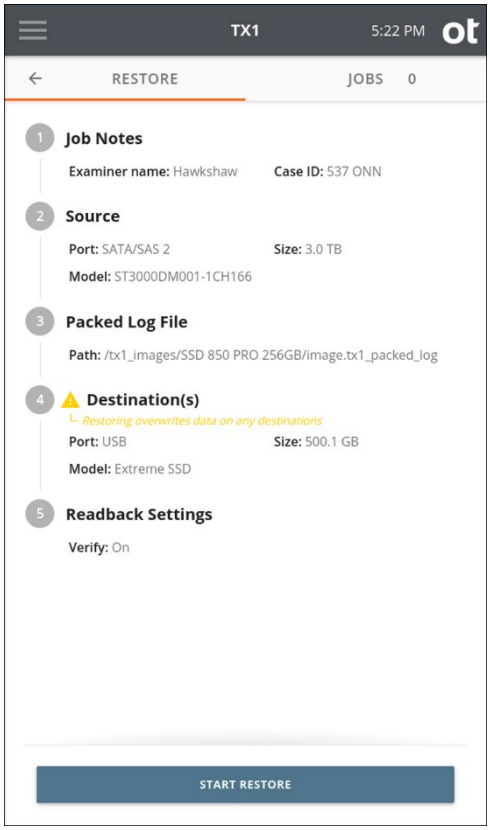

6. Tap the **Start Restore** button at the bottom of the screen. A **Job Status** modal is displayed.

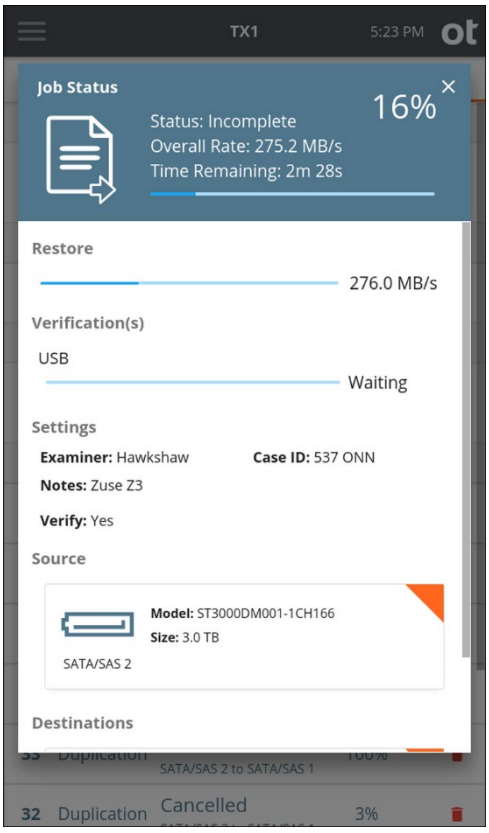

To cancel the **Restore** operation, tap the **Cancel** button from the Jobs summary screen.

When the **Restore** operation completes, the results are displayed on-screen. The log for the completed job can be viewed by tapping on the **View Log** link on the right side of the top **Job Status** screen header or through the side navigation menu.

The hash values calculated while reading back the data from the source logical evidence file are captured in the forensic logs. These appear as Restoration MD5 and Restoration SHA1 values in the **Image Source** section of the log.

4.9 Viewing sources and destinations

Tap the **Sources** or **Destinations** button on the **Home** screen to display the list of connected drives.

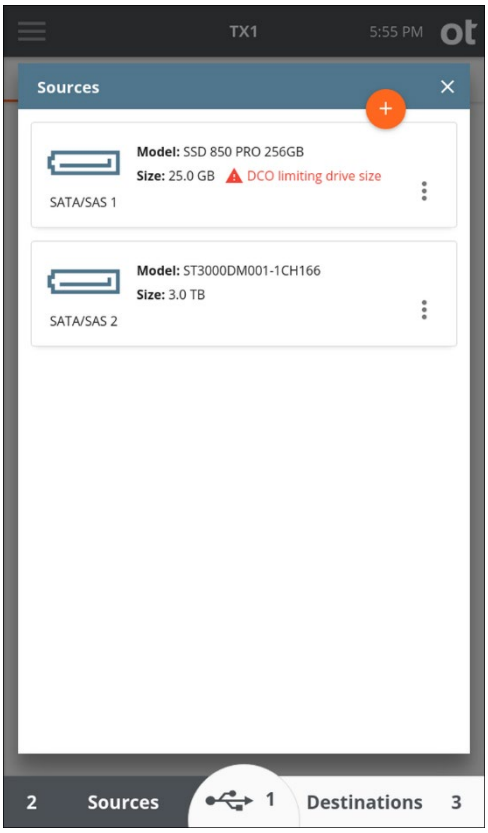

Tap a drive row to view the drive details, to access **Media Utilities**, or tap the options icon ⋮ located on the right side of the drive row to view more options.

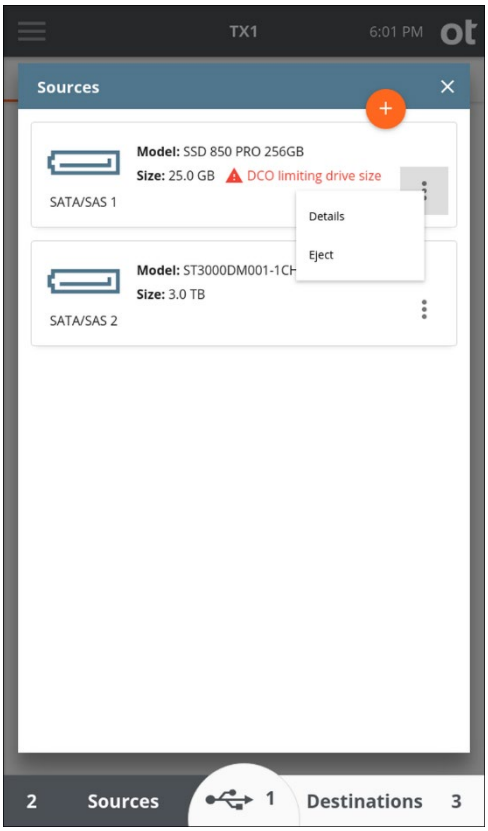

The top blue section of the drive details screen displays the physical drive interface as well as the drive model, serial number, and size.

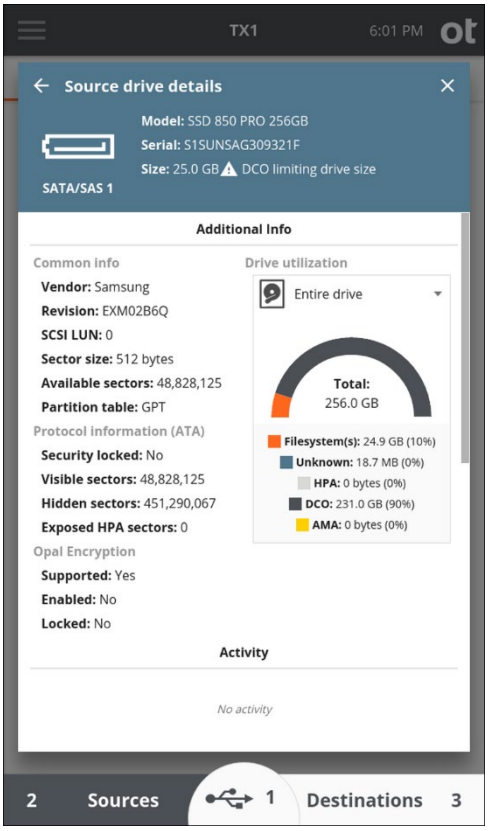

Note: Partition read errors appear in the top blue section. When such errors are detected, the drive can be physically imaged to allow further analysis in a higher-level forensic tool such as EnCase Forensic. Even though TX1 could not parse or detect partitions and filesystems, there is still forensically valuable info that can be carved or extracted using EnCase.

The **Additional Info** section displays common information for all drive types, protocol information for a subset of drive types, and drive utilization information.

In the **Drive Utilization** area, if the drive has one or more filesystems, tap the **Entire drive** menu to display a list of filesystems.

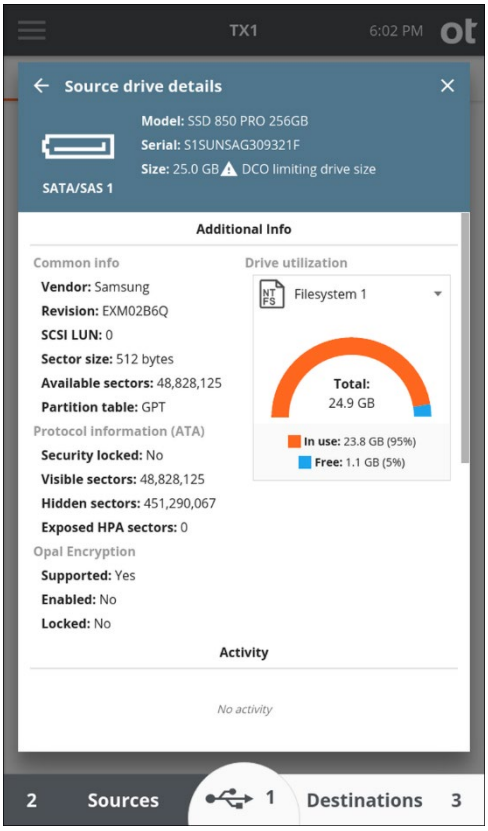

Select a filesystem to display **In use** and **Free** utilization information. Note that changing this selection from Entire drive to one of the detected filesystems only changes the utilization information displayed in this specific sub-area of this screen. All other information on this screen reflects the entire drive, and the media utilities below will act on the entire drive.

Note: All the information in the section above also pertains to the USB Accessory drives. The USB Accessory button only appears in the bottom section of the **Home** screen when a USB Accessory drive is connected and detected by TX1.

To help users identify which source drives have already been acquired, a green checkmark is shown in the bottom right portion of the drive tile for any drives that have been used in a previous, successful duplication job (clone or image), even in cases of a disconnected or reconnected drive. For active (incomplete) duplication jobs, a hollow checkmark will appear which will turn green once the job has successfully completed. While this acquisition indication method is useful in general, it was added with the advent of automated acquisition as a means of helping users keep track of which drives have been acquired in the fast-paced environment that happens when jobs are automatically

started upon source drive connection. The screenshot below shows an example of this new acquisition indication method.

Note: The green checkmark that indicates a given drive has been successfully acquired will persist in the user interface until TX1 is power-cycled, at which time no drives will indicate they were acquired (despite the fact that they may have been in a previous TX1 session). Also, if an acquired drive is removed and then reinstalled on TX1 within the same session/power-cycle, that drive will show as having been acquired.

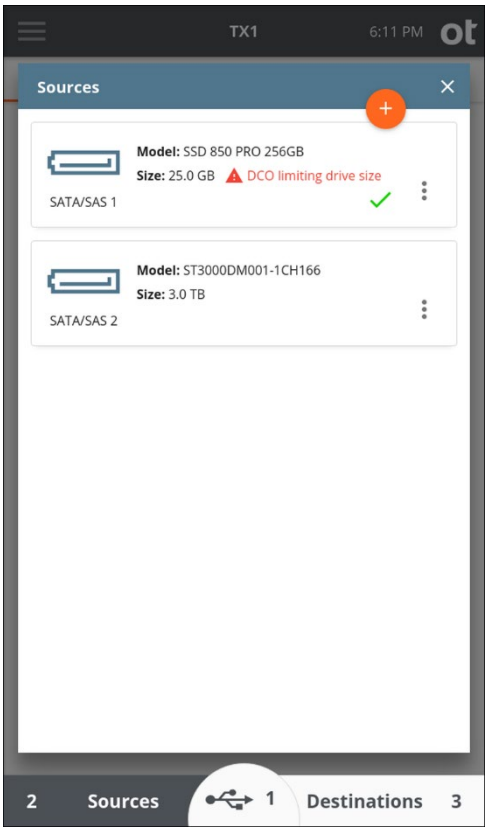

4.9.1 Encryption detection

TX1 automatically detects certain types of encryption present on attached drives. This detection is possible for software-based encryption types that have known header signatures. TX1 can also detect the hardware-based Opal encryption method. TX1 can detect the following types of encryption:

- APFS
- Apple FileVault 2
- BestCrypt
- BitLocker
- BitLocker To Go
- Check Point Full Disk Encryption

- GuardianEdge Encryption (Plus, Anywhere, Hard Disk Encryption)
- LUKS
- McAfee Drive Encryption (SafeBoot)
- Opal
- Sophos Safeguard (Enterprise and Easy/Ultimaco)
- Symantec Endpoint Encryption
- Symantec PGP Disk
- WinMagic SecureDoc Full Disk Encryption

Encryption detection information is always shown in the drive tile for a given drive, regardless of the viewing location within the user interface and what type of encryption is present (whole disk or partition based). For example, even if only one partition on a given drive is encrypted with BitLocker, a BitLocker encryption warning message will appear in the top-level drive tile regardless where it is viewed. In situations where partitions are broken out under a drive tile (which occurs for any function that involves filesystem operations, like browsing and logical image setup), more granular information is presented in the partition tiles under each drive, making it obvious which of the drive's partitions has encryption.

Note: In addition to detecting BitLocker encryption, TX1 can also unlock and perform any supported operations on a drive/partition encrypted with BitLocker. See "BitLocker encryption" on page 49 for more information.

Encryption information is also provided in the header of the **Drive Details** view, in the **Content Breakdown** view, and in the forensic logs.

The sample screenshots below show the various places where encryption detection information appears. These examples reflect a drive with multiple partitions, with only one of them having BitLocker encryption.

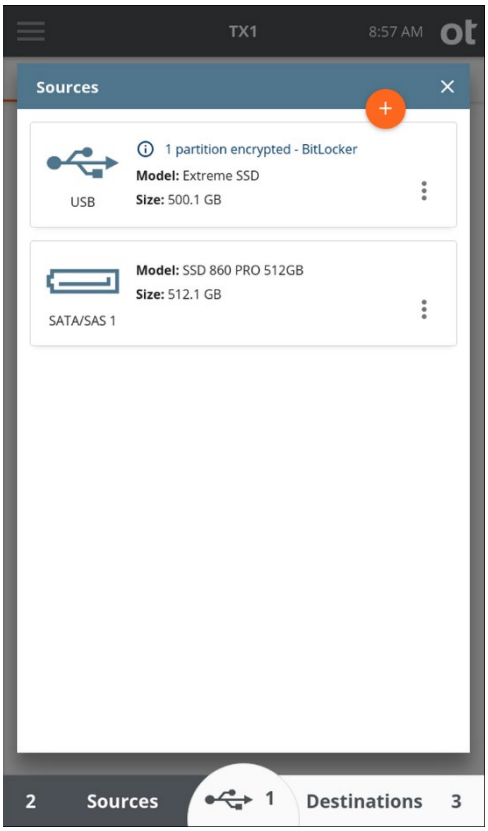

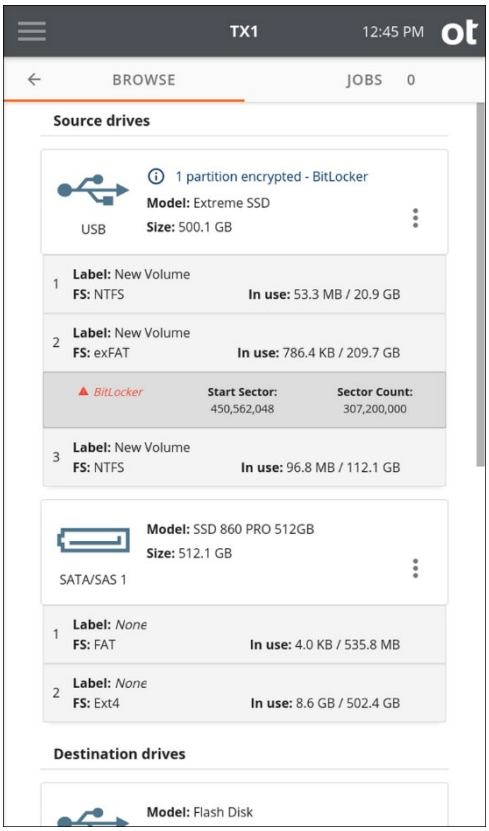

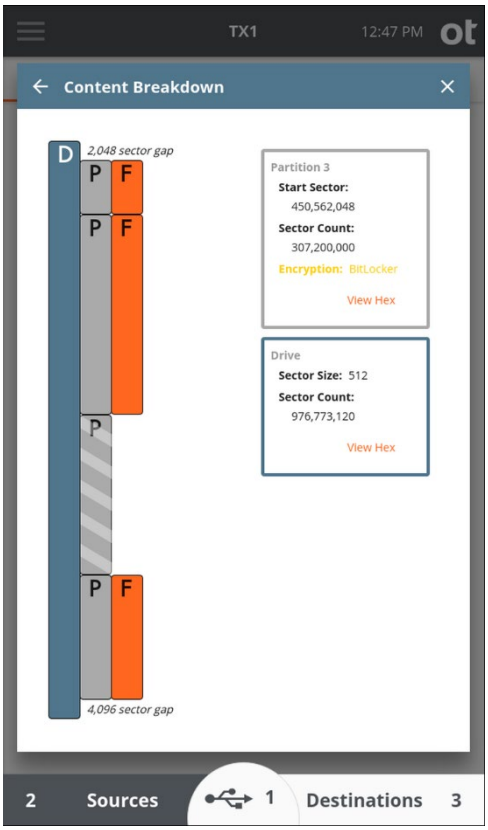

4.9.1.1 Opal encryption

Opal encryption is a unique, hardware-based encryption method that is managed by the controller on the drive with only minimal host system interaction. Opal is an industry standard created by the Trusted Computing Group (TCG) consortium that defines, among other things, the interface protocol to these types of hardware encrypted drives. These are commonly referred to as self-encrypting drives (SEDs) as the host system does little more than provide a front-end interface to enable the encryption and unlock a previously encrypted drive. The control system on the drive is responsible for encrypting/decrypting all stored data on the drive and controlling access to it.

TX1 can detect encrypted Opal SEDs and warn of the presence of Opal encryption in various places in the user interface and forensic logs. Some Opal SEDs can also be unlocked via the TX1 Opal Unlock Media Utility. Once unlocked, an Opal SED can be read from (or written to, in the case of a destination drive). A detected locked Opal drive appears like this:

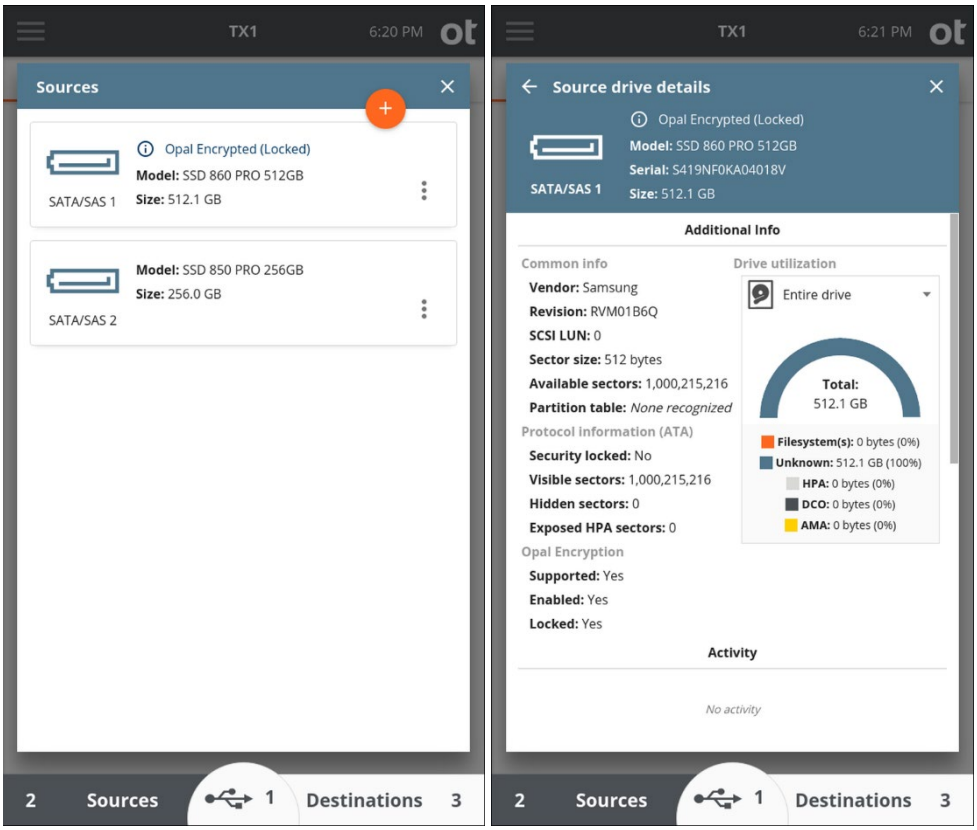

See “Encryption unlock” on page 48 for information related to unlocking Opal SEDs.

Note that Opal drives that have not had their encryption enabled will behave as regular, non-encrypted drives.

Note: Docking station type devices that have Opal drives in them must support ATA command pass-through for TX1 to properly detect the presence of Opal encryption and allow it to be unlocked. Docking stations that do not support ATA command pass-through may present locked Opal media as all zeros with no indication of Opal encryption being present in the TX1 user interface. Use caution when acquiring any docking station-based media. If you suspect a drive in a docking station is Opal-encrypted, but is not being presented that way in the TX1 user interface, removing the drive from the enclosure and connecting it directly to TX1 may yield the desired outcome.

4.9.1.2 Apple Core Storage and FileVault 2

TX1 can detect Apple Core Storage partitions and the presence of FileVault 2 encryption. FileVault 2 encryption will be indicated in both the user interface and the forensic log. The screenshots below show how FileVault 2 encryption will be indicated in the TX1 user interface.

Note: There are many possible ways to configure Apple Core Storage volumes with FileVault 2 encryption. TX1 has been designed to catch them all, but it is possible that

nuanced Core Storage configurations exist that would prevent unequivocal FileVault 2 detection. In those cases, TX1 will revert to a warning message that indicates Core Storage has been detected and that FileVault 2 is possible.

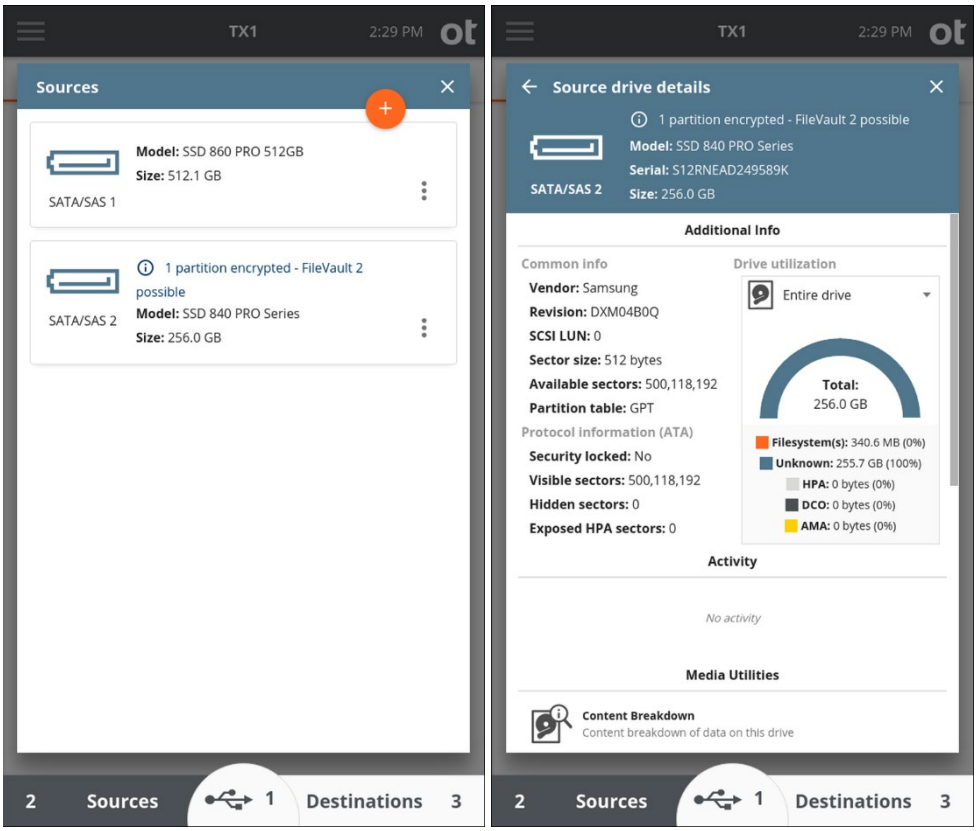

4.9.2 RAID detection

In addition to encryption detection, TX1 will automatically detect drives that were originally part of certain types of RAID systems. RAID detection is based on the existence of known signatures in specific locations on these drives. Due to the nature of the RAID specifications over time and industry adoption nuances (including many proprietary systems that do not follow all the specifications, and mixes of hardware and software-based RAID systems), it is not possible to comprehensively identify all RAID system drives. Sometimes only the primary drive in a RAID set is detectable or has detailed metadata describing the RAID setup. However, some RAID system drives can be reliably detected, and knowing about them can be of forensic value.

TX1 will detect drives from the following RAID system types:

- Intel RST (BIOS)
- SNIA DDF
- Linux MD
- Adaptec HostRAID ASR

- Highpoint (HPT37X HPT45X)
- Intel Software RAID
- JMicron JMB36x
- LSI Logic MegaRAID
- NVidia NForce
- Promise FastTrack
- Silicon Image Medley
- VIA Software RAID

RAID detection information is always shown in the drive tile for a given drive, regardless of the viewing location within the user interface and what type of RAID is detected. Any available detailed RAID information is provided in the Drive Details view, but that information depends on which type of RAID system the drive is from. Whatever information is available and displayed in the **Drive Details** screen is also captured in the forensic log for any job that involves a detected RAID drive.

Note: RAID information is detected on both source and destination drives. In the case of a clone of an SNIA DDF RAID source drive, there is a technical detail to be aware of that is of forensic importance. This RAID type stores its RAID identification information relative to the end of the drive. If the clone job did not utilize the Trim feature (to add a DCO to the destination to make its size match that of the source drive), then the new destination clone will not be detected as an SNIA DDF RAID drive. For this reason, *it is highly recommended to use the Trim feature when cloning any RAID source drive*. Also, note that a related issue can occur in the case of an SNIA DDF RAID drive that was repurposed for use as a standard destination drive. In that case, if the destination drive was not wiped before use, it will inaccurately show as an SNIA DDF RAID drive, since the original RAID identification information would still be stored at the end of the drive. Besides being good standard practice, for SNIA DDF RAID detection purposes*, it is highly recommended that destination drives be wiped before use as a forensic destination drive*.

The sample screenshots below show how two different types of RAID drives are shown in the **Sources** drive list and the **Drive Details** screens.

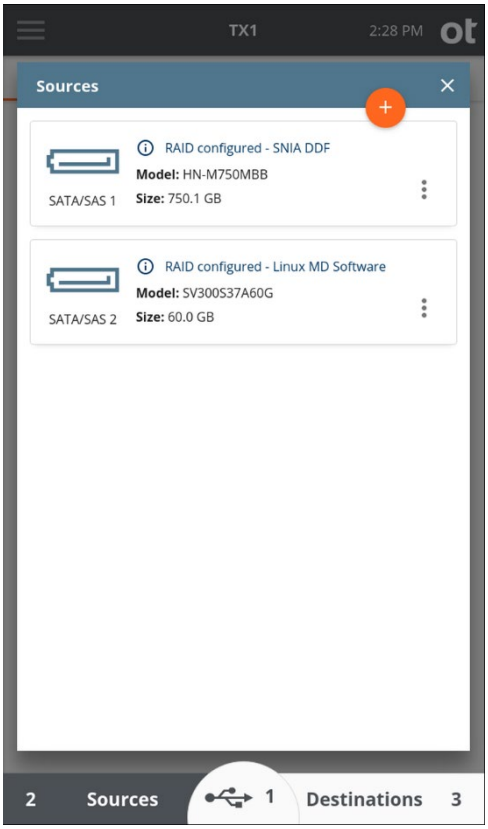

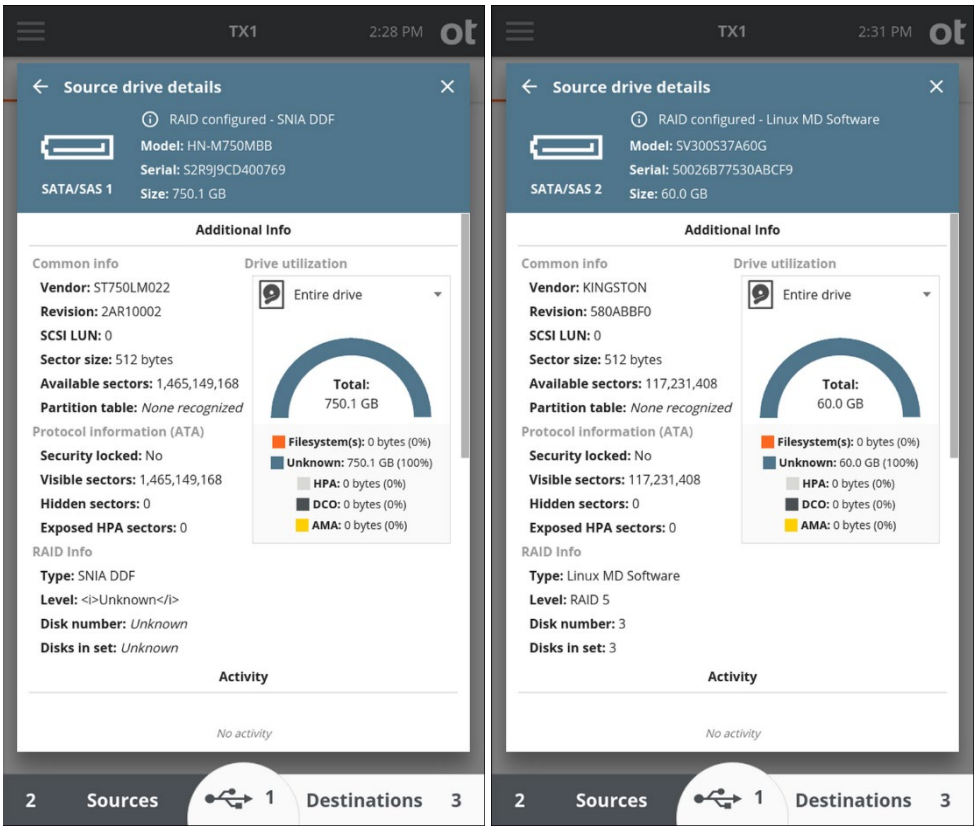

4.10 Logs module

TX1 generates a detailed log for all forensic jobs and most media utility operations. The detailed information captured in the logs will depend on the job type. A summary of the information captured for an image-based duplication job is shown below. See the sample logs at the end of this section for some specific job log examples.

- **Status** – Overall job status (Incomplete, Ok, Error/Failed, Cancelled) as well as date/time stamps, username, job notes, and TX1 unit and firmware version information.
- **Source** – Source drive details, including overall drive information (interface type, make/model number, firmware version, serial number, HPA/DCO/AMA related information, RAID and encryption information, size/layout information, and the partition table type), partition details, and, if present and supported by TX1, filesystem specific information.
- **Acquisition Results** – Details about the acquisition aspects of the job, including block start and count numbers, acquisition hash values, and read error information.
- **Configuration** – Job configuration information, such as the output file format type, segment size, and whether or not compression was enabled. If logical image searches are used, the specific search criteria is listed in this section.

- **Image Destination** – Destination drive details, including readback verification hash values (if enabled for the job), overall drive information (interface type, make/model number, firmware version, serial number, HPA/DCO/AMA related information, RAID and encryption information, size/layout information, and the partition table type), partition details, and filesystem specific information.
- **Failure Summary** – If a failure occurred during the job, this section will be shown and will include a failure reason and code. Note that the failure code is not intended to be meaningful to the end user. In cases where customer support is required to resolve a job failure situation, the failure code should be noted and included in the incident report. This information is helpful to the support team and will help in determining the root cause of the failure.

To access the **Logs** module, tap the side navigation menu icon  in the upper left corner, then tap **Logs**.

The **Log List** displays a summary of each log, including the job type, log creation date and time (essentially the job start time), and the status of the job (Started, Ok, Paused, Cancelled).

To export all logs, tap the **Export All Logs** button on the bottom of the **Log List** screen. Then select a drive with a recognized file system to which to export the logs. Note that logs can be exported to any writable media connected to TX1, including those physically connected to the front USB accessory ports or the right-side destination ports, and any network based media (iSCSI target or CIFS shares) that was mounted as a destination.

To delete all the logs stored on this TX1 for previously completed (inactive) jobs, tap the **Delete Inactive** button on the bottom of the **Log List** screen. When completed, the only logs remaining on the TX1 will be for any jobs that were active when the delete action was initiated.

Note: Resuming a paused job requires the original job log to be present. If there are any paused jobs in the **Recent** jobs list and the log for that job is deleted while it is paused, the job will not be resumable and will need to be started over.

Tap on a job row to view the detailed forensic log for that job.

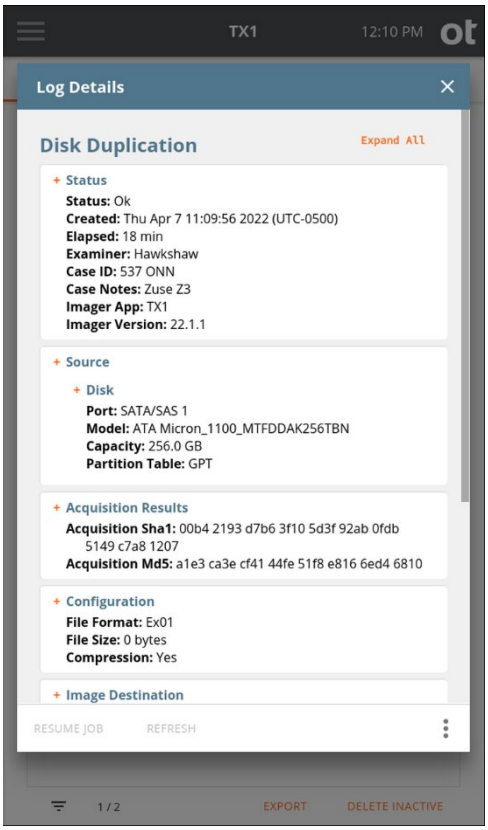

Before showing some specific sample logs later in this section, let us cover the basic log management options available from the bottom of the log details screen.

- **Resume Job** – If highlighted (orange), then the job is resumable. See “Pausing and resuming a duplication job” on page 94 for more details.
- **Refresh** – This allows the user to refresh the displayed information if the detailed log is being viewed during an active job.
- The following items are available by tapping the Options icon at the bottom of the screen:
  - **Delete** – This will delete the currently viewed job log.
  - **Export** – This allows for exportation (saving) of only the currently viewed log.
  - **Download** – Remote web interface users can download the currently viewed log. This option is unavailable when accessing TX1 locally.
  - **View Text** – Logs are stored in both HTML and text formats, with HTML being the default. Tapping the **View Text** button will switch the displayed log to the legacy text log view. When viewing a log in text format, this button changes to **View HTML** to allow switching back to that default view.
  - **View Job** – Tapping this button will display the **Job Status** screen, regardless of the job’s state.

4.10.1 HTML logs

TX1 will store forensic logs in both text and HTML file formats, in the same location as the forensic image files for a given job. While the core forensic information is the same between the two formats, HTML allows for styling and organization of the log data, such as bolding, coloring, and grouping items into collapsible sections. HTML is the default format for viewing logs on TX1, but this view can easily be switched back to the legacy text view via the ⋮ menu at the bottom of any log details screen.

As shown in the HTML log above, there are plus signs next to the section headers. This indicates there is additional information available to view in that section. Each piece of log information was categorized as critical or supplementary, and only the critical information is shown when a section is collapsed. Expanding any section will show all the available information. In that expanded view, the critical information is highlighted with bold field descriptions, while the supplementary information is shown in light gray. Note that specific pieces of log information may be considered supplementary in one situation but critical in another. For example, the encryption information for a given source drive will be considered supplementary if the drive has no encryption but will become critical if encryption is detected.

The initial state for any HTML log will be to show all fields collapsed with only the critical information displayed. While individual sections can be toggled between showing all the information or just a summary (by hitting the + or – symbol next to the section header, respectively), there is a button at the top right side of the log details screen that will allow all sections to be expanded or collapsed at one time.

Error messaging in the HTML logs has some unique functionality as well. Any error conditions will show in red text as critical information in the summarized view. Expanding the section with an error condition will show more detailed information on the error status, including the cause of the error.

4.10.2 Sample logs

Two sample logs are shown below: one from a successful duplication and one from a failed standalone verification. Note that, while the HTML log view (as shown in the section above) is the default when viewing logs via the TX1 user interface, the sample logs below are shown in legacy text format as it provides a better view for showing all the log information at once.

4.10.2.1 Log 1

```
------------------------Start of TX1 Log Entry-------------------------
Task: Disk Duplication
Status: Ok
Created: Thu Apr  7 11:09:56 2022 (UTC-0500)
Started: Thu Apr  7 11:09:56 2022 (UTC-0500)
Closed: Thu Apr  7 11:27:39 2022 (UTC-0500)
Elapsed: 18 min
Username: User1
Examiner: Hawkshaw
Case ID:
  537 ONN
Case Notes:
  Zuse Z3

Imager App: TX1
Imager Ver: 22.1.1
TX1 S/N: 000ecc58010012
-----------------------------Source Disk-------------------------------
Interface: SATA
Port: SATA/SAS 1
Model: ATA Micron_1100_MTFDDAK256TBN
Firmware revision: M0MU031
Serial number: 18271D61EABC
SCSI LUN: 0
Opal Encryption
  Supported: Yes
  Locked: No
Capacity in bytes: 256,060,514,304 (256.0 GB)
Block Size: 512 bytes
Block Count: 500,118,192
Partition table: GPT
Partition 1
  Start Sector: 2,048
  Sector Count: 1,048,576
  Partition Encryption: None detected
    Filesystem
      Type: FAT
      Block Size: 4,096 bytes
      Total Blocks: 130,812
      Free Blocks: 129,473
      Total bytes: 535,805,952 (535.8 MB)
      In Use bytes: 5,484,544 (5.4 MB)
Partition 2
  Start Sector: 1,050,624
  Sector Count: 487,346,176
  Partition Encryption: None detected
    Filesystem
      Type: Ext4
      Block Size: 4,096 bytes
      Total Blocks: 59,699,623
      Free Blocks: 56,239,509
      Total bytes: 244,529,655,808 (244.5 GB)
      In Use bytes: 14,172,626,944 (14.1 GB)
ATA-specific info
  Power-ON Block Count: 500,118,192
  HPA Block Count: 500,118,192
  DCO Block Count: 500,118,192
  AMA Block Count: 500,118,192
Tableau Encrypted: No
Whole disk encryption: None detected
Error granularity: 32,768 bytes
```

```
--------------------------------Imaging---------------------------------
Automated Job: No
Output file format: Ex01
Chunk size in bytes: 0 (0 bytes)
---------------------------Image Destination----------------------------
Interface: SATA
Port: SATA/SAS 2
Model: ATA ST3000DM001-1CH166
Firmware revision: CC43
Serial number: S1F0WNPD
SCSI LUN: 0
Capacity in bytes: 3,000,592,982,016 (3.0 TB)
Block Size: 512 bytes
Block Count: 5,860,533,168
Partition table: GPT
Partition 1
  Start Sector: 2,048
  Sector Count: 5,860,530,176
  Partition Encryption: None detected
    Filesystem (Target Filesystem)
      Type: NTFS
      Block Size: 4,096 bytes
      Total Blocks: 732,566,271
      Free Blocks: 730,426,960
      Total bytes: 3,000,591,446,016 (3.0 TB)
      In Use bytes: 8,762,617,856 (8.7 GB)
ATA-specific info
  Power-ON Block Count: 5,860,533,168
  HPA Block Count: 5,860,533,168
  DCO Block Count: 5,860,533,168
  Supports AMA: No
Tableau Encrypted: No
Whole disk encryption: None detected
Folder: /tx1_images/2022_04_07_11_09_56/
File name base: image
Verification Status: Finished OK
    Verification Sha1: 00b4 2193 d7b6 3f10 5d3f 92ab 0fdb 5149 c7a8 1207
    Verification Md5: a1e3 ca3e cf41 44fe 51f8 e816 6ed4 6810
-------------------------Duplication Results---------------------------
LBA Range Duplicated: Entire Source Disk
Total recoverable errors: 0
Total unrecoverable errors: 0
Acquisition Sha1: 00b4 2193 d7b6 3f10 5d3f 92ab 0fdb 5149 c7a8 1207
Acquisition Md5: a1e3 ca3e cf41 44fe 51f8 e816 6ed4 6810
------------——---------End of TX1 Log Entry--------------------------
```

4.10.2.2 Log 2

```
-------------------------Start of TX1 Log Entry-------------------------
*** CAUTION: THE OPERATION RECORDED IN THIS LOG DID NOT COMPLETE NORMALLY

Task: Image Readback Verification
Status: Error/Failed
Created: Thu Apr  7 11:27:59 2022 (UTC-0500)
Started: Thu Apr  7 11:27:59 2022 (UTC-0500)
Failed: Thu Apr  7 11:28:26 2022 (UTC-0500)
  Destination unreadable - 0xfb1ecc28838d027a
Closed: Thu Apr  7 11:28:26 2022 (UTC-0500)
Elapsed: 27 sec
Username: User1
Examiner: Hawkshaw
Case ID:
  537 ONN
Case Notes:
  Zuse Z3

Imager App: TX1
Imager Ver: 22.1.1
TX1 S/N: 000ecc58010012
--------------------------Image Destination---------------------------
Interface: SATA
Port: SATA/SAS 2
Model: ATA ST3000DM001-1CH166
Firmware revision: CC43
Serial number: S1F0WNPD
SCSI LUN: 0
Capacity in bytes: 3,000,592,982,016 (3.0 TB)
Block Size: 512 bytes
Block Count: 5,860,533,168
Partition table: GPT
Partition 1
  Start Sector: 2,048
  Sector Count: 5,860,530,176
  Partition Encryption: None detected
    Filesystem (Target Filesystem)
      Type: NTFS
      Block Size: 4,096 bytes
      Total Blocks: 732,566,271
      Free Blocks: 728,812,825
      Total bytes: 3,000,591,446,016 (3.0 TB)
      In Use bytes: 15,374,114,816 (15.3 GB)
ATA-specific info
  Power-ON Block Count: 5,860,533,168
  HPA Block Count: 5,860,533,168
  DCO Block Count: 5,860,533,168
  Supports AMA: No
Tableau Encrypted: No
Whole disk encryption: None detected
Folder: /tx1_images/2022_04_07_11_09_56/
File name base: image
File format: Ex01
-------------------------Failure Summary--------------------------
Reason for failure: Destination unreadable
Failure code: 0xfb1ecc28838d027a
-------------------------End of TX1 Log Entry-------------------------
```

There are three encryption related lines in a log for each drive that was part of the job, as follows:

- **Opal Encryption:** This section of the log has two sub-fields: **Supported (Yes/No)** and **Locked (Yes/No).**
- **Tableau Encrypted:** This field identifies if the drive has been encrypted by TX1. The options for this field are: **No**, **Locked**, and **Unlocked**.
- **Whole disk encryption:** This field will be populated with the specific type of third-party whole disk encryption that TX1 was able to detect. The options for this field are: **None detected**, **BitLocker**, **BitLocker To Go**, **Symantec PGP Disk**, **LUKS**, **BestCrypt**, **McAfee Drive Encryption (SafeBoot)**, **Sophos Safeguard**, **Winmagic SecureDoc**, **GuardianEdge Encryption**, **Symantec Endpoint Encryption**, and **FileVault 2**. Note that **FileVault 2** cannot be conclusively detected using standard signature inspection, but the existence of **Core Storage** can be detected. TX1 indicates that **FileVault 2** encryption is possible when a **Core Storage** partition is detected.

Note that partition information is also provided in the logs, including **Partition Encryption** status (type, if present, or **None detected**).

If TX1 detects any bad sectors on the source drive, it adds a section at the end of the job log. This additional section lists the sector address and the number of sectors of each unreadable region of the source drive. As an example, the following forensic log read error entry means that an error was encountered in at least one of the 64 sectors starting at sector offset 234,567: `Error # 1: Read error (source), address=234567, length=64`

Note: The default error granularity setting is **Standard**, which will result in a minimum chunk of 32kB of source data (64 sectors for a 512B sector drive) that will get skipped and filled with zeros upon completion of the attempted reads (assuming no reads were successful). If this condition is encountered, consider changing the error granularity setting to be **Exhaustive**, which will result in repeated read attempts of the error region with decreasing sector sizes. This will maximize the amount of recoverable data and minimize the sectors that get skipped and filled with zeros.

If error retries are enabled and TX1 is able to successfully read sector data after an initial read error is encountered, the **Total recoverable errors** count shown in the **Duplication Results** area will reflect the number of original read errors encountered. The **Total unrecoverable errors** count will reflect read errors for which no retry attempts were successful.

It is a best practice to export and delete logs from TX1 after each case. TX1 will store 100 logs before overwriting logs (starting with the oldest log). A warning will be provided before any logs are overwritten. Once a log is deleted or overwritten, the data is unrecoverable.

4.10.3 Filtering logs

TX1 can store up to 100 forensic logs. To make it easier to view, export, and delete specific logs of interest, a **Filter Logs** feature has been provided. To filter the log list, simply tap the log filter icon at the bottom left side of the log list screen. The available filter parameters are shown in the screenshot below, and they include case information (**Examiner name**, **Case ID**, and **Notes**) and drive information (**Model**, **Vendor**, **Serial number**, and CIFS/iSCSI share information). Each of the filter fields can be populated manually by typing the desired value in the field(s) of interest. Alternatively, the fields can be auto-populated from case information default values or from information from an attached drive or network share. To auto-populate these fields, simply tap the orange **Use Case Info Defaults** and/or Use Connected Drive buttons. Note that you can also auto-populate these fields using the buttons and then manually override specific fields to match your desired filter parameters. The screenshot example below used the **Use Connected Drive** option to fill in the filtering information for the desired drive.

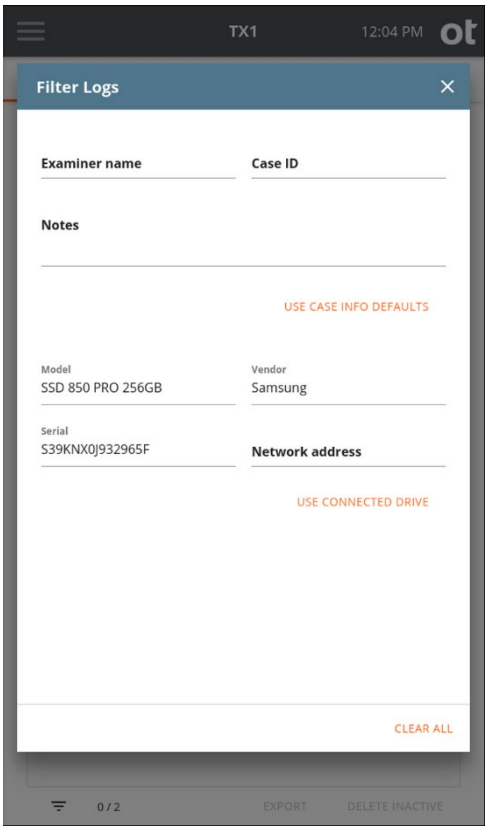

Note: When using the **Network address** field to filter the log list, CIFS share paths may be used (full or partial) as may an iSCSI IQN or target IP address. Also, when shares are mounted using nicknames, the underlying IP address information is not stored in the logs and thus cannot be used for filtering.

Once the desired filter parameters are set, simply close the **Filter Logs** window to see the list of logs that match your chosen parameters. The number of logs that matched your filter parameters is shown next to the filter icon in the bottom of the log list window. This subset of logs can be viewed on the unit, exported to external media, or deleted. To export or delete the filtered log list, simply tap the appropriate button on the bottom right of the log list screen and follow the prompts. As with the full log list, when deleting a filtered list of logs, only logs for previously completed (inactive) jobs will be deleted.

To remove all filtering and see the entire log list again, either navigate away from the filtered log list and then reselect the main log list, or tap the filter icon again and then tap **Clear All** at the bottom right of the Filter Logs screen and close that screen.

4.11 Remote web interface

TX1 offers the ability to use any network connected TX1 unit remotely via a web browser. If your TX1 has a user setup with remote access privileges and a valid IP address, remote access is as simple as opening a browser window and entering the IP address or hostname of TX1 into the address field. A login screen will always appear first on the remote browser, and, after credentials are entered and validated, the TX1 user interface will appear and be fully usable as if you were working locally on the unit itself. Note that local and remote users can work independently on the same TX1 unit. However, TX1's resources are shared between all logged in users, so jobs and other changes made from one connection type (local or remote) will be seen by all users and will affect what other users can do.

The remote web interface can be used to perform nearly all functions available on the local user interface, with the following differences:

- For security reasons, the remote login screen requires typing in both the username and password. On the local unit, a pulldown list of available usernames is provided. Usernames and passwords are case sensitive.
- When using the **Browse** function remotely, any individual file from any mounted drive with a supported filesystem can be downloaded to the remote device. After selecting the desired file in the remote **Browse** window, simply tap the **Download** button at the bottom right, and the file will be downloaded to your remote system. This is a convenient way to view any file from a TX1-mounted drive using the remote computer's applications, which can help with evidence triage decisions.
- When accessing logs from the remote web interface, you can download targeted log files. Select a log from the Log Details view, click the options icon ⋮ and select **Download**. The **Download** button is not displayed when accessing TX1 locally.
- The time shown in the top-right corner of the remote view represents the time zone of the remote computer/browser, not the time zone of TX1. However, the logs for any remotely initiated jobs will reflect the time zone as set on TX1.
- The **Factory Reset** feature is not available to remote users.
- **Network Configuration** settings are viewable remotely but can only be changed directly on TX1.

- Locking the system (PIN lock) from any location (local or any remote user instance) will lock all the active screens.

4.11.1 SSL certificate setup and installation

TX1 uses SSL certificates to ensure secure communication during remote sessions. By default, TX1 generates a new self-signed certificate on power up if one does not already exist on the system. These self-signed certificates expire after one year, and TX1 will automatically generate a new one before expiration. Users can manually initiate a new TX1 self-signed certificate or install their own SSL certificates.

To either generate a new self-signed certficate or install a user certificate with TX1:

1. Open the side navigation menu and click **Network settings**. **Current status**, **Configuration**, **Custom hostname**, **802.1X settings**, **CA certificate**, **Client certificate**, and **HTTPS certificate** details are displayed.

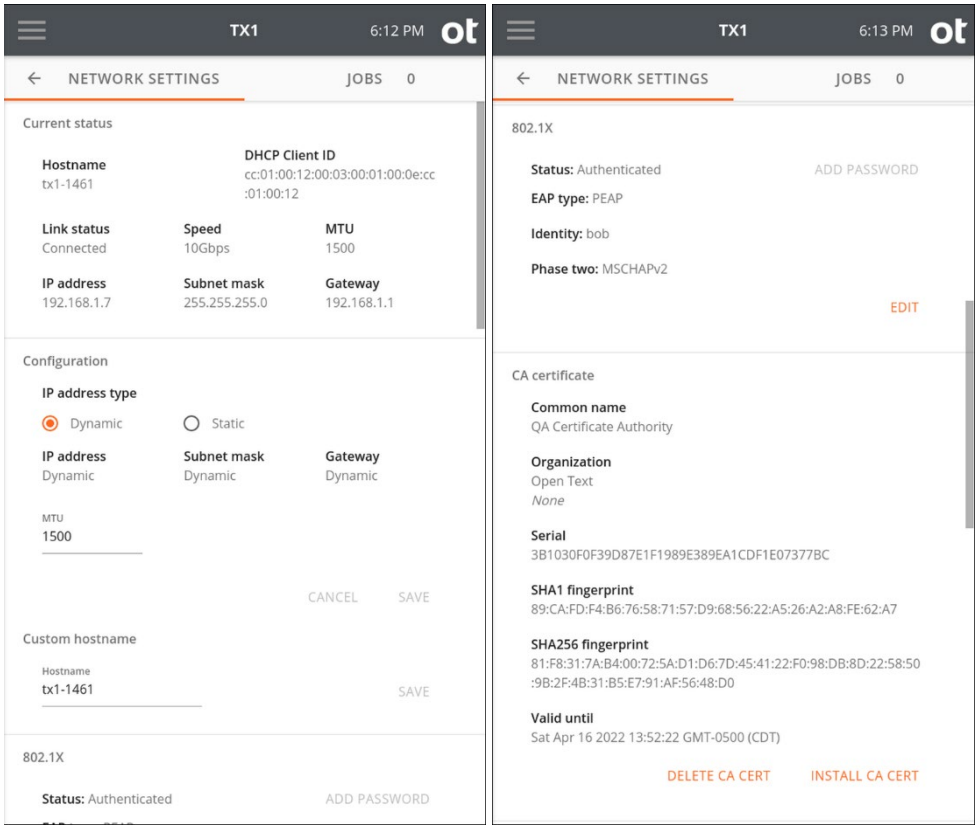

2. Scroll to the bottom of the screen.

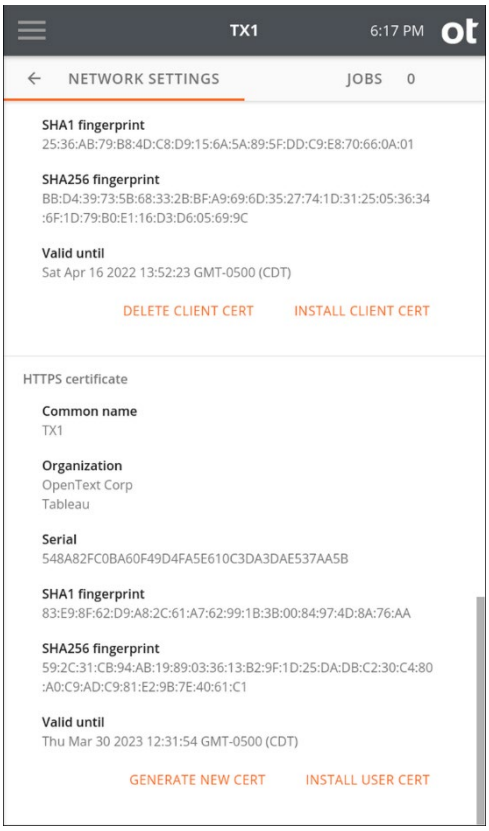

3. You have two options:

   - To create a new self-signed certificate:

     1. Tap **Generate New Cert**. A warning modal appears asking you to confirm generation of a new self-signed SSL certificate and system reboot.

     2. After confirming, the system reboots.

   - To install your own SSL certificate on TX1:

1. Tap **Install User Cert**.

     2. Select the desired drive/filesystem from the list, and a file browser window is displayed.

     3. Navigate to and select the `.pem` SSL certificate file. A warning modal appears asking you to confirm installation of the SSL certificate and system reboot.

     4. Select **Install**. The system reboots.

Once you install your own certificate, TX1 will retain it in the event of reboot or power disruption. Manually generating a TX1 self-signed certificate will overwrite your own certificate and return TX1 to the default state of generating a new self-signed certificate upon annual expiration.

Note: TX1 remote user interface has been validated for use with the following web browsers: Chrome (v86.0.4240.75), Firefox (v81.0.2), and Safari (v13.2.1). Older and newer versions of these browsers and even other browsers may work fine but have not been validated at this time.

4.11.2 Accessing TX1 remotely

To access TX1 remotely, a valid username and password are required. The user account must also have remote access enabled.

To access a TX1 remotely:

1. On the **User Management** screen, confirm the user is set up on the system with a password and remote access is enabled. In the example below, default user User1 is shown. Before enabling remote access, a password must be added by tapping the **Add Password** button. Once a password has been entered, toggle the **Allow Remote** switch, as shown below.

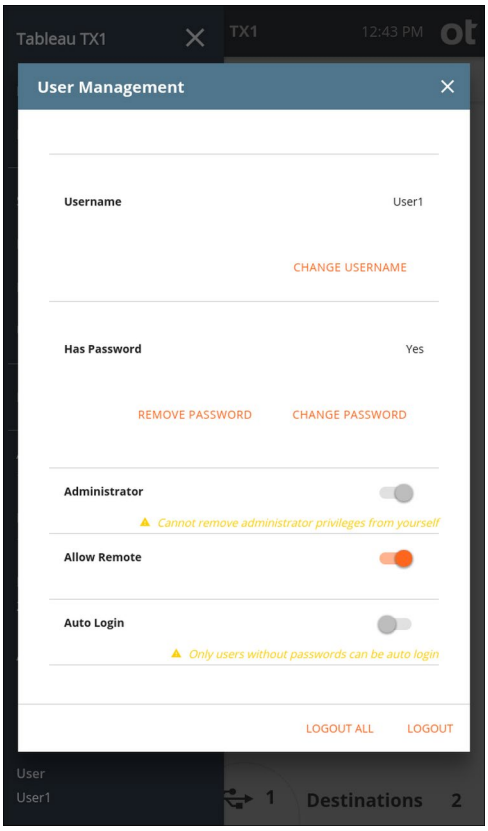

2. Connect TX1 to your local area network with an Ethernet cable.
3. Obtain the assigned IP address from the side navigation menu, as shown below.

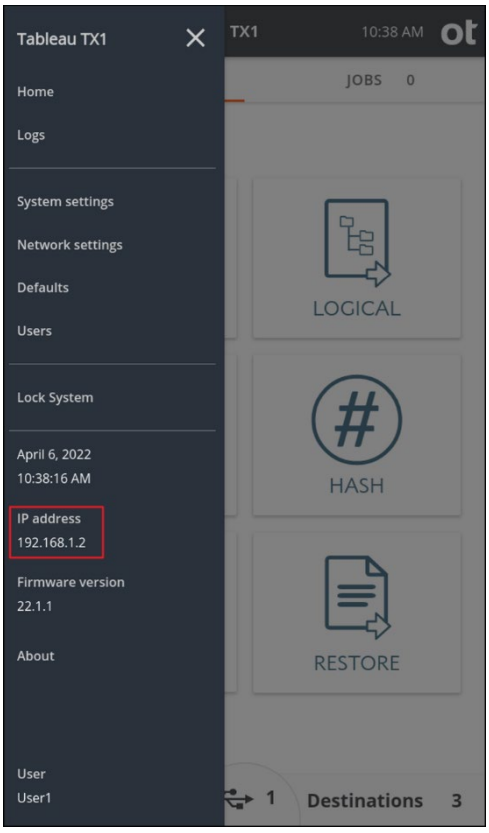

4. Open a web browser of your choice on a computer/device that is connected to the same local area network as TX1, type TX1's IP address (or hostname) into the address field in the web browser (same field where you would enter a web page URL), and press Enter. The remote UI login screen should be shown in your browser.

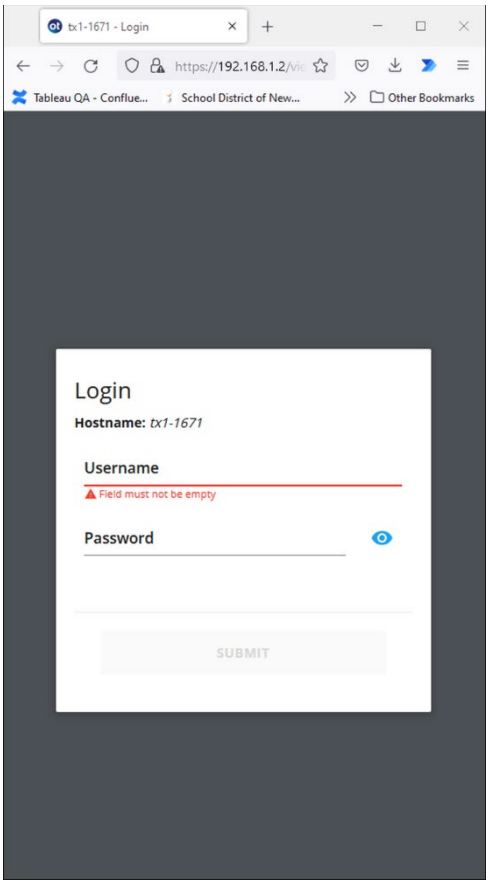

Note: The SSL certificate will show as invalid for TX1 at this time. An exception will need to be made to view the remote user interface.

5. Enter the username and password, and the TX1 user interface should appear in your browser, as shown below.

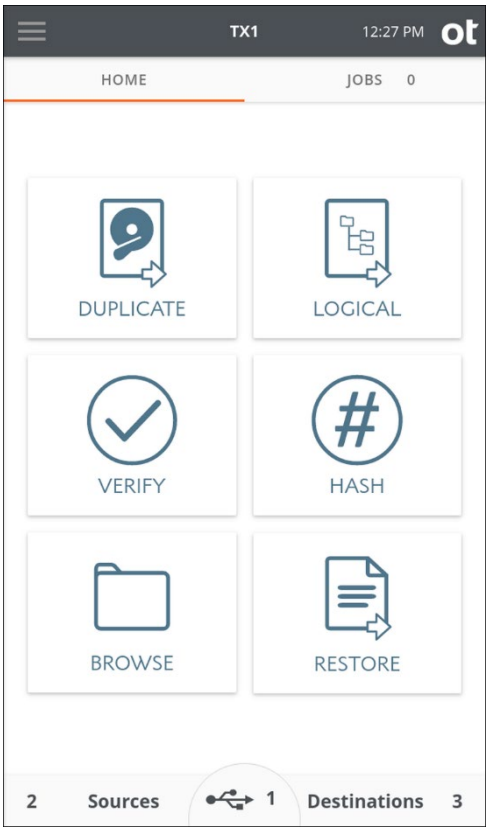

Note: If any TX1 units in your network environment have a custom hostname defined (in the **Network Settings** screen), that hostname will be displayed in the browser tab when the unit is accessed remotely. This allows for easy identification of the specific TX1 associated with each browser window, which is helpful in a lab environment with multiple TX1 units being accessed remotely from the same browser.

5 Adapters

This chapter describes the add-on drive adapters available for TX1, which extend imaging capabilities in an easy to connect and use manner.

Important:

The PCIe interface on TX1 does not support hot-swapping. PCIe adapters with drives installed must be attached to TX1's PCIe port before powering up the unit. Also, removing a PCIe adapter from TX1 or removing a drive from a PCIe adapter before powering down the device will cause system instability and could possibly damage the

drive, adapter, or TX1. TX1 must be powered down before removing the drive from any Tableau PCIe adapter or removing a PCIe drive/adapter from the TX1 PCIe port.

5.1 PCIe SSD adapters

Tableau PCIe SSD adapters enable the acquisition of PCIe based SSDs of various types via TX1's PCIE source port. The following adapters are available individually or as part of an adapter kit:

- PCIe card SSD adapter – TDA7-1
- PCIe m.2 SSD adapter – TDA7-2
- PCIe adapter for Apple SSDs 2013+ (through 2016) – TDA7-3
- PCIe u.2 SSD adapter cable – TDA7-4

Visit the Tableau product website to learn more about available Tableau PCIe adapters.

5.2 PCIe IDE adapter (TDA7–5)

The PCIe IDE Adapter (TDA7-5) enables the acquisition of IDE drives via TX1's PCIe source port. The IDE adapter kit (sold as an add-on to the TX1 kit) includes the IDE power and signal cables used to connect the IDE drive to the adapter. The PCIe cable used to connect the adapter to TX1 is included with the TX1 kit.

5.2.1 Using the TDA7-5

1. With TX1 powered off, attach the TDA7-5 to the TX1 by connecting it to the source PCIe port on the left side of TX1, using a Tableau TC-PCIE cable (4" or 8").
2. Connect an IDE drive using the TC6-2 ribbon cable (blue connector goes to the TDA7-5) and TC2-8-R2 power cable.

3. Power on TX1.
4. The **Sources** drive counter will increment by one to let you know your IDE drive is connected. Tap the Sources tab to view IDE drive details.

5.3 PCIe FireWire adapter (TDA7–9)

The PCIe FireWire adapter (TDA7-9) enables the acquisition of FireWire drives via the TX1's PCIe source port. The PCIe FireWire adapter kit (sold as an add-on to the TX1 kit) includes 9-pin and 6-pin FireWire adapter cables which are used to connect the FireWire media to the adapter. The PCIe cable used to connect the adapter to TX1 is included with the TX1 kit.

Note that this adapter was originally created to replace the T9 Forensic FireWire bridge. TX1 has a native FireWire port which makes this particular PCIe adapter only useful in cases where many FireWire drives need to be acquired.

Note: There is a 9-pin 1394b port and a 6-pin 1394a port on TDA7-9. If media are attached to both FireWire ports of the adapter, both will show as available drives in the TX1 user interface, which will allow for both to be used simultaneously for any available drive actions (imaging, hashing, etc.).

5.4 Apple Target Disk Mode acquisition adapters

TX1 is designed to acquire Apple computers that support Target Disk Mode. This can be done via three different Apple computer connection interfaces: USB-C, FireWire, and Thunderbolt 2. Depending on the Apple computer interface connection, different adapters and/or cables are required to allow for connection to the source side of TX1.

Note: Beginning in December, 2017, some Apple devices started using a new secure enclave interface to their integrated SSDs, which has created challenges for forensic examiners. The core of the secure enclave interface is the T2 chip, which sits between the internal memory devices and anything that needs access to that memory. When provided with the proper credentials using proprietary Apple interface commands, the T2 chip will unlock the encrypted data and make it available to the requestor. However, without the proper credentials or without a mechanism to provide the proper credentials to the Apple device, the exposed memory will be encrypted. As of this writing, Tableau devices do not support the special mechanisms to expose Apple memory devices protected by a T2 chip, nor a way of providing credentials to enable acquisition of the unencrypted data. The Tableau development team is actively pursuing support for this scenario, and a future update that provides such support will be announced as soon as it is available.

5.4.1 USB-C to USB-A adapter cable

Newer Apple computers have a USB-C connector that can communicate via either USB or Thunderbolt protocols. The Apple system will detect the type of connection and automatically adjust how it communicates with the device. A USB-C to USB-A style interface cable capable of USB 3.0 or USB 3.1 Gen 1 speeds (5 Gbps) will allow connection between the Mac computer and the TX1 USB source port. These adapter cables are widely available from a variety of commercial retailers, including Apple. A USB-C to USB-A adapter cable is shown below.

5.4.2 FireWire adapter cable

TX1 has a source side FireWire 800 port, making acquisition of a Macintosh computer via its FireWire connection as easy as plugging in an interface cable. Different Macintosh computers have different FireWire connectors, so check the connector type before you begin. FireWire 800 connectors are common on many legacy Mac systems. A FireWire 800 9-pin to 9-pin cable is shown below.

5.4.3 Thunderbolt 2 adapter cable

Adapting from a Thunderbolt 2 connector on a Mac to TX1 requires two separate adapters/cables. A Thunderbolt 2 to FireWire 800 (9-pin) adapter is used along with the same FireWire 800 (9-pin to 9-pin) cable shown above to connect between the Thunderbolt 2 port on the Macintosh and the TX1 FireWire 800 port. A Thunderbolt 2 to FireWire 800 adapter is shown below.

6 Specifications and troubleshooting

6.1 Specifications

Connectors: Source Side	
SATA/SAS	Two SATA/SAS (6 GBPS) signal connectors
USB	One USB 3.1 Gen 1 (5 GBPS) Standard-A connector
FireWire	One 1394b “FireWire 800” signal connector

PCIe	One PCIe (10 GBPS) adapter connector
Drive Power	Two 3M-style 4-pin power connectors for SATA/SAS drive power
Connectors: Destination Side	
SATA/SAS	Two SATA/SAS (6 GBPS) signal connectors
USB	One USB 3.1 Gen 1 (5 GBPS) Standard-A connector
TX1-S1	One TX1-S1 (Two SATA/SAS 6 GBPS) signal connector
Drive Power	Two 3M-style 4-pin power connectors for SATA/SAS drive power
Connectors: Miscellaneous	
Ethernet	One 10 Gbps Ethernet Connection (Source or Destination); auto-negotiable to 1Gbps
USB	Two USB 3.1 Gen 1 (5 Gbps) Standard-A Connectors
SD Card	One SD Card Connector for Device Firmware
DC Input	One Barrel Connector for use with Tableau TP6 Power Supply
User Interface	
LCD	7.0 in. graphic LCD (600 x 1024 Resolution) with capacitive touch-screen
Power Button	One On/Off power button
Indicators	
Power Indicator	White LED indicates TX1 is powered on
Status Indicator	Multi-color LED indicates TX1 job status
Speaker	Audio tones indicate job completion and errors
Physical / Environment	
Power	55 Watts typical operating power (not including external drive power) 130 Watts maximum allowed power draw (including external drives)
DC Input	24 VDC (Nominal)
DC Output (per drive)	+5/12V @2A (Spin-up) +5/12V @1A (Continuous)

Dimensions	9.5 in. (L) x 6.5 in. (W) x 2.625 in. (H)
Weight	35 oz (980 g)
Storage Temperature Range	-20 to 70° Celsius
Operating Temperature Range	0 to 40° Celsius ambient (room temperature)
Relative Humidity	Up to 90% (Non-condensing)
Warranty	
TX1 Unit	Three Years Parts and Workmanship from Date of Purchase
TX1-S1 and TX1 Accessories	One Year Parts and Workmanship from Date of Purchase

6.2 Troubleshooting common problems

This section covers the following troubleshooting issues and solutions:

- Power supply issues
- Thermal issues
- Problems with drive detection
- Problems detecting Apple devices in the Target Disk Mode
- Long time to complete locally initiated firmware update
- Replacing the backup battery for the real-time clock

6.2.1 Power supply issues

The power supply provided with TX1 is capable of powering TX1 and nearly all combinations of drives. TX1 also employs staggered power sequencing for the source and destination drives. With staggered sequencing, power is first provided to one drive as it spins up, then to the second drive as it spins up, and so on. This feature prevents large current demand spikes at initial power-up, which helps to ensure reliable operation even with a heavily loaded system. Due to staggered power sequencing, it is normal to hear the source and destination drives spin up separately.

During power-on initialization and self-test, TX1 checks the output voltages of the power supply. If the voltage is below the minimum specification, TX1 displays a warning.

If you are having difficulty turning TX1 on, check the status of the DC power LED on the TP6 power supply connector to ensure that it is on, which indicates TX1 is connected to power.

6.2.2 Thermal issues

TX1 is constantly monitoring the operating temperature of key components inside the unit. While it was designed to have plenty of operating temperature margin, conditions that affect the airflow or the effectiveness of the cooling system inside the unit can occur. Depending on the severity of the issue, overheating can possibly cause performance issues and/or physical damage to the unit.

Should any of the key components become overheated for any reason, a warning will be provided in the top navigation bar (dark gray bar across the top of the screen). The first level warning will be a yellow triangle that indicates temperatures are on the rise. Clicking the yellow triangle will provide an informative message, but the unit will otherwise operate normally in this condition indefinitely. If conditions do not improve, the warning triangle may eventually turn red, indicating a major thermal issue with the unit. In that scenario, you will automatically be prompted to immediately shut down the unit to prevent permanent damage. You will have the option to ignore the shut down message, but this is strongly discouraged.

If the yellow warning triangle is present, please check to make sure that all airflow openings on the unit are open and free of obstruction. This includes the inlet vents on both sides of the unit and the fan outlet vent in the rear. If there are no obstructions to these airflow vents, then please contact OpenText Customer Support at your earliest convenience for further guidance.

If the red warning triangle is present, please immediately shut down the unit via the automatic shut down prompt. Check that there are no obstructions to the inlet or outlet air vents, and let the unit cool down for a few minutes before attempting to use it again. If the red warning triangle condition returns, please immediately shut down the unit and contact OpenText Customer Support for further guidance.

6.2.3 Problems with drive detection

When using a product like TX1, the most common problem you may encounter is a failure to achieve drive detection. Most drive detection problems are the result of improper cabling. The following table lists the most common drive detection problems and corrective actions.

Problem	Corrective Action
General drive detection	Check the power and signal cable connections between TX1 and the drive to ensure that all connectors are properly seated. If TX1 still does not detect the drive, cycle the power to attempt to detect the drive during a fresh start-up sequence.
PCIe SSD is not detected	Due to the nature of the PCIe bus not supporting hot-plugging, TX1 can only detect drives connected to the source PCIe port during start-up. Therefore, connect the PCIe SSD to the appropriate Tableau PCIe adapter and

Problem	Corrective Action
	cable, and connect the PCIe adapter and cable to TX1, while TX1 power is off.
IDE drive is not detected	Ensure that the blue end of the IDE signal cable faces the IDE adapter, and that the IDE drive is configured for Master or Single Drive mode. Also, since the TX1 IDE adapter connects through the PCIe port, it must be connected while TX1 is powered off as with other TX1 PCIe adapters.
SATA or SAS drive is not detected	Use only the Unified SATA/SAS cables provided by Tableau (TC4-8-R3 or TC4-8-R2). With some SATA drives, the SATA connector may be loose. Ensure the cable is seated properly in the SATA connector of the drive.
USB drive is not detected	While there are a variety of reasons a USB drive may not detect on TX1, one specific type of USB drive was not detectable on TX1 prior to the 20.2 firmware update. Some USB drives (especially proprietary, self-encrypting thumb drives like Kingston's IronKey) expose a small CDFS volume to the host system instead of the main data volume. This CDFS volume typically includes product literature and, more importantly, an application that allows for entry of a key/password/passphrase on the host system. While TX1 cannot run these unlocking applications (as they are designed for x86 based systems), as of the 20.2 firmware version, it will at least detect such a drive and report its type to the user in the drive tile.

Tableau has tested TX1 with an extensive in-house library of different drives spanning many years of drive development, but there may still be compatibility issues with some drives. Tableau issues firmware updates to address most compatibility issues. If your drive is not recognized by TX1, check the Tableau download webpages (https://security.opentext.com/tableau/download-center) to see if any firmware updates are available for TX1.

If there are no firmware updates available to resolve your detection issue, please contact your Tableau reseller or OpenText Customer Support to report your issue or ask for further assistance.

6.2.4 Problems detecting Apple devices in target disk mode

Accessing drives inside Apple devices is most simply accomplished by putting the Apple device into target disk mode (TDM) and then connecting it to the USB source port on your TX1 using a quality USB-C to USB-A adapter cable. Problems with these types of

connections are not uncommon. The table below lists the most common issues seen when trying to mount and acquire an Apple computer in TDM.

Problem	Corrective Action
No detection on TX1; Apple computer boots to normal user login screen (not in target disk mode)	Pressing the “T” key on the keyboard during bootup is how Apple computers are put into target disk mode. This mode is indicated by connection logos moving around the screen that show the different ways an external device can be connected to the target Apple computer (dependent on the available connections on that computer). If the Apple device does not enter target disk mode after pressing the “T” key during bootup, it may be defective and alternative methods for memory acquisition may be required (such as drive removal, if supported, or third party services that specialize in hardware based forensic data retrieval). With the advent of the T2 secure enclave interface, it’s also possible that the System Preferences on the Apple device are set to block remote booting, which disables entry into target disk mode. If possible, change the System Preferences to allow remote booting, and then retry the TDM boot sequence.
No detection on TX1; Apple device is in TDM, but more than the USB logo is shown on the Apple screen	This condition indicates that the Apple device is not properly connected to a powered-on TX1. This is most commonly caused by the use of a defective cable between the Apple computer and TX1. See “Apple Target Disk Mode acquisition adapters” on page 171 for recommended cables. If you are already using the recommended cable and it has been used successfully in the past, try gently bending the cable at each end or in the middle to see if the logo pattern on the Apple screen changes.
No detection on TX1; Apple device is in TDM with only the USB logo on the screen□When connecting an Apple computer to TX1, the order of operations matters.	The recommended steps are as follows: • Power-on the Apple computer while pressing the “T” key until only the USB logo appears on the screen. • Power on TX1 and ensure it is fully booted to the Home screen. • Connect the interface cable to the Apple computer first, and then to the USB-A connector on the source side of TX1.
No detection on TX1; Apple device is in TDM and all cabling is solid	As mentioned in “Apple Target Disk Mode acquisition adapters” on page 183 above, it’s possible that the Apple computer you are attempting to acquire is a newer version with a T2 security chip. TX1 does not currently provide a means of detecting or unlocking drives from a T2 Apple computer.

Problem	Corrective Action
TX1 detects two drives even though only one Apple computer is connected	This indicates the presence of two drives internal to the Apple computer. Certain iMac models have the option of a small SSD paired with a large HDD. Those two drives can be configured on the Mac to act independently or as a virtual Fusion drive setup (one virtual drive that makes use of the two physical drives). In either configuration, TX1 will show the two drives as separate drives, each of which can be acquired independently.

If you are still having trouble detecting your Apple computer in TDM, please contact OpenText Customer Support to report your issue or ask for further assistance.

6.2.5 Long time to complete locally initiated firmware update

The ability to update TX1's firmware from the unit via any mounted media (local or network based) has made firmware updates more flexible and convenient. However, this process can take a long time when used with original/obsolete TX1 SD cards. In Tableau testing, updating an original TX1 SD card via the user interface took approximately 13 minutes compared to approximately 2 minutes 30 seconds with a new model SD card. During the update process, the screen will show a progress bar to let you know the process is still active, but the bar moves so slowly with the original (now obsolete) SD cards that it may appear to be hung. Note that there is nothing wrong with the original SD cards after a local firmware update. Such an updated card can be used with confidence.

Please be aware of this delay with the original SD cards. Should this lengthy firmware update time become an issue for you or your organization, please contact OpenText Customer Support for resolution options.

Note that the original SD cards are easy to identify as they have the Guidance Software "G" logo in the bottom right corner of the label. New SD cards have the OpenText logo.

6.2.6 Real-time clock data retention issue

Under normal operating conditions, the real-time clock on your TX1 should retain the time and date settings for the life of the product. If the time and/or date setting is not being retained after power cycles, there could be an issue with the battery inside the unit. We do not recommend opening your TX1 for any reason, including battery replacement. If you notice an issue with the time and/or date setting not being retained, please contact OpenTextCustomer Support to report your issue and ask for further assistance.

About OpenText

OpenText enables the digital world, creating a better way for organizations to work with information, on-premises or in the cloud. For more information about OpenText (NASDAQ/TSX: OTEX), visit opentext.com.

Connect with us:

OpenText CEO Mark Barrenechea's blog

Twitter | LinkedIn