

## OpenText™ Forensic TX2 Imager

### **User Guide**

This guide presents a wide range of technical information and procedures for using the OpenText™ Forensic TX2 Imager.

ISTXII260100-UGD-EN-1

---

## **OpenText™ Forensic TX2 Imager**

### **User Guide**

ISTXII260100-UGD-EN-1

Rev.: 2026-Jan-30

**This documentation has been created for OpenText™ Forensic TX2 Imager 26.1.**

It is also valid for subsequent software releases unless OpenText has made newer documentation available with the product, on an OpenText website, or by any other means.

#### **Open Text Corporation**

275 Frank Tompa Drive, Waterloo, Ontario, Canada, N2L 0A1

Tel: +1-519-888-7111

Toll Free Canada/USA: 1-800-499-6544 International: +800-4996-5440

Fax: +1-519-888-0677

Support: <https://support.opentext.com>

For more information, visit <https://www.opentext.com>

#### **© 2026 Open Text**

Patents may cover this product, see <https://www.opentext.com/patents>.

#### **Disclaimer**

##### **No Warranties and Limitation of Liability**

Every effort has been made to ensure the accuracy of the features and techniques presented in this publication. However, Open Text Corporation and its affiliates accept no responsibility and offer no warranty whether expressed or implied, for the accuracy of this publication.

---

# Table of Contents

<b>1</b>	<b>Preface .....</b>	<b>7</b>
1.1	Drive capacity and transfer rate measurement conventions .....	7
<b>2</b>	<b>Overview .....</b>	<b>9</b>
2.1	Kit contents .....	13
2.2	Navigating the touchscreen display .....	13
2.2.1	HOME screen .....	14
2.2.2	Side navigation menu .....	16
2.2.3	JOBS tab .....	16
2.2.4	Job status .....	18
2.2.5	Quick Reference Guide .....	19
2.3	Reading the status LEDs .....	19
2.4	Interpreting audio feedback .....	20
2.5	On-screen warnings .....	20
2.6	USB keyboard and mouse support .....	20
<b>3</b>	<b>Configuring the OpenText Forensic TX2 Imager .....</b>	<b>21</b>
3.1	Startup sequence .....	21
3.2	Configuration settings .....	22
3.2.1	System settings .....	22
3.2.2	Network settings .....	23
3.2.2.1	Configuring 802.1X network authentication .....	26
3.2.2.2	HTTPS certificate setup .....	27
3.2.3	Default settings .....	27
3.2.4	User management .....	29
3.2.4.1	Resetting a forgotten administrator user password .....	31
3.2.5	Locking the system .....	31
3.2.6	Updating the firmware .....	32
3.3	Media utilities (traditional media) .....	34
3.3.1	Eject .....	34
3.3.2	Content Breakdown .....	35
3.3.3	Reconfigure .....	37
3.3.3.1	Remove HPA/DCO/AMA .....	37
3.3.3.2	Wipe .....	38
3.3.3.3	Enable Tableau Encryption .....	41
3.3.3.4	Format Filesystem .....	42
3.3.4	Encryption Unlock .....	42
3.3.4.1	Opal encryption .....	43
3.3.4.2	BitLocker encryption .....	44
3.3.4.3	APFS encryption .....	45
3.3.5	Removing drive capacity limiting configurations .....	46

3.3.5.1	Volatile HPA removal .....	47
3.3.5.2	Non-volatile HPA/DCO/AMA removal .....	48
3.3.6	Blank checking .....	48
3.3.7	Browse Filesystem .....	49
3.3.8	SMART data .....	50
3.3.9	Export .....	51
3.4	Connecting drives .....	51
3.4.1	Source drives .....	52
3.4.2	Destination drives .....	53
3.4.3	Accessory drives .....	53
3.4.4	Drive detection .....	53
3.5	Turning off your unit .....	55
<b>4</b>	<b>Using the OpenText Forensic TX2 Imager .....</b>	<b>57</b>
4.1	Navigating OpenText TX2 features and options .....	57
4.2	Preconditions checking .....	58
4.3	Duplicating .....	58
4.3.1	Cloning .....	58
4.3.2	Physical imaging .....	59
4.3.3	Performing a duplication .....	60
4.3.3.1	Files created during disk-to-file duplication .....	66
4.3.4	Using the Automated Acquisition mode .....	67
4.3.5	Duplication over a network .....	73
4.3.5.1	Adding an iSCSI target .....	73
4.3.5.2	Adding a CIFS share .....	74
4.3.6	Pausing and resuming a duplication job .....	75
4.3.7	Spanning a destination drive .....	78
4.4	Hashing .....	79
4.5	Logical imaging .....	82
4.5.1	Performing a logical image acquisition .....	83
4.5.2	Include/exclude criteria .....	88
4.5.2.1	File type .....	90
4.5.2.2	Path .....	91
4.5.2.3	Folder .....	93
4.5.2.4	File size .....	93
4.5.2.5	File date .....	94
4.5.2.6	Hash database .....	94
4.5.3	About the logical imaging process .....	95
4.5.3.1	Logical image job status .....	95
4.5.3.2	Files created during logical imaging .....	96
4.5.3.3	Logical image verification .....	96
4.5.3.4	Advanced logical imaging setup .....	97

4.5.4	File MIME types .....	98
4.5.5	Folders .....	99
4.5.6	Source file metadata .....	100
4.6	Verifying .....	101
4.7	Browsing .....	102
4.7.1	Viewing text and image files .....	102
4.8	Restoring .....	103
4.9	Mobile backup acquisition .....	104
4.9.1	Connecting and detecting mobile devices .....	105
4.9.2	Mobile device details .....	106
4.9.3	Mobile device media utilities .....	107
4.9.4	Performing a mobile backup acquisition .....	108
4.9.5	Files created during mobile backup acquisition .....	111
4.9.5.1	Files created during iOS device backup acquisition .....	111
4.9.5.2	Files created during Android device backup acquisition .....	115
4.10	Viewing sources and destinations .....	118
4.10.1	Encryption detection .....	119
4.10.1.1	Opal encryption .....	120
4.10.1.2	Apple Core Storage and FileVault 2 .....	121
4.10.2	RAID detection .....	122
4.11	Logs module .....	123
4.11.1	HTML logs .....	125
4.11.2	Sample logs .....	125
4.11.3	Filtering logs .....	129
4.12	Remote web interface .....	130
4.12.1	SSL certificate setup and installation .....	131
4.12.2	Remote access .....	132
<b>5</b>	<b>Adapters .....</b>	<b>135</b>
5.1	PCIe SSD adapters .....	136
5.2	PCIe IDE adapter (TA7-5) .....	137
5.2.1	Using the TA7-5 .....	137
5.3	PCIe FireWire adapter (TA7-9) .....	137
5.4	Adapting SAS drives .....	138
5.5	Apple Target Disk Mode acquisition adapters .....	138
5.5.1	FireWire adapter cable .....	138
5.5.2	Thunderbolt 2 adapter cable .....	139
<b>6</b>	<b>Specifications and troubleshooting .....</b>	<b>141</b>
6.1	Specifications .....	141
6.2	Troubleshooting common problems .....	143
6.2.1	Power supply issues .....	143
6.2.2	Thermal issues .....	144

6.2.3	Problems with drive detection .....	144
6.2.4	Problems detecting Apple devices in target disk mode .....	146
6.2.5	Real-time clock data retention issue .....	147
<b>7</b>	<b>Appendix A: Logical Imaging Engine hash databases .....</b>	<b>149</b>
7.1	Hash database format specifications .....	149
7.2	Example script 1: create_othd.py .....	151
7.3	Example script 2: describe_othd.py .....	151

# Chapter 1

## Preface

This guide presents a wide range of technical information and procedures for using the OpenText™ Forensic TX2 Imager. It includes the following chapters:

- “[Overview](#)”: Provides general information about this product, as well as unpacking, starting up, navigating the touchscreen display, and reading the LEDs.
- “[Configuring the OpenText Forensic TX2 Imager](#)”: Provides system overview information, as well as procedures for configuring and connecting this product.
- “[Using the OpenText Forensic TX2 Imager](#)”: Provides detailed information and procedures for its operation.
- “[Adapters](#)”: Describes the adapters that extend the drive acquisition options and destination drive capabilities of the OpenText Forensic TX2 Imager.
- “[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), go to OpenText 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 defining 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.

OpenText Forensic Equipment products report drive capacities and transfer rates according to the industry standard *powers of ten* convention. In OpenText TX2 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.



## Chapter 2

# Overview

The OpenText Forensic TX2 Imager is a powerful, yet intuitive, forensic imager that offers superior local and networked imaging performance with no compromises. The touchscreen user interface is easy to use and provides a familiar user experience, similar to modern tablets and smartphones.

This product 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 (Gen 3 x4), USB (3.2 Gen2), SATA, SAS, FireWire, IDE, and network shares (iSCSI and CIFS).  
 **Note:** PCIe, SAS, FireWire, and IDE adapters (sold separately) are required to image these drive types.
- Output to PCIe (Gen 3 x4), USB (3.2 Gen2), SATA, and network shares (iSCSI and CIFS).
- Support for cableless, toolless, SATA destination drive connections via the optional drive bay (TX2-S1), which also provides drive cooling.
- Compatible with the following drive adapters: TA7-1, TA7-2, TA7-3, TA7-4, TA7-5, TA7-7, TA7-9, TDA3-1, TDA3-2, TDA3-3, and TKDA3-LIF.
- Clearly labeled and color-coded source (write-blocked) and destination (read/write) ports.
- 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 use a file hash database to search for specific files on a source drive. This can be used to eliminate files of non-forensic interest (for example, operating system files) to reduce the evidence set or to find specific files of forensic interest such as known illegal images.
- Two high performance 10-gigabit Ethernet ports (with 1 GbE auto-negotiation), which allow connectivity to two different local-area networks and/or network-attached storage devices. OpenText supports Ethernet port bonding (that is, link aggregation), which combines its two Ethernet links into a single logical interface, to increase bandwidth and/or provide redundancy.
- Browser-based remote user interface to any number of network-connected OpenText TX2 units, with the ability to directly download selected files to the remote system.
- The ability to export locally attached media as an iSCSI share for remote, network-based acquisition.

- The ability to enable and configure 802.1X network authentication, to strengthen network access security.
- The ability to acquire backup files from iOS- and Android-based mobile devices (including tablets).
- The ability to mount USB Still Image Class devices that support PTP (Picture Transfer Protocol) and MTP (Media Transfer Protocol), such as mobile devices and digital cameras, to enable on-screen browsing/triage and logical acquisition of any exposed files.
- The ability to automatically acquire drives connected to the imager's source ports, based on predefined job settings.
- The ability to duplicate a source drive to up to four destination drives (locally connected and/or network shares).
- The ability to simultaneously run multiple jobs of any type (clone, physical image, or logical image) to any available destination media/shares.
- The ability to resume a job that failed due to running out of space on a destination, by adding a new destination drive (known as destination drive spanning).
- Automatic assessment of available system resources, to balance active/queued jobs for maximizing jobs efficiency.
- The ability to manually reorder queued jobs and start them regardless of resource availability.
- 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 prior to physical removal.
- The ability to power down the system after the last active job is complete.
- User management – create, delete, and manage user profiles, including support for PIV Smart Card multi-factor authentication via YubiKey devices.
- Superior data transfer rates even while performing calculations of MD5, SHA-1, and SHA-256 hash values on multiple active jobs.
- Industry-leading OpenText Tree Hashing support, which allows for massive parallelization of data block hashing and unmatched imaging performance.
- The ability to readback verify sequentially-hashed acquisition file sets, in a block-based and parallel manner, to maximize the readback verification performance.
- 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 detect/acquire multiple namespaces on NVMe SSDs.

---

- The ability to detect and warn of the presence of detached NVMe namespaces and allow their attachment to enable acquisition of otherwise obscure evidence.
- The ability to unlock Opal-compliant self-encrypting drives (SEDs) and BitLocker encrypted drives/partitions, to enable unencrypted source media acquisition/ browsing and as an alternative to VeraCrypt based 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 the touchscreen user interface.
- Gallery View with one-touch scrolling, to allow rapid viewing/triage of image type files on a per folder basis.
- 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.
- Convenient and configurable destination drive *Reconfigure* utility that allows for removing HPA/DCO/AMA, wiping, encrypting, and/or formatting all in one job.
- 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 removal, DCO/AMA shelving, and trim support for the creation of a destination DCO or AMA.
- The ability to update the 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.
- The ability to put the unit into *stealth mode* for situations where bright LCD screens and loud job alerts may be undesirable.
- Regular and free firmware updates, available on OpenText My Support.
- User interface localization support for German, English, Spanish, French, Korean, Portuguese, Russian, Turkish, and Chinese languages, including virtual keyboard support for user inputs.

This product can operate as a standalone device (as shown in the following image), or with a destination imaging bay (TX2-S1).



The following image shows the product's left (source) side (write blocked).



The following image shows the product's right (destination) side (read/write).



## 2.1 Kit contents

OpenText TX2 ships in a boxed kit (with custom foam) that includes the following items.

Model #	Quantity	Description
TX2	1	OpenText Forensic TX2 Imager
TP8	1	Provides power to TX2, the optional TX2-S1 Drive Bay, and the attached 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-R4	4	8" unified SATA data and power adapter cable
TC-PCIE4-8	2	8" PCIe Gen3+ adapter cable. For use with OpenText Forensic PCIe Gen3+ adapters. For more information, see <a href="#">"Adapters" on page 135</a> .
TCA-USB3-AC	2	4" USB 3.2 Type A to Type C adapter cable
TPKG-VCT-5	1	Five-piece Velcro cable tie kit
TPKG-CLOTH	1	Micro fiber screen cleaning cloth
TX2-QRG	1	<i>Quick Reference Guide</i>

Do not discard the OpenText TX2 foam packaging, as it is designed to fit several industry-standard hard-sided carrying cases (for example, the *Pelican 1500*). If you received this kit in the cardboard box shipped by OpenText, you can reuse the stacking foam inserts in your own hard-sided case.

## 2.2 Navigating the touchscreen display

Use the touchscreen display to navigate within the user-friendly interface and choose or modify options. Use the touchscreen keyboard or an external USB keyboard to enter alphanumeric text, when prompted. For more information, see ["USB keyboard and mouse support" on page 20](#).

## 2.2.1 HOME screen

The **HOME** screen includes icons for initiating the following functions:

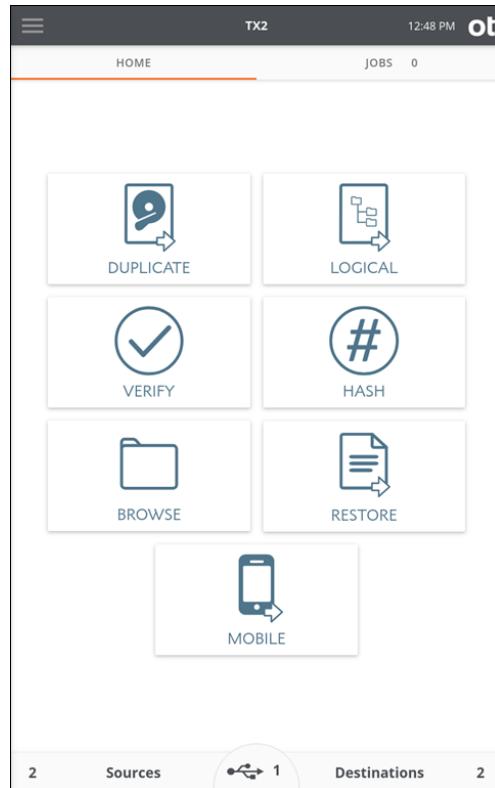
- [Duplicate](#)
- [Logical](#)
- [Verify](#)
- [Hash](#)
- [Browse](#)
- [Restore](#)
- [Mobile](#)

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.

The following image shows the **HOME** screen, with two source drives, two destination drives, and one accessory drive attached.



Tap one of the icons on the **HOME** screen 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 return to the previous screen or to the **HOME** screen.

The top navigation bar provides buttons to quickly access the side navigation menu , the **HOME** screen, and to view the current time. The **TX2** model name in the top navigation bar links to the **HOME** screen.

 **Note:** In the event of unusual cooling conditions, a warning triangle will be shown in the top navigation bar, near the time indicator. This warning is not displayed during standard operating conditions. For more information, see “[Troubleshooting common problems](#)” on page 143.

The main screen includes the following tabs:

- **HOME** tab: shows the Home screen or one of the many other screens.
- **JOBS** tab: Always shows the Jobs summary screen and the current number of total active and queued jobs.

## 2.2.2 Side navigation menu

Click the menu icon  in the top navigation bar to display the side navigation panel, which provides a menu of additional options and information. For more information about this menu, see “[Configuration settings](#)” on page 22.

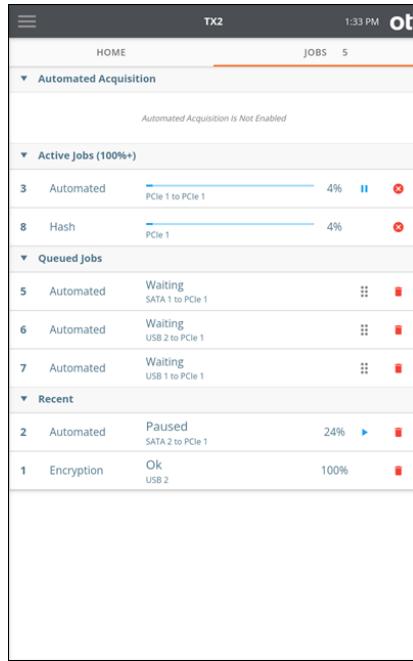
## 2.2.3 JOBS tab

The JOBS tab provides a convenient way to keep track of your active jobs, queued jobs, and recent work done. It includes a counter in the top tab area. When OpenText TX2 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 tab has the following areas.

<b>Automated Acquisition</b>	This area indicates if the <b>Automated Acquisition</b> mode is turned on. When it is turned on, the count of any duplication jobs run during that automated session will be displayed. For more information about this valuable, time-saving feature, see “ <a href="#">Using the Automated Acquisition mode</a> ” on page 67.
<b>Active Jobs</b>	This area displays any active jobs. The imager will automatically decide when to start jobs or enqueue them, based on available system resources. Queued jobs start as soon as system resources become available, or when the source or destination resource for which the job is waiting becomes available. The automated resource assessment system can be overridden by manually dragging a job from the queue into the <b>Active Jobs</b> area. As jobs are completed, they are moved to the <b>Recent</b> jobs area. The detailed <b>Job Status</b> screen for any given job can be viewed by tapping on the job tile shown in the JOBS tab.
<b>Queued Jobs</b>	This area displays any queued jobs. Jobs are queued in the order in which they were entered.  You can reorder queued jobs by dragging individual jobs, to place them in the order you prefer. On the touchscreen, press and hold the drag icon  of the job you want to move, then drag the job to the desired position in the queue and release. Queued jobs may also be dragged into the <b>Active Jobs</b> area, to force them to start, regardless of available system resources.
<b>Recent</b>	This area displays any recently completed jobs.   <b>Note:</b> This area is cleared out upon power-cycling the imager. A complete and non-volatile list of all jobs can always be found in the side navigation bar, by tapping the <b>Logs</b> button.
<b>Media Utilities</b>	This area displays media utility operations and appears only when such an operation is active. Media utility operations will move to the <b>Recent</b> area when they are complete. Media utility jobs cannot be queued like forensic jobs.

The following image shows an example of the acquisition jobs.



In this example, there are two active jobs – a hash and an automated duplication job. There are three 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 or for manually starting a queued job. In the **Recent** area there are two jobs, one that completed successfully and one that was paused. Since the pending job count includes both active and queued jobs, the **JOBS** tab counter reads 5.

In this example, *Job 2* has moved from the **Active Jobs** area down to **Recent** (due to being paused), which allowed *Job 8* to start as shown in the **Active Jobs** area. For more information regarding that feature, see “[Pausing and resuming a duplication job](#)” on page 75.

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 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*.

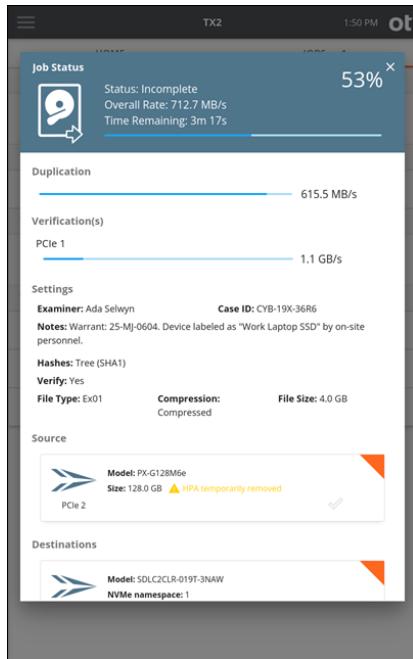
Tapping the **Play** button  (or tapping the **Resume Job** button in the header of the **Job Status** screen) will display a **Resume Duplication** screen.

For more information about pause and resume, see “[Pausing and resuming a duplication job](#)” on page 75.

## 2.2.4 Job status

After a job starts, its **Job Status** screen is automatically displayed. This 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.

The following image shows an example of an active **Job Status** screen.



**Note:** A solid orange triangle in the corner of a drive tiles indicates that the drive is currently being used by 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 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 header area at the top of the status screen.

Clicking the **View Log** link displays the detailed log for that job. You can 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 screen. For example, if you 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:** The **Job Status** screen can also be viewed for completed jobs that are shown in the **Recent** area of the **Jobs** tab. Tapping the job row from the **Recent** jobs area will display the final **Job Status** screen for that completed job, including a link to the job log.

## 2.2.5 Quick Reference Guide

This product ships with a *Quick Reference Guide*, which illustrates the firmware update procedure, drive connections, status LEDs, power button, cable/adapter recommendations, and tips for getting started. Keep this card handy as you familiarize yourself with this product.

## 2.3 Reading the status LEDs

### Power supply LED

The TP8 power supply DC cable has a blue LED ring near the end of the barrel connector, which indicates that the power supply is connected to an AC power source and functioning properly.

### On/Off indicator LED

The illuminated power switch is located in the top-left corner of your device and it displays a white LED when the unit is on.

### Activity LED

The multi-color activity LED is located in the lower-right corner of your device. The following table provides details for interpreting the status of this LED.

Status	Activity LED
The unit is booting up.	White
A power issue is detected.	Blinking white
The unit is idle.	Off
An operation is in progress.	Blue
An operation completes successfully.	Blinking green
An operation fails.	Blinking red

### Network interface LED

The dual Ethernet connector is located on the back of your device and it has two LEDs for each port. The following table provides details for interpreting the status of these network interface LEDs.

Status	Link Status/Activity LED (Left side of the unit, looking at the connectors from the rear of your device; yellow in color)	Link Speed LED (Right side of the unit, looking at the connectors from the rear of your device)
No Link	Off	Off
1 Gbps Link	On/Blink	Green
10 Gbps Link	On/Blink	Orange

## 2.4 Interpreting audio feedback

This product 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 sound volume in the side navigation menu, from **System Settings**.

## 2.5 On-screen warnings

When necessary, this product provides 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 and mouse support

You can plug a standard USB keyboard and/or mouse into either of the USB accessory ports on the front of your device. Some users find it more convenient to use an external keyboard/mouse to enter data, instead of using the touchscreen and on-screen virtual keyboard. Language-localized keyboard support is limited to the virtual keyboard only.



**Note:** If you prefer to use a wireless keyboard/mouse, plug the USB wireless adapter into either of the front USB accessory ports, and it should automatically pair with the keyboard/mouse and start working. There are many vendors of wireless keyboards/mice, and some may not be compatible with OpenText TX2. Contact OpenText Customer Support for wireless keyboard/mouse recommendations.

## Chapter 3

# Configuring the OpenText Forensic TX2 Imager

This chapter describes how to configure the OpenText Forensic TX2 Imager.

### 3.1 Startup sequence

1. When you turn on the unit, the activity LED (in the lower-right corner of the device) will immediately turn on (and display white color), which indicates that the unit is starting the boot sequence.
2. After a short time, the fan will turn on and then an OpenText branded splash screen will appear during the rest of the boot sequence.
3. Once booted, the unit displays the **HOME** screen (or the user login screen, if the auto-login is turned off), and then it sequentially powers on the drive ports.

Once all of that is complete, your OpenText TX2 is ready for use.

For situations or locations where drawing attention to your presence is undesirable, OpenText TX2 can be put into *stealth mode*. This mode will deactivate all audio alerts, deactivate the status LED, and set the LCD brightness to a minimum setting.

#### To turn on the stealth mode:

1. Open the SSD access door on the bottom of the unit and change the first DIP switch (labeled 1 on the DIP switch block) to the OFF position.
2. Replace the SSD access door.

The unit will power up in stealth mode until the switch is returned to its default (ON) position.



**Note:** Do not change the state of any of the other DIP switches when turning on the stealth mode. While switch 1 is the only assigned switch at initial release of this product, other switches may be activated in future firmware updates. Leaving all other switches in their default (ON) state will prevent undesirable unit behavior after future firmware updates.

## 3.2 Configuration settings

The 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, 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 operational **Defaults settings** 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 serial number, network MAC address, firmware build ID and version, firmware SHA-256 value, and copyright and licensing information. This area also includes the firmware update utility.
- **User:** Access the **User Management** screen for the current user.

### 3.2.1 System settings

Tap **System settings** to display the **System Settings** screen.

The **System Settings** screen allows you to configure 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 restores all system settings to their factory defaults and deletes all user-generated data. This includes job logs, 802.1x certificates, PIV Smart Card (YubiKey) certificates, SSL certificate, saved logical image searches, user configuration, and bookmarked CIFS/iSCSI logins. It is recommended that you make notes regarding any of that information and export all logs to an external device, before initiating a factory reset.

Tap the toggle buttons to activate or deactivate 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**. Once the area expands, tap a setting value to reveal and select from a drop-down list.

To set the OpenText TX2 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 (<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 the OpenText TX2 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** screen.

This screen includes the following areas:

- **Port Selection** allows you to select an Ethernet port from the list of available options. If port bonding is desired, tap **Add Bond** in this area to initiate this action.



**Note:** All of the settings on this screen are done on a per Ethernet port basis or per the single bonded network (if port bonding is utilized). Before configuring any network settings, make sure to select the appropriate Ethernet port or Bonded Network in the **Port Selection** area.

- **Current status** displays network-related information and the current connection status.
- **Configuration** allows you to set the **Bond Mode** and **Bond Transmit Policy** (if applicable), **IP address**, **MTU** (maximum transmission unit) value, and **Custom hostname**.

The following **Bond Mode** options are available.

Bond Mode	Best for	Description	Customer value	Use case	Bond Transmit Policy	Switch Port Aggregation
<b>Active Backup</b>	Field work, unattended imaging.	Primary port only is actively used for traffic. If the active port fails, another port becomes active. Does not require switch configuration.	Ensures uninterrupted imaging.  No network configuration required.  Simple and reliable.	<b>Overnight or unattended imaging.</b>  If a network cable or switch port fails during a long acquisition, imaging continues automatically without interruption.	N/A	N/A
<b>Balance XOR</b>	Managed switch environments.	Traffic is distributed across ports using a hash-based algorithm.	Load balancing + fault tolerance.  Predictable traffic distribution.	<b>Structured network environments.</b>  Labs with managed switches configured for port aggregation that want both performance and redundancy.	Yes	Yes/ Ethernet Channel
<b>802.3ad</b>	Enterprise labs, NAS storage.	Uses IEEE 802.3ad (LACP) to combine multiple Ethernet ports into one logical high bandwidth link.	Redundancy + performance.  Enterprise-grade networking.	<b>High-speed forensic lab or data center.</b>  OpenText TX2 connects to a managed switch and NAS/SAN storage where large evidence files must be transferred quickly and reliably.	Yes	Yes, with Link Aggregation Control Protocol (LACP)
<b>Balance TLB</b>	Standard labs, no switch changes.	Adaptive Transmit Load Balancing (TLB) distributes outgoing traffic across multiple ports based on load. Incoming traffic uses a single active port.	Improved performance.  No special switch configuration needed.  Ideal for standard lab networks.	<b>Busy forensic lab.</b>  Multiple imaging jobs are writing data at the same time. OpenText TX2 balances outgoing traffic across ports to reduce bottlenecks.	N/A	N/A

Bond Mode	Best for	Description	Customer value	Use case	Bond Transmit Policy	Switch Port Aggregation
Balance ALB	Multi connection environments.	Adaptive Load Balancing (ALB) balances both outgoing and incoming traffic across multiple Ethernet ports.	Higher overall throughput. No switch-side configuration required. Efficient use of available network links.	<b>Multi connection workflows.</b> Handling multiple acquisitions or simultaneous transfers where balanced send and receive traffic improves responsiveness.	N/A	N/A

The following **Bond Transmit Policy** options are available.

Bond Transmit Policy	Description
Layer 2	It is recommended to use this option when OpenText TX2 and the target servers are on the same network segment (also referred to as broadcast domain).
Layer 2 and 3	It is recommended to use this option when OpenText TX2 and the target servers are on different network segments (or broadcast domain).

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 1,500. If your device is attached to a network that supports jumbo frames, change the MTU value to 9,000, which may allow for much faster network transfer speeds.



**Note:** The maximum allowed MTU settings on this device is 9,000. Attempting to manually set the **MTU** to a value higher than 9,000 will result in no change to the setting. This prevents slow network performance due to mismatched MTU settings.

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

- **802.1X** allows you to [configure the 802.1X network authentication](#).
- **HTTPS certificate** allows you to [setup the HTTPS certificate](#).

### 3.2.2.1 Configuring 802.1X network authentication

OpenText TX2 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 configure your device for use on a network with 802.1X authentication:

1. Tap **Edit** in the bottom right of the 802.1X settings area, to choose one of the three EAP types (**TLS**, **TTLS**, or **PEAP**).
2. Tap **Identity** to enter your 802.1X identity (required).
3. One or more certificates (depending on the EAP type and other settings) may need to be loaded onto your device before attempting to authenticate on the network. The certificate loading process is straightforward:
  - a. Store the required certificates on a USB memory device, and then insert that device into an OpenText TX2 USB Accessory port.
  - b. 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 screen will appear, allowing navigation to the appropriate memory device and certificate file.
- c. Select the desired certificate file from the browser and tap the **Install** button.

4. Each EAP type has additional requirements and configuration settings depending on the type selected, as follows:
  - **TLS:** The OpenText TX2 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 OpenText TX2, to turn on the 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 OpenText TX2, to turn on the 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 on the top navigation bar, and an **Add Password** button becomes active in the settings summary area.

5. Tap the navigation bar icon or **Add Password** to enter an 802.1X passphrase/password.

 **Note:** 802.1X passphrases are required to decrypt encrypted private keys. These passphrases can be between 4 and 1,023 characters in length.

6. Tap **SUBMIT** to begin the authentication procedure.

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

### 3.2.2.2 HTTPS certificate setup

OpenText TX2 generates an SSL certificate on startup. You can use this certificate, manually generate a new certificate, or install your own certificate.

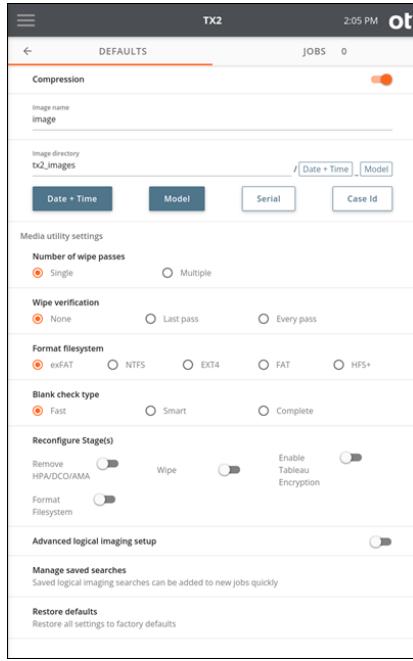
The bottom area of the **Network Settings** screen shows the current SSL certificate information and provides options for manually generating a new OpenText TX2 certificate or installing a custom certificate. For more information about SSL certificate options, see ["Remote web interface" on page 130](#).

### 3.2.3 Default settings

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

Several different entry methods are used for the settings on this screen, including direct data entry, sliders, option buttons, and an **Image Directory** name builder area. As shown in the following image, the **Date + Time** and **Model** directory path element boxes have been selected (in that order), so therefore the image directory path is `tx2_images/<Date+Time>`.

 **Note:** The source drive 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. OpenText TX2 checks for this situation at job startup, and, if it detects a file name conflict, it automatically appends sequential numbers to the subsequent filenames. 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. For more information about basic and advanced setup modes, see [“Logical imaging” on page 82](#).

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 OpenText TX2 units. For more information about search setup and utilization, see [“Logical imaging” on page 82](#).

Default settings can be restored to their factory set values. 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 credentials to limit access to the available devices. Also, with the addition of remote access capability, the ability to set user credentials has become a security requirement. OpenText TX2 ships with a default user named *User1*, which has administrator rights but no authentication and no remote access. *User1* also has **Auto Login** turned on (which is only possible for users with no set password). These defaults make this device 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 add remote access capability to any user (even default *User1*), tap the **Users** item from the side navigation menu. A list of defined users is displayed on the left tab (default *User1* is the base user). 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. If a user with no administrator rights is logged in, the **Users** button on the side navigation menu will be dimmed and unusable.

Tapping on any user in the list will show a **User Management** screen which contains the available options for each user.

On this screen, an administrator can change any user's authentication configuration (username, password, multi-factor authentication), set administrator rights, allow remote access, allow tree hashing enablement, and enable auto login.



**Note:** Remote access is allowed only for users that have authentication setup. Auto login is allowed only for users that do not have authentication setup. Turning off administrator rights for a currently logged-in administrator is not allowed.

While using the **User Management** screen to set up new and edit existing users, you will be prompted to re-enter your password for each change to the credentials for a given user. This ensures no unauthorized changes can be made. For other User Management configuration changes (not related to authentication) there is a 30-second timeout that allows additional changes to be made without password re-entry.

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, a **Logout** button is also visible. This button logs out only that instance (local or remote) of that user.

To create a new user, tap the **Create New User** button at the bottom-right of the **Users** list screen and enter the desired username. If authentication is not required for the new user, tap **Submit** and the **User Management** screen for that user will appear, allowing you to complete that user's configuration. If authentication is required, it can be set up on this initial user creation screen. OpenText TX2 supports two authentication methods, as follows:

- Password based authentication: An initial alphanumeric, case-sensitive password may be entered to allow user access.
- PIV Smart Card via YubiKey authentication: After configuring the YubiKey device for the new user (for more information, see <https://www.yubico.com/>), insert it into any USB port on your OpenText TX2. Within the new user creation screen, enter the PIN for the Smart Card.



**Note:** OpenText TX2 requires validation of the PIV Smart Card certificate chain up to the root certificate. Depending on your Smart Card configuration, that root certificate may reside on the user's YubiKey device (known as a self-signed configuration) or it may need to be separately installed onto your OpenText TX2. If you enter the Smart Card PIN during user configuration and a PIV Smart Card certificate chain (up to a root certificate) cannot be found, a failure message will appear, and the credentials will not be entered/changed for that user. To install a PIV Smart Card certificate chain (.pem file type), insert a USB drive into any OpenText TX2 port that contains the certificate file(s) (up to and including the root certificate), and then tap the **Manage Smart Card Validation Certificates** button on the main **Users list** screen. At the bottom of the Certificates screen, tap the **Install From File** button. A browse window will appear that will allow you to navigate to the certificate file(s) on the attached drive. Once located, tap on the desired certificate file(s) and then tap the **Install** button at the bottom right of that screen. Once this step is successfully completed, then the Smart Card PIN may be entered for the associated individual user(s) on their User Management screen.

Regardless of the chosen authentication method, once the appropriate credentials are configured, tap the **Submit** button at the bottom-right of the **Create New User** screen and the **User Management** screen for that user will appear, allowing you to complete that user's configuration.



**Note:** 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.

For more information about setting up users, contact OpenText Customer Support.

### 3.2.4.1 Resetting a forgotten administrator user password

This section provides instructions for resetting a forgotten administrator user password.



**Note:** If a non-administrator user has forgotten their password, the administrator can log in and change the non-administrator user's password. The only time a factory reset is required to recover a unit is if the designated administrator has forgotten their password.

**To reset a user password:**

1. Enter an incorrect password five times, consecutively.
2. Perform a factory reset, as prompted by the system.



**Important**

A factory reset will delete all user-sourced data. This includes all job logs, settings (system, network, and defaults), user management configuration, and network share bookmarks. For more information about the factory reset, see ["System settings" on page 22](#).

After completing the factory reset, all system settings are restored to their factory defaults. No password is required in this state as the default user (*User1*) will be configured for auto-login with no password.

### 3.2.5 Locking the system

OpenText recommends locking your device while unattended, to ensure that no settings are changed and your active/queued jobs are not altered in any way.

**To lock your system:**

1. Tap **Lock System** on the side navigation menu.
2. In the screen that appears, enter the desired six-digit personal identification number (PIN), then tap **Submit** to begin the locking process.



**Note:** The  button on 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.

The system will prompt you for a second entry of the same PIN to confirm the desired digits have been entered.

3. Enter the same PIN again.

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.

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

**To unlock the system:**

- Enter the current PIN and then tap **Submit** in the lower right corner of the screen.

### 3.2.6 Updating the firmware

The OpenText TX2 firmware is stored on an m.2 NVMe SSD located on the bottom of the unit, behind a removable access door.



**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 using the web interface. It is recommended that all active users on a given device (local and remote) collaborate to ensure no jobs are running or will be started before a firmware update is initiated.

To update your device firmware, go to OpenText My Support <https://support.opentext.com> and log in (or register) to access firmware package files for your OpenText Forensic TX2 Imager. Locate and download the most recent firmware package file (.tx2\_pkg) and then select one of the following methods.

**To update the firmware using the local package file:**

1. Copy the desired OpenText TX2 firmware package file to a USB drive or to a network share that will be accessible to your device.
2. Tap **Firmware version** (or **About**) on the side navigation menu, and then tap **BROWSE FOR FIRMWARE**.
3. In the file browse window, tap the desired drive or network file share and then use the **Browse** window to navigate to the folder that contains the firmware package file.
4. Select the desired firmware package file (.tx2\_pkg) and then tap **SELECT** in the bottom-right corner.

The browse window will automatically close, showing the name of the package file that was selected on the **About** screen.

5. **Optional** Tap **HASH FIRMWARE** to generate a hash of the selected firmware package file. This hash value can then be compared to the hash value from the source download page (from <https://support.opentext.com> ) providing confidence that you will be updating your unit with the exact desired firmware package file.
6. Once you are confident you want to proceed with the update, tap **UPDATE**.  
The firmware update process starts.



**Note:** Depending on the state of the sub-system firmware packages on the unit, your device may perform multiple reboots/power-cycles. Do not interact with the unit or power it down during this time. Once the login screen or Home screen appears (depending on user configuration), your updated device is ready to use.

#### To update the firmware using the remote user interface:

1. Establish a remote user interface connection to the OpenText TX2 to be updated.
2. On the remote user interface, open the side navigation menu and select **About**.
3. Tap the **SELECT FILE** button in the **Upload device firmware** area.  
This launches a file browser window on your host system, allowing you to navigate to and select the desired firmware package file.
4. Check the name of the selected package file in the **Upload device firmware** area, then tap **UPLOAD DEVICE FIRMWARE** to initiate the firmware file upload.  
A progress bar appears, indicating that the file upload is in progress.



#### Caution

Navigating away from this page when a firmware file upload is in progress will stop the firmware update process.

Once the file has been fully uploaded, OpenText TX2 will automatically update its local firmware and reboot the unit.



**Note:** Depending on the state of the sub-system firmware packages on the unit, your device may perform multiple reboots/power-cycles. Do not interact with the unit or power it down during this time. Once the login screen or Home screen appears (depending on user configuration), your updated device is ready to use.

Updating the firmware with either of these methods will leave all previously stored user data (settings, logs, saved searches, HTTPS/802.1x certificates, etc.) intact. To wipe all user data from the system drive, perform a Factory Reset which is available at the bottom of the **System Settings** screen.

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 verification that the proper firmware version is running and that it has not been altered.

## 3.3 Media utilities (traditional media)

Accessible from the Sources, USB Accessories, or Destinations buttons at the bottom of the **Home** screen (and all locations that provide drive lists), OpenText TX2 provides the following media utilities for all traditional media types (mobile devices excluded):

- **Eject**
- **Content Breakdown**
- **Reconfigure** (destination/accessory only) – Includes HPA/DCO/AMA removal, Wipe, Enable Tableau Encryption, Format Filesystem
- **Encryption Unlock** (source or destination; Tableau, Opal, BitLocker, APFS)
- **HPA/DCO/AMA Disable** (ATA source drives only)
- **Blank Check**
- **Browse Filesystem**
- **SMART** (ATA drives only)
- **Export** (iSCSI target)



**Note:** Mobile devices connected to OpenText TX2 have host system interactions and capabilities that are very different from traditional media devices (HDDs, SSDs). The media utilities listed in this section are specific to traditional media devices. See “[Mobile backup acquisition](#)” on page 104 for information specific to that type of job.

### 3.3.1 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 OpenText TX2. 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. Ejection of media being used in an active job will not be allowed until the job is complete.



**Note:** OpenText TX2 supports PCIe drive hot-swap. Ejection of PCIe drives is required prior to removal from a powered-on system. Failure to do so can result in unpredictable system behaviors.

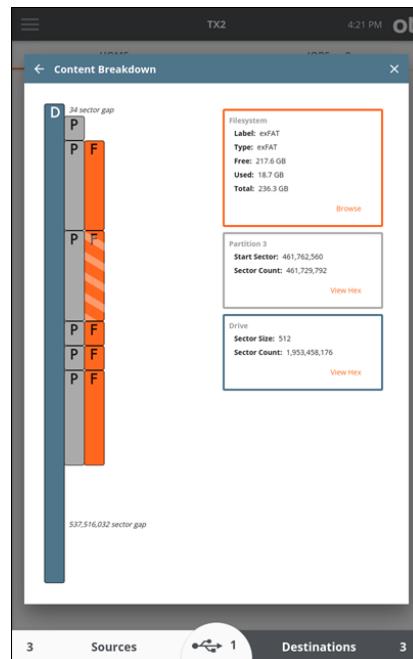
In addition to quiescing the drive for system removal, ejecting will issue an ATA spin down command to drives that may support it. Spinning down rotating hard disk drives is recommended to minimize the chance of platter damage upon physical removal of the drive from the system. Some drives do not support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

### 3.3.2 Content Breakdown

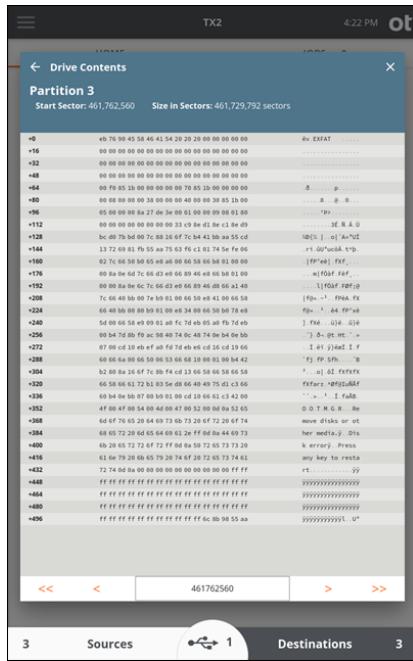
The **Content Breakdown** utility is available for all mounted drives and offers a unique perspective on drive geometry that is quite intuitive and powerful. This screen provides a physical map style overview of the selected drive and all its partitions and filesystems. Tapping on each element (rectangular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

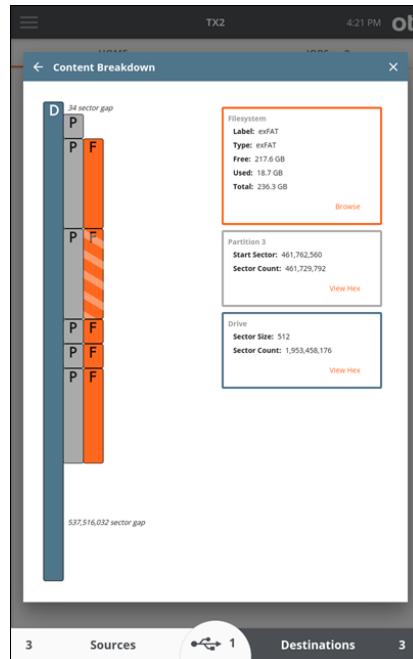


To view the raw hex data for a given drive/partition, tap the **View Hex** button within the information box. The following image shows a sample hex view window.



Each row shows 16 bytes of hex data, followed by each byte's ASCII equivalent. The number with the plus sign (located at the beginning 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 single arrow) or one sector backward (left single 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. A specific sector number can also be entered in the box at the bottom of the screen. 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, and free/used/total space) and a **Browse** button that opens the OpenText TX2 standard browse modal, as shown in the following image.



### 3.3.3 Reconfigure

The **Reconfigure** utility provides an easy way to perform the following actions on a destination or accessory drive all in a single job:

- **HPA/DCO/AMA removal**
- **Wipe**
- **Enable Tableau Encryption**
- **Format Filesystem**

#### 3.3.3.1 Remove HPA/DCO/AMA

If a destination or accessory drive has an HPA, DCO, or AMA drive capacity limitation set, it will be removed when this **Reconfigure** option is selected. For detailed information regarding these types of drive configurations, see “[Removing drive capacity limiting configurations](#)” on page 46.



**Note:** This specific function (within the **Reconfigure** utility on destination and accessory drives) is a non-volatile (or permanent) HPA/DCO/AMA removal.

### 3.3.3.2 Wipe

The **Wipe** portion of **Reconfigure** provides three wipe types for destination and accessory drives.



**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 TX2-S1 drive bay (which has active cooling) or an external drive cooler or fan when wiping media on your device, to help prevent thermal damage to drives.

<b>Overwrite</b>	This method writes known pattern data to every accessible region of a drive. This can be done with one pass or multiple passes. Verification is optional.
<b>Secure Erase</b>	This method is only available for ATA-based SSDs that support the command. OpenText TX2 only issues the <b>Secure Erase</b> command to the drive, and the rest is done by the controller on the drive. Verification is not supported with this method, because the state of the post-wipe data on the drive is drive manufacturer-specific.
<b>Sanitize</b>	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 – <b>Overwrite</b> and <b>Block Erase</b> . The OpenText TX2 only issues the <b>Sanitize</b> command to the drive, and the rest is done by the controller on the drive. Verification is optional.  <b>Note:</b> An active <b>Sanitize</b> 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. In general, 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 the OpenText TX2'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 allows the support of post-wipe verification. If you experience **Sanitize** wipe verification failures, contact OpenText Customer Support to report the specific make and model of the drive.

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

Option	Description
Overwrite - One Pass	<p>OpenText TX2 writes 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.</p> <p>Verification is optional.</p> <p> <b>Note:</b> 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.</p>
Overwrite - Multiple Pass	<p>OpenText TX2 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 0xFFFF.</p> <p>Verification is optional. If turned on, it can be configured to verify after each wipe pass or after only the last pass.</p> <p> <b>Note:</b> 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.</p>
Secure Erase (SSD only)	<p>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).</p> <p>Due to the indeterminate nature of the post-wipe data state, OpenText TX2 does not offer verification for Secure Erase wipes.</p> <p>Due to known issues with inconsistent and unreliable Secure Erase support on rotating drives (HDDs), OpenText TX2 only supports this feature on SSDs.</p> <p> <b>Note:</b> 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.</p> <p>OpenText TX2 will force the removal of any detected HPA/DCO/AMA configurations prior to starting a Secure Erase wipe, except for USB connected media. It cannot remove HPA/DCO/AMA configurations for USB-connected media, which means Secure Erase is not supported in that situation.</p>

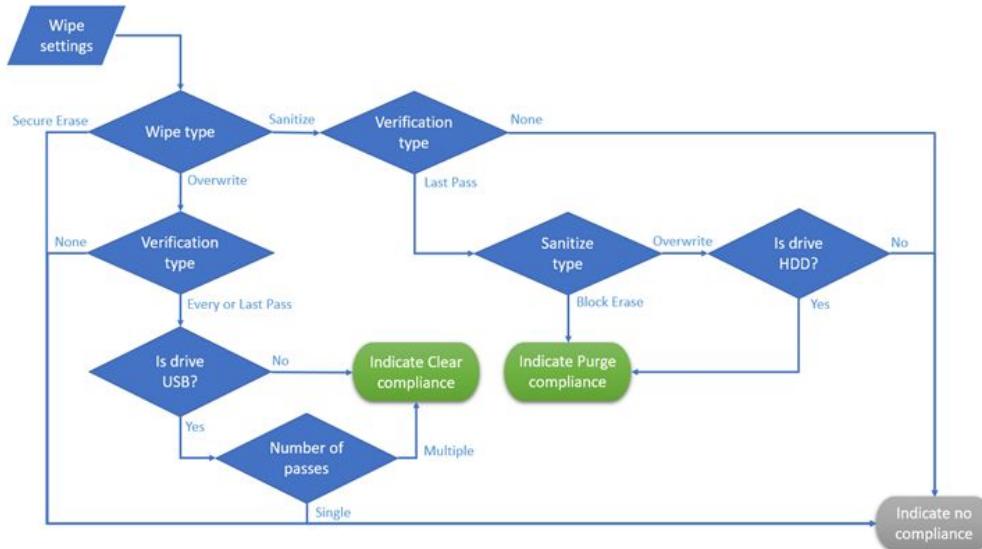
Option	Description
Sanitize – Overwrite	<p>The ATA and SCSI Sanitize – Overwrite command instructs the drive to overwrite all drive data (in both storage and on-drive cache) with zeros. This feature is typically implemented on HDDs, but is available on some SSDs.</p> <p> <b>Note:</b> For SSDs that support 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.</p> <p>OpenText TX2 will force the removal of any detected HPA/DCO/AMA configurations prior to starting a Sanitize – Overwrite wipe, except for USB connected media. It cannot remove HPA/DCO/AMA configurations for USB-connected media, which means Sanitize - Overwrite is not supported in that situation.</p>
Sanitize - Block Erase (SSD only)	<p>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.</p> <p> <b>Note:</b> Sanitize – Block Erase will erase all user-accessible drive space as well as over-provisioned space and any other space reserved by the drive's internal controller.</p> <p>OpenText TX2 will force the removal of any detected HPA/DCO/AMA configurations prior to starting a Secure Erase wipe, except for USB-connected media. It cannot remove HPA/DCO/AMA configurations for USB-connected media, which means Secure Erase is not supported in that situation.</p>

### NIST 800-88r1 Compliance

In the 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**. OpenText TX2 indicates compliance with these data sanitization methods when the appropriate wipe settings are selected – as indicated by the appropriate box turning green with a checkmark (at the top of the **NIST 800-88 r1 Compliance** area of the screen). Conversely, tapping either the **Clear** or **Purge** button (when this functionality is turned on) will automatically select compliant wipe settings.

OpenText TX2 can purge drives that support conformant Sanitize commands. The Purge **option** is deactivated if the **Sanitize** command is not supported by the drive.

The following flowchart depicts how OpenText TX2 determines wipe setting conformity to NIST 800-88 r1. This flowchart reflects handling of drives with no HPA/DCO/AMA settings.



### 3.3.3.3 Enable Tableau Encryption

OpenText TX2 can encrypt destination and accessory drives using a password-based XTS-AES whole disk encryption. This Tableau-based encryption is compatible with the OpenText Tableau Forensic TD2u Duplicator, OpenText Forensic TD4 Duplicator, OpenText Forensic TX1 Imager, and the open source *VeraCrypt* utility. This section only covers the encryption formatting option. For information regarding unlocking a drive with pre-existing Tableau encryption, see ["Encryption Unlock" on page 42](#).

 **Note:** The encryption formatting process overwrites the destination/accessory drive, so remember to copy any pre-existing drive data that is of value prior to encrypting the drive.



#### Caution

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

To remove encryption from a drive, connect the drive to the OpenText TX2 as a destination or accessory and wipe the drive.

 **Note:** The **Reconfigure** media utility is not available for use on unlocked Tableau-encrypted drives. If you wish to perform any other Reconfigure functions on a drive that you intend to encrypt, those should all be done in one Reconfigure job.

### 3.3.3.4 Format Filesystem

To perform a duplication to or save logs to a drive, you must format the destination or accessory drive with a recognizable file system. OpenText TX2 supports formatting destination/accessory drives with the following file system formats: **exFAT, NTFS, EXT4, FAT, or HFS+.**

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



#### Notes

- When FAT is selected as the filesystem type for a destination drive format, OpenText TX2 will format the drive as FAT32. However, job logs (including the format log) and all user interface elements will show this as FAT. That is because OpenText TX2 supports reading from all FAT formats (12, 16, and 32), and identifying them all as FAT is considered acceptable and accurate for filesystem identification purposes.
- OpenText TX2 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 Encryption Unlock

OpenText TX2 can unlock drives/volumes that have been encrypted with Tableau encryption, APFS, BitLocker, and Opal. APFS encryption support is limited to source drives, but Tableau, BitLocker, and Opal encrypted media can be unlocked regardless of which port they are connected to.

Whether dealing with Tableau, APFS, BitLocker, or Opal encryption on a given drive/volume, the same **Encryption Unlock** media utility is used to unlock it. A pull-down 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.

#### To unlock an encrypted drive/volume:

1. Select the encrypted entity you want to unlock, enter the password (or BitLocker recovery key), and then tap **UNLOCK**.

If successful, the progress bar turns green and a confirming message box appears.

2. Tap **OK** to close the message box and then close the **Encryption Unlock** screen to access the other functions with the now unlocked drive/volume.

Once unlocked, each drive/volume can be used with any supported operations, including browsing, imaging (physical, logical, or mobile), and any applicable media utilities.

While unlocking Tableau, APFS, BitLocker, and Opal encryption is simple and done using the same **Encryption Unlock** media utility, there are some notable differences

in how OpenText TX2 handles these types of encryptions, that warrant special consideration, as covered in the following sections.

### 3.3.4.1 Opal encryption

OpenText TX2 can unlock Opal Self Encrypting Drives (SEDs) that have had their encryption enabled in a Linux environment, as described in [“Encryption Unlock” on page 42](#). 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 OpenText TX2. 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. OpenText TX2 uses the Linux `SEUTIL` 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 `SEUTIL` function uses. Attempting to use a known password for such drives using OpenText TX2 will result in failed unlock attempts. If you encounter such a situation, contact OpenText Customer Support.

An additional consideration for Opal drives is a unique configuration that exposes a Shadow MBR. This Shadow MBR can be enabled by drive/system manufacturers to initially identify the drive as a small, non-encrypted volume, which overrides the actual MBR information.

A typical use case for this configuration is to enable system manufacturers to request credentials from a user before revealing the actual MBR information on the drive. Regardless of the use case, it is important to be able to identify situations where only the Shadow MBR is revealed, to make it clear that the entire drive contents are not being seen. OpenText TX2 will detect when an Opal Shadow MBR is enabled and clearly inform of its presence. The lock icon will show in the affected drive tile in the **Sources** list, and the presence of an Opal MBR will be explicitly called out in the **Drive Details** screen. Note that the Shadow MBR configuration is essentially a unique form of a locked Opal drive, therefore unlocking the Opal encryption on OpenText TX2 will disable the Shadow MBR (regardless of the underlying encryption state) and make the full, unencrypted drive contents available for triage/acquisition. Also, Opal encryption unlock (including Shadow MBR disablement) is a volatile change, meaning that the drive will revert to its original configuration after it is power-cycled.



#### Caution

Docking station devices that contain Opal drives must support ATA command pass-through for the OpenText TX2 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 OpenText TX2 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 your device display, removing the drive from the enclosure and connecting it directly to OpenText TX2 may yield the desired outcome.

### 3.3.4.2 BitLocker encryption

OpenText TX2 can unlock drives and partitions that are encrypted with Microsoft BitLocker, as described in [“Encryption Unlock” on page 42](#). 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 your device 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 your device.

Unlike an Opal SED, a BitLocker drive can be physically imaged (E01, Ex01, DD, and DMG) or cloned in its encrypted state. Such evidence can then be used with forensic investigation tools, such as OpenText 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. While OpenText TX2 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. The OpenText TX2 method to unlock a disabled BitLocker drive/partition is similar to the method described in this section, except that the **Password/Recovery Key** field will be disabled. Instead, OpenText TX2 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. OpenText TX2 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.

### 3.3.4.3 APFS encryption

OpenText TX2 can unlock drive volumes that are encrypted with APFS, as described in ["Encryption Unlock" on page 42](#). 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 OpenText TX2 handles the various encryption methods that it can unlock – APFS, BitLocker, Opal, and Tableau encryption. The following table summarizes how each of these encryption types will appear or be used in various operations, in both their locked and unlocked states.

		Operation				
		Logical Image (Source)	Physical Image/ Clone (Source)	Wipe (Dest)	Format (Dest)	Blank check (Source/ Dest)
<b>APFS</b>	<b>Locked</b>	n/a (no filesystems to image from)	Full drive will be imaged/ cloned; encrypted state	Clears drive starting at sector/ block 0	n/a (no APFS support on destination)	Checks full drive starting at sector/ block 0
	<b>Unlocked</b>	Selected files/ folders will be imaged	Full drive will be imaged/ cloned; encrypted state	Clears drive starting at sector/ block 0	n/a (no APFS support on destination)	Checks full drive starting at sector/ block 0
<b>BitLocker</b>	<b>Locked</b>	n/a (no filesystems to image from)	Full drive will be imaged/ cloned; encrypted state	Clears drive starting at sector/ block 0	Will overwrite existing formatting, including BitLocker	Checks full drive starting at sector/ block 0
	<b>Unlocked</b>	Selected files/ folders will be imaged	Full drive will be imaged/ cloned; encrypted state	Clears drive starting at sector/ block 0	Will overwrite existing formatting, including BitLocker	Checks full drive starting at sector/ block 0
<b>Opal</b>	<b>Locked</b>	n/a (no reads possible)	n/a (no reads possible)	n/a (no writes possible)	n/a (no writes possible)	n/a (no reads possible)

		Operation				
		Logical Image (Source)	Physical Image/ Clone (Source)	Wipe (Dest)	Format (Dest)	Blank check (Source/ Dest)
	<b>Unlocked</b>	Selected files/ folders will be imaged	Full drive will be imaged/ cloned; unencrypted state	Clears drive starting at sector/ block 0	Formats drive starting at sector/ block 0	Checks full drive starting at sector/ block 0
Tableau	<b>Locked</b>	n/a (no filesystems to image from)	Full drive will be imaged/ cloned; encrypted state	Clears drive starting at sector/ block 0	Not allowed	Not allowed
	<b>Unlocked</b>	Selected files/ folders will be imaged	Only the unlocked encryption container contents will be imaged/ cloned; unencrypted state	Clears only the unlocked encryption container leaving encryption intact	Formats unlocked encryption container only leaving encryption intact	Checks contents of unlocked encryption container only

### 3.3.5 Removing 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. The ACS-3 (ATA/ATAPI Command Set 3) specification update introduced the concept of Addressable Maximum Address (AMA). Newer drives may support this method of limiting the reported drive capacity.

OpenText TX2 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:** Removing drive capacity limiting configurations applies to both source and destination drives. For source drives it is a stand-alone media utility, but for destination drives it is part of the **Reconfigure** utility. For details regarding HPA/DCO/AMA removal for destination drives, see “[Reconfigure](#)” on page 37.

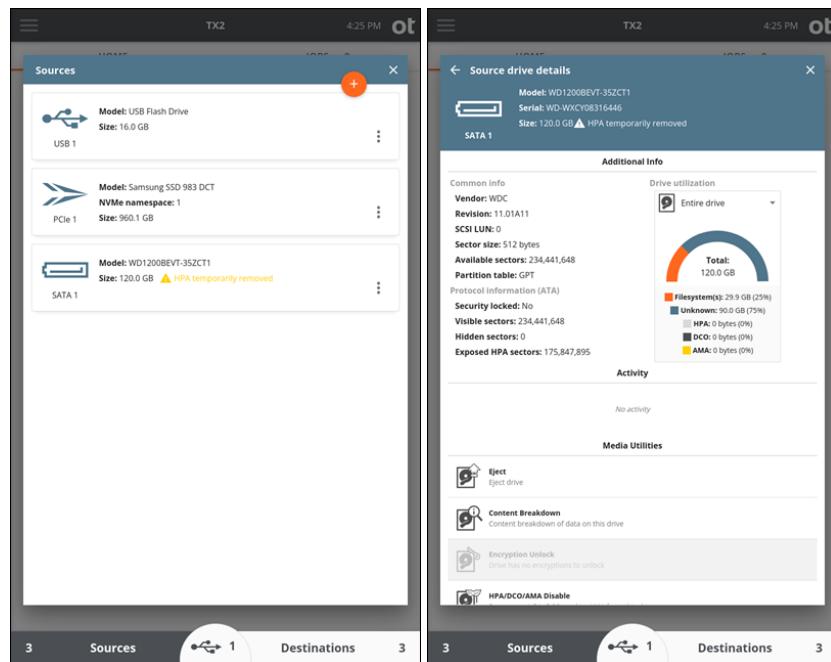
These methods (HPA/DCO and AMA) are mutually exclusive. A drive that supports HPA/DCO does not support AMA, and a drive that supports AMA does not support HPA/DCO. Also, while HPA and DCO are related features for a given drive, HPA has a unique attribute (volatile, or temporary, removal) that distinguishes it from DCO and AMA. For that reason, this section covers the volatile HPA removal as a

separate topic before addressing the non-volatile (permanent) removal of HPA/DCO or AMA.

### 3.3.5.1 Volatile HPA removal

An HPA can be removed 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 OpenText TX2 (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), your device automatically removes HPA on any drive connected to one of its source ports. Since DCO and AMA settings can only be removed on a permanent basis, your device does not automatically remove them on connected source drives.

In the case of an automatic, volatile HPA removal from a connected source drive, a message stating “HPA temporarily removed” will be shown in any affected drive tiles (**Sources** drive list) and in the header section of the associated drive details screen, as shown in the following images.



Referring to the drive details image, 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 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.

OpenText TX2 never makes automatic changes to any drive capacity-limiting configurations on destination drives. It 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, your device will not override that decision. For details regarding HPA/DCO/AMA removal for destination drives, see “[Reconfigure](#)” on page 37.

### 3.3.5.2 Non-volatile HPA/DCO/AMA removal

The **HPA/DCO/AMA Disable** media utility permanently removes the HPA, DCO, or AMA configurations on the source drive. 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 DCO or AMA configured, a red warning message is displayed on the drive tile indicating DCO or AMA is limiting the drive size. Permanently removing a DCO (and any HPA on that drive) or AMA is done using the **HPA/DCO/AMA Disable** utility found in the Media Utilities section of the source drive details screen. Drive details can be viewed through the **Sources** and **Destinations** areas of the main **Home** screen or from the **Select a Source** or **Select a Destination** areas during duplication job setup.



**Note:** Removing drive capacity limiting configurations applies to both source and destination drives. For source drives, it is a stand-alone media utility (**HPA/DCO/AMA Disable**), but for destination drives it is part of the **Reconfigure** utility. For details regarding HPA/DCO/AMA removal for destination drives, see “[Reconfigure](#)” on page 37.

OpenText TX2 also provides the ability to “shelve” a DCO or AMA, which will remove a source drive DCO or AMA for the purposes of evidence acquisition and then put the same DCO or AMA back after the job is complete. See “[Duplicating](#)” on page 58 for more details on shelving a DCO.

### 3.3.6 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
<b>Fast</b>	Quickly checks to determine if the drive appears to be blank, by reading in and checking the sectors in the <i>Master Boot Record</i> , the <i>Primary GPT</i> , and the <i>Secondary GPT</i> .
<b>Smart</b>	Performs the <b>Fast</b> 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.

Option	Description
<b>Complete</b>	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.

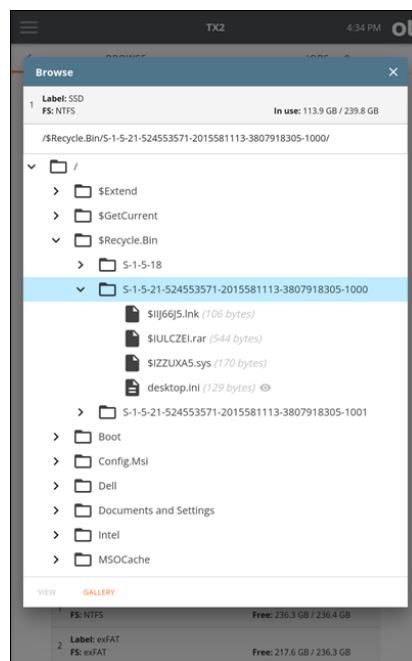
A sector is considered blank if it contains only the same repeated 2-byte pattern. Any non-repeating pattern is considered to be non-blank. However, each individual sector may contain different repeating patterns. If any sector is found to not be blank, the drive is not considered blank, and the blank 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.

### 3.3.7 Browse Filesystem

The **Browse** function provides an easy way to view the contents of a recognized filesystem on any attached drive, whether it is connected locally or via the network interface (iSCSI or CIFS). 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.

The following image shows a sample **Browse** screen.



The first row of the **Browse** window provides basic information about the selected filesystem (label, type, and used space). 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.

The main browse window (under the path name), shows the complete filesystem tree, including all folders/files contained on the drive or share.

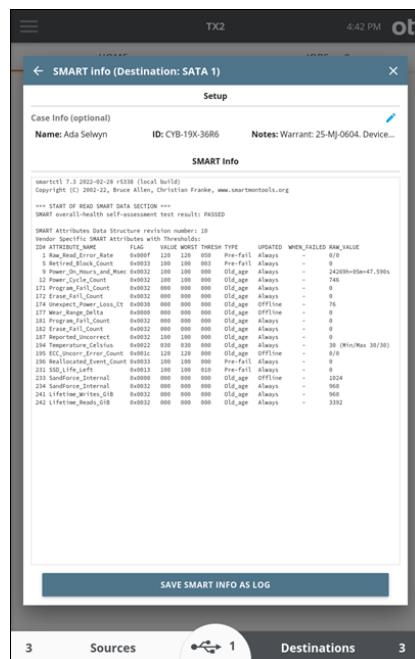
In the browser portion of the window, you can scroll up and down the list of folders/files and tap individual folders to drill down to the desired level, to expose the names of individual files located on the drive. The size of each file is shown at the end of the filename. Many users will find this 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 OpenText TX2. For more information about this feature, see “[Viewing text and image files](#)” on page 102.

### 3.3.8 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, as shown in the following image.



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

### 3.3.9 Export

This media utility allows the user to export any source, accessory, or destination drive as an iSCSI target in read-only mode. This makes the drive available to be read by a remote user on any IP-based network via the Ethernet connection on the rear of your device, which can be useful for evidence file transfer purposes.

During the export process, your device 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 OpenText TX2, by exporting 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
<code>iqn.2015-10.com.guid:</code>	Base identification number/domain for all OpenText Forensic Equipment products (as provided by the iSCSI naming authority).
<code>sn-000ecc5801000c.</code>	The serial number of the specific OpenText TX2 in use, which serves as a unique identifier within the domain.
<code>sata.1</code>	OpenText TX2 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 <b>Suffix</b> entry field on the <b>iSCSI Export</b> setup screen.

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 the **Remove Export** button in the **iSCSI Export** screen.

## 3.4 Connecting drives

This section provides instructions for safely connecting drives to an OpenText TX2.

TX2 operates as a standalone device or with the TX2-S1 drive bay.

### To connect the drive bay to TX2:

1. Ensure TX2 is powered off.
2. Align TX2 into place on top of the drive bay, and slide it back to lock it into place.

### 3.4.1 Source drives

Users can connect one or more drives to the OpenText TX2 source (left), write-blocked side interfaces: USB 3.2 Gen 2 Type C (x2), PCIe Gen 3x4 (x2), SATA Gen 3 (x2).



**Note:** While all OpenText TX2 device ports support hot swapping, it is highly recommended that all drives be ejected from the system before removing them. This is especially true of PCIe drives, which have a rather nuanced relationship with system level software.

An OpenText Forensic PCIe adapter (sold separately) is required to acquire PCIe drives. The OpenText Forensic IDE-PCIe adapter (TA7-5, sold separately) is required to acquire IDE drives. The OpenText Forensic FireWire-PCIe adapter (TA7-9, sold separately) is required to acquire FireWire media. For more information about adapters available for OpenText TX2, see ["Adapters" on page 135](#).

To acquire SAS drives, an OpenText T6u Forensic SAS bridge (sold separately) may be used on any source USB port. Alternatively, there are many high quality and reliable commercially available SAS-USB adapters that can be used.

OpenText TX2 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 Type C, FireWire, and Thunderbolt. Commercially available adapters are required to convert these interfaces to USB for connecting to OpenText TX2. For more information about the required adapters, see ["Adapters" on page 135](#) or contact OpenText Customer Support.

OpenText TX2 can acquire backup files from Apple and Android mobile devices. This done by directly connecting the mobile device to one of the OpenText TX2 USB Type C ports. For more information, see ["Mobile backup acquisition" on page 104](#).

OpenText TX2 also provides two 1/10 Gbps RJ-45 Ethernet connections, which allow the use of network-based acquisition of iSCSI physical media (clones and physical or logical images) and mounted CIFS shares (logical images). For more information, see ["Duplication over a network" on page 73](#).

Source drives are listed in the user interface in the order of the OpenText TX2 physical port layout: iSCSI/CIFS, USB 1, USB 2, PCIe 1, PCIe 2, SATA 1, SATA 2.

### 3.4.2 Destination drives

Users can connect one or more drives to the OpenText TX2 destination (right) side interfaces: USB 3.2 Gen 2 Type C (x2), PCIe Gen 3x4 (x2), SATA Gen 3 (x2).

Two additional SATA Gen 3 destination ports are available through use of the OpenText Forensic TX2-S1 Drive Bay (sold separately). This is a drive docking bay that attaches to the bottom of the imager, allowing for cable-free attachment of up to two 2.5" or 3.5" SATA destination drives. TX2-S1 also includes built-in fans to help keep drives cool during acquisition jobs.

OpenText TX2 also provides two 1/10 Gbps RJ-45 Ethernet connections, which allow the use of network-based iSCSI physical media and CIFS shares as destinations, to store clone or image file data. For more information, see ["Duplication over a network" on page 73](#).

Destination drives are listed in the user interface in the order of the OpenText TX2 physical port layout: iSCSI/CIFS, USB 1, USB 2, PCIe 1, PCIe 2, SATA 1, SATA 2.

### 3.4.3 Accessory drives

Users can connect up to two USB drives to the Accessory USB 3.2 Gen 1 Type C ports on the front of OpenText TX2. These drive interfaces are mostly used for saving stored log files, loading HTTPS/802.1x certificates, or updating the firmware. They may also be used for other purposes, including attachment of a physical (wired or wireless) keyboard and/or mouse.



#### Caution

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

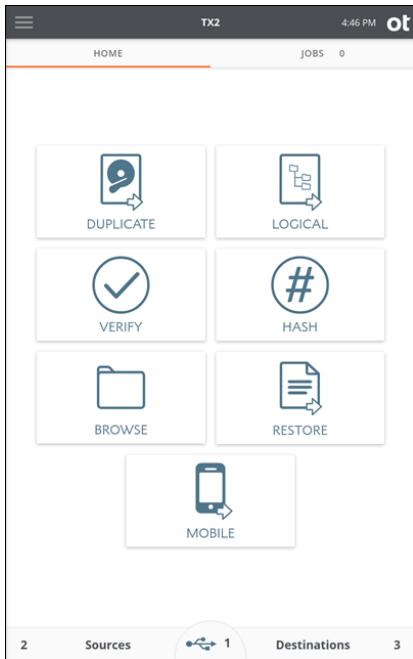
Accessory drives are listed in the user interface in the order of the OpenText TX2 physical port layout, with the leftmost USB drive always on top, and the rightmost USB drive on the bottom.

### 3.4.4 Drive detection

After booting, OpenText TX2 begins powering on and detecting connected drives, sequentially. Source and destination drive counts are displayed in the **Sources** and **Destinations** buttons at the bottom of the **HOME** screen, indicating the number of detected source and destination drives. 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.

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

The following screenshot shows the **HOME** screen and highlights the source, USB accessory, and destination drive counts.



OpenText TX2 can detect PCIe NVMe drives that have multiple namespaces defined. NVMe namespaces are similar to SCSI Logical Units (LUNS) in that there can be multiple of them configured on a single physical drive and each appears as a distinct drive in the imager's user interface. Each namespace will have a distinct drive tile (as shown in the **Sources** and **Destinations** drive lists) and will be treated as an individual drive by the imager. It is important to note that namespaces can be "detached" from the controller on some NVMe drives. A detached namespace would normally not be exposed to a user when its parent drive is connected to a typical host computer, but OpenText TX2 will detect the presence of any such detached namespaces and inform the user of their presence, by way of a dedicated drive tile in the **Sources** and **Destinations** drive lists. When you see a drive tile with the *Unattached NVMe namespace* warning message, tap the tile and a media utility modal will appear, allowing you to attach the namespace. Enter case information (optional) and tap the **ATTACH NAMESPACE(S)** button to begin the attachment process. If successful, the namespace will show as a regular drive tile and that namespace will be usable for any forensic activity (browsing, logical imaging, physical imaging, etc.).

Your device 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. Your device cannot run these proprietary applications, as they are typically made for x86-based Windows systems, therefore OpenText TX2 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, OpenText TX2 will detect these drives and report their type.



**Note:** Mobile devices connected to OpenText TX2 have unique detection interactions compared to traditional media devices (HDDs, SSDs). Once connected and detected, they will be displayed in the **Sources** list, but the similarities with traditional media end there. For information specific to that type of job, see [“Mobile backup acquisition” on page 104](#).

## 3.5 Turning off your unit

### To turn off your unit:

1. Push the power button in the top left corner of the unit.  
The shutdown options will be different for an idle unit versus one that has active/queued jobs.
2. For an idle unit, confirm the request by tapping the **SHUTDOWN** button, or keep the unit powered up by tapping the **CANCEL** button.
3. For a unit with active/queued jobs, the choices are expanded to include an option to wait for jobs to complete before the unit is turned off. This feature is convenient for running a job overnight or over a weekend with the unit unattended, as it will help reduce power consumption and unnecessary runtime on any attached drives.

To turn off your unit when the current job(s) is (are) complete, push the power button in the top left corner of the unit, and then tap the **WAIT FOR JOBS** button.

After all active/queued jobs complete, the unit will power itself off. This will work for any job type.



**Note:** If this method of powering down is used, there is no need to eject any attached drives before shutting down the unit. Using this proper shutdown method allows the software time to complete any active tasks and eject drives prior to turning the unit off. Forcing the unit to power off by pulling the power cord or by holding down the power button is not recommended, as it may corrupt any existing partition/filesystem information.



## Chapter 4

# Using the OpenText Forensic TX2 Imager

This chapter provides detailed procedures and information for using OpenText TX2.

## 4.1 Navigating OpenText TX2 features and options

The OpenText TX2 user interface includes the following elements:

- **Home** screen: Provides access to the following functions:
  - **Duplicate**
  - **Logical**
  - **Verify**
  - **Hash**
  - **Browse**
  - **Restore**
  - **Mobile**
  - **Sources, Accessory Drives, and Destinations**
    - **View connected drive detail**
    - **Access media utilities**
- **JOBS** screen: Provides the Job summary list and job details/status.
- **Side navigation menu**: Provides access to the following functions:
  - **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, OpenText TX2 automatically checks for preconditions. Some preconditions produce warnings, and you can choose to continue or cancel after viewing each one. Some preconditions are gating; they require mitigation or that the duplication process be stopped.

## 4.3 Duplicating

For each active job, OpenText TX2 duplicates one source drive to up to four destination drives simultaneously. The destinations can be any combination of clone and/or image jobs, and a mix of network shares and/or locally attached drives may be used.

There is no predefined limit on the number of jobs that can be active at the same time. However, the system automatically monitors available processor resources as job requests are added and determines whether to start them or queue them. This helps to ensure maximum efficiency for all the requested jobs, by minimizing job context switching while simultaneously ensuring that processor resources are being fully utilized.

While it is recommended to let OpenText TX2 determine when to start jobs, this system may be manually overridden. To start a job that the system has decided to enqueue, tap and hold the **drag** icon near the queued job tile and drag it into the **Active Jobs** section. If reordering jobs is the goal, it is recommended to pause the lower priority active job(s) before manually activating a higher priority job.



**Note:** This section focuses on whole disk duplication operations. See “[Logical imaging](#)” on page 82 for details on that alternative acquisition method. See “[Mobile backup acquisition](#)” on page 104 for information specific to that type of job.

### 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, a yellow warning is displayed, 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.

At the beginning of a clone job, your device 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. The device issues a warning when a sector size mismatch is detected. This condition will need to be rectified before the clone job can be started.

### 4.3.2 Physical imaging

A physical image, also known as disk-to-file duplication, copies the entire source drive to a series of files (sometimes called segments) on the destination drive(s). OpenText TX2 supports *EnCase* file formats **Ex01** and **E01** and raw file formats **DD** and **DMG**. Compression is supported and turned on by default with **Ex01** and **E01** file formats. File sizes from 4 GB per segment to Unlimited are supported. Smaller segments create more directory entries and Unlimited creates one large file segment.



**Note:** Due to filesystem addressing limitations, FAT32 formatted destinations have a maximum file size of 2 GB. If such a destination is detected by the imager, the output file size is automatically set to 2 GB with no option to change it.

When imaging, the destination media must first be formatted with a recognized filesystem. Format destination drives by selecting the **Reconfigure** media utility in the drive details screen. For more details, see “[Reconfigure](#)” on page 37. The drive details screen can be accessed through the **Destinations** button on the **Home** screen or through the **Select Destinations** screen during the 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 in cases where the data is mostly incompressible such as encrypted data.



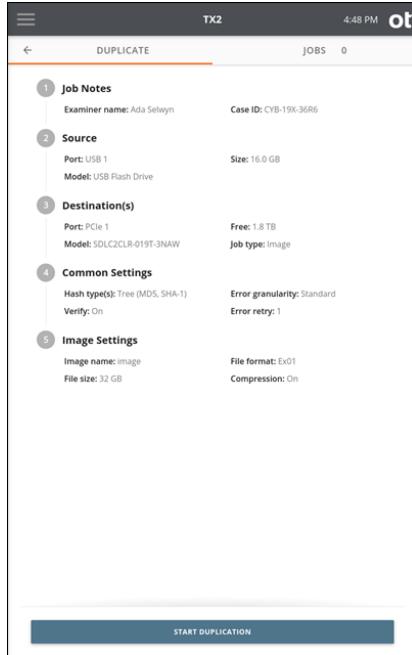
**Note:** Use caution when attempting to copy a source drive to a same size or smaller destination drive. Image file formatting adds overhead and, when coupled with incompressible data (such as encrypted data), a larger destination drive may be needed.

### 4.3.3 Performing a duplication

#### To perform a duplication:

1. Follow the steps listed in “[Connecting drives](#)” on page 51 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 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 **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 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 the modal.

If a different source is desired, go back into the **Select a source** screen by tapping on the **2** or **Source** heading from the **Duplication** stepper.



**Note:** 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 drive that has 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 on the drive tile, to see more drive detail and access any available media utilities.

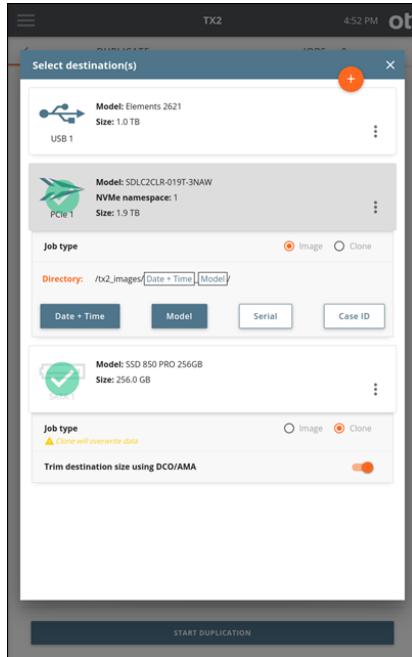
OpenText TX2 also allows shelving of a DCO or AMA for source drives. When enabled for a given source drive, the **Shelve DCO/AMA** feature will, after the duplication job is started, remove the DCO/AMA, complete the acquisition of the entire source drive, and then attempt to re-apply the original DCO/AMA 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. Note that this **Shelve DCO/AMA** feature is only available for SATA drives that support DCO/AMA. For more details regarding DCO/AMA support, see [“Removing drive capacity limiting configurations” on page 46](#).



**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 under the Drive tile.

Select the **Clone** button to clone to the destination or select the **Image** 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 appear dimmed 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 on 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. For more information, see “[Duplication over a network](#)” on page 73.

6. Review all messages displayed in the **Source** and **Destination(s)** headings.

For imaging jobs (not clones), if you selected only one destination drive which is smaller than the source, the *Image may not fit on selected destination(s)* message is displayed in the **Source** heading, and the *Spanning may be necessary* message is displayed in the **Destination(s)** heading. In this case, if the job ends up failing due to running out of destination space, it can be resumed by connecting a new destination drive. This is known as destination drive spanning.

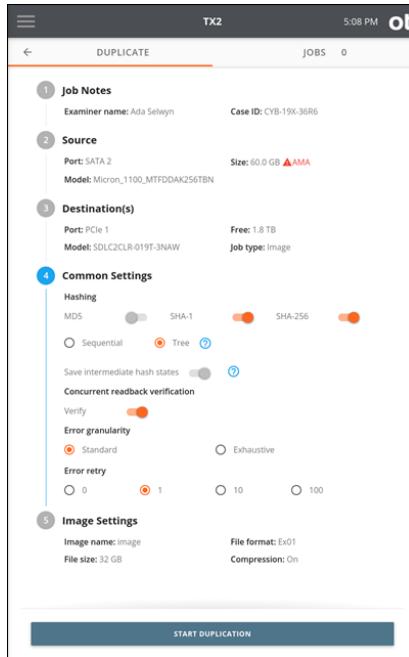


**Note:** To take advantage of this destination drive spanning feature, only one destination may be used in the original job configuration and the image file size setting (see step 8) must not be set to **Unlimited**.

If you selected multiple destination drives or if your image file size setting is set to **Unlimited**, the *Spanning will not be possible* message is displayed in the **Destination(s)** heading. To make such a setup spannable, return to step [step 5](#) and select only one destination drive for this job and/or change the image file size settings to any option other than **Unlimited**.

For more information, see “[Spanning a destination drive](#)” on page 78.

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



- a. **Hashing:** Select up to two hash types (of MD5, SHA-1, and SHA-256), then select one of the following hash methods:

- **Sequential Hashing:** This is the traditional hashing method that has been used in the digital forensics industry for many years. It hashes each byte of data in a sequential manner which makes this process the bottleneck in most forensic data acquisition jobs.
- **Tree Hashing:** This new hashing method breaks data into fixed size chunks and then calculates the selected hash(es) (MD5, SHA-1, and/or SHA-256) of each of those chunks. These intermediate chunk hashes can be calculated independently which allows for massive parallelization of what used to be a slow, sequential process. Then those stored intermediate hash values are sequentially hashed with the same selected algorithm(s) (MD5, SHA-1, and/or SHA-256) to generate the tree hash value for the entire source data set. OpenText TX2 keeps track of the order of each of the chunk hash values, which ensures reliable and predictable final tree hash values that will be repeatable and verifiable with the OpenText TX2, OpenText Forensic (formerly EnCase Forensic), and any forensic tool that supports this hashing method.

! **Important**

The hash value calculated during a tree hash job and that of a sequential hash job for the same exact source data set will be different. The OpenText Forensic TX2 Imager's tree hash implementation (patent pending) is based on the publicly available Sakura tree hashing framework ([https://link.springer.com/chapter/10.1007/978-3-319-07536-5\\_14#preview](https://link.springer.com/chapter/10.1007/978-3-319-07536-5_14#preview)). OpenText's specific implementation of this framework, as defined in the *OpenText Tree Hashing Specification*, allows the implementation of this tree hashing method by other forensic tool providers, subject to any applicable license agreements. OpenText Forensic TX2 Imager can verify any tree-hashed forensic file set (E01, Ex01, DD, and DMG) or clone it creates. Additionally, OpenText Forensic (formerly EnCase Forensic) version 25.2 (or later) is able to verify tree-hashed E01 and Ex01 file sets created by an OpenText Forensic TX2 Imager.

OpenText is actively working to proliferate this technology by making it easy for other tool vendors to implement it in their digital forensics tools.

To see more information about *OpenText Tree Hashing* and to view the specification document, tap the help button  available for the Tree Hash setting in either the Common Settings area of a Duplicate job setup screen or in the Job settings area of the Defaults settings screen. If you are using the remote user interface, there is an additional option to download the specification document from the help message window. For more information about *OpenText Tree Hashing*, or to discuss collaboration with other tool vendors, contact your OpenText Forensic Equipment distributor or an OpenText Forensic salesperson.



**Note:** The E01 file format does not support SHA-256 hash values (sequential or tree), and thus OpenText Forensic TX2 Imager will not

allow that selection. For all other image file output types (Ex01, DD, DMG) and for clone jobs, any two hashes can be selected from the three supported algorithms (MD5, SHA-1, SHA-256).

- b. **Concurrent readback verification** is turned on as a system default, which will automatically initiate readback verification as soon as possible after evidence data is written to the destination. In the case of a clone job, readback verification will begin shortly after the first write is made to the destination drive(s). For all other image output types (E01, Ex01, DD, and DMG) verification will begin as soon as the first segment file is complete.



**Note:** To make the most of concurrent imaging/verification, it is recommended to select the smallest segment file size of 4 GB. This will allow OpenText TX2 to begin readback verification as early as possible.

- c. **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 32 KB 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.

- d. **Error retry** defines the number of times OpenText TX2 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.

8. 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**.



### Notes

- Regarding the file format setting, 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. A warning is displayed in the **Image Settings** area when this condition is detected and **DMG** is selected as the file format.

- Regarding the files size setting, OpenText TX2 will begin readback verification (if turned on during the job setup) as soon as the first segment file is complete. To make the most of this concurrent imaging/verification feature, it is recommended to select the smallest segment file size of 4 GB. This will allow OpenText TX2 to begin readback verification as early as possible.
- **Compression** is by default turned on for **Ex01** and **E01** file formats, but it can be turned off, if desired. Compression is not supported for **DD** and **DMG** file formats.

9. Once you are satisfied with your settings and drive selection, tap the **Start Duplication** button.

A **Job Status** screen is automatically displayed.

#### 4.3.3.1 Files created during disk-to-file duplication

When performing an image, OpenText TX2 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 `/tx2_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**. For more information, see “[Performing a duplication](#)” on page 60.

[filename] is the base image filename and is defined in “[Physical imaging](#)” on page 59 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).

OpenText TX2 generates a .LOG file for each image job. It also creates a `.tx2_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 the Automated Acquisition mode

OpenText TX2 can be set up to run in **Automated Acquisition** mode, which will start a duplication job (physical image type) when any source drive is detected by the system. Custom job parameters can be set to the user's liking when turning on 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 OpenText TX2 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 following steps are focused on Automated Acquisition mode setup. For basic duplication job setup information, see ["Performing a duplication" on page 60](#).



**Note:** Automated Acquisition is available for duplication (physical image type) 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. Automated Acquisition jobs do not support clone type duplication jobs because subsequent automated clone jobs would overwrite the destination, effectively deleting the previous automated clone acquisition.

##### To set up the 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 **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 to make data entry easier. All information entered in the **Job Notes** section will be used for each of the ensuing automated jobs.
3. To set the device 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:
  - **Acquire currently connected:** If this option is turned on, any source drives currently connected to and detected by the system will have jobs automatically started when the Automated Acquisition mode is fully turned on.

- **Acquire newly connected:** If this option is turned on, 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 removed prior to the start of the job, and then replaced after the job is complete. See [“Removing drive capacity limiting configurations” on page 46](#) for more information regarding the Shelve AMA/DCO feature.

4. Your selection of Automated Acquisition mode will be confirmed by the **Automated Acquisition** icon near 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 under the **Drive** tile.

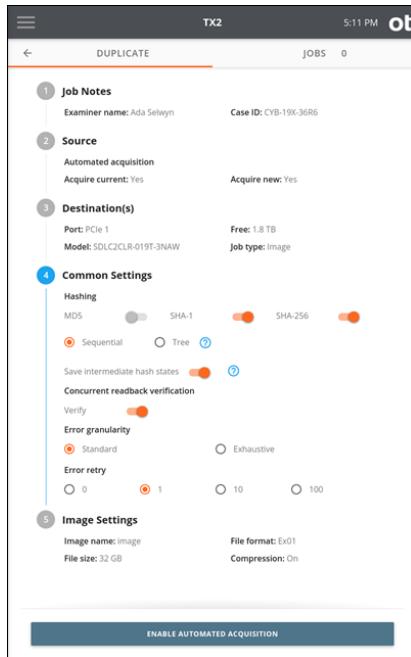
 **Note:** Make sure to select **Image** as the job type for each destination drive. While it is possible to set a **Clone** job type here, that is not an option for Automated Acquisition jobs.

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 appear dimmed 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 on 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.



a. **Hashing:** Select up to two hash types (of MD5, SHA-1, and SHA-256), then select one of the following hash methods:

- **Sequential Hashing:** This is the traditional hashing method that has been used in the digital forensics industry for many years. It hashes each byte of data in a sequential manner which makes this process the bottleneck in most forensic data acquisition jobs.
- **Tree Hashing:** This new hashing method breaks data into fixed size chunks and then calculates the selected hash(es) (MD5, SHA-1, and/or SHA-256) of each of those chunks. These intermediate chunk hashes can be calculated independently which allows for massive parallelization of what used to be a slow, sequential process. Then those stored intermediate hash values are sequentially hashed with the same selected algorithm(s) (MD5, SHA-1, and/or SHA-256) to generate the tree hash value for the entire source data set. OpenText TX2 keeps track of the order of each of the chunk hash values, which ensures reliable and predictable final tree hash values that will be repeatable and verifiable with the OpenText TX2, OpenText Forensic, and any forensic tool that supports this hashing method.

**! Important**

The hash value calculated during a tree hash job and that of a sequential hash job for the same exact source data set will be different. The OpenText Forensic TX2 Imager's tree hash implementation (patent pending) is based on the publicly available Sakura tree hashing framework ([https://link.springer.com/chapter/10.1007/978-3-319-07536-5\\_14#preview](https://link.springer.com/chapter/10.1007/978-3-319-07536-5_14#preview)). OpenText's specific

implementation of this framework, as defined in the *OpenText Tree Hashing Specification*, allows the implementation of this tree hashing method by other forensic tool providers, subject to any applicable license agreements. OpenText Forensic TX2 Imager can verify any tree-hashed forensic file set (E01, Ex01, DD, and DMG) or clone it creates. Additionally, OpenText Forensic (formerly EnCase Forensic) version 25.2 (or later) is able to verify tree-hashed E01 and Ex01 file sets created by an OpenText Forensic TX2 Imager.

OpenText is actively working to proliferate this technology by making it easy for other tool vendors to implement it in their digital forensics tools.

To see more information about *OpenText Tree Hashing* and to view the specification document, tap the help button  available for the Tree Hash setting in either the Common Settings area of a Duplicate job setup screen or in the Job settings area of the Defaults settings screen. If you are using the remote user interface, there is an additional option to download the specification document from the help message window. For more information about *OpenText Tree Hashing*, or to discuss collaboration with other tool vendors, contact your OpenText Forensic Equipment distributor or an OpenText Forensic salesperson.



**Note:** The E01 file format does not support SHA-256 hash values (sequential or tree), and thus OpenText Forensic TX2 Imager will not allow that selection. For all other image file output types (Ex01, DD, DMG) and for clone jobs, any two hashes can be selected from the three supported algorithms (MD5, SHA-1, SHA-256).

- b. **Concurrent readback verification** is turned on as a system default, which will automatically initiate readback verification as soon as possible after evidence data is written to the destination. In the case of a clone job, readback verification will begin shortly after the first write is made to the destination drive(s). For all other image output types (E01, Ex01, DD, and DMG) verification will begin as soon as the first segment file is complete.



**Note:** To make the most of concurrent imaging/verification, it is recommended to select the smallest segment file size of 4 GB. This will allow OpenText TX2 to begin readback verification as early as possible.

- c. **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 32 KB 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.

- d. **Error retry** defines the number of times OpenText TX2 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. 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**.



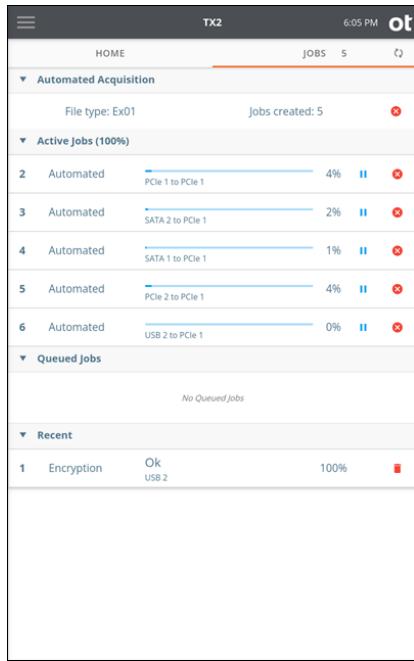
### Notes

- Regarding the file format setting, **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. A warning is displayed in the **Image Settings** area when this condition is detected and **DMG** is selected as the file format.
- Regarding the files size setting, OpenText TX2 will begin readback verification (if turned on during job setup) as soon as the first segment file is complete. To make the most of this concurrent imaging/verification feature, it is recommended to select the smallest segment file size of 4 GB. This will allow OpenText TX2 to begin readback verification as early as possible.

**Compression** is by default turned on for **Ex01** and **E01** file formats, but it can be turned off, if desired. Compression is not supported for **DD** and **DMG** file formats.

8. Once you are satisfied with your settings and drive selections, tap **ENABLE AUTOMATED ACQUISITION**.

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, near the job count indicator. 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, 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 turned off in step 3, 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)** near the checkmark on the drive tile that indicates a drive was acquired (or is being acquired, in the case of a hollow checkmark). This helps you keep track of which jobs were manually or automatically initiated.

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 **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

OpenText TX2 has two, independent, copper (Base-T) 10-Gigabit Ethernet (GbE) ports. These connections allow for a superior network imaging performance when combined with a properly configured 10 GbE network infrastructure setup and destination storage server, such as a Storage Area Network (SAN) or Network Attached Storage (NAS) or a server configured with an iSCSI target or CIFS share. Duplication/Logical Imaging from network-based iSCSI targets and CIFS shares is also supported. These Ethernet ports also support legacy 1 Gigabit Ethernet connections, and the connection speed will auto-negotiate based on network capabilities.

#### Notes

- OpenText TX2 must be connected to a network before attempting to mount and use iSCSI targets or CIFS shares.
- The two Ethernet ports on OpenText TX2 are independent without support for NIC Teaming (also known as Link Aggregation). A typical use case is to have one connection to a LAN and one to a directly-attached NAS device. For more information about network port configuration, see [“Network settings” on page 23](#).

#### 4.3.5.1 Adding an iSCSI target

##### To add an iSCSI target as a source or destination drive:

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 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:** In the user interface, passwords can be shown as either plain text or hidden. Tap the eye icon  near 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.
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** link under the iSCSI drive tile, enable or disable the desired **Username** and **Password** values to be saved, and then tap **SAVE AS BOOKMARK**.

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 regular 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, server hostname, or fully qualified domain name (FQDN) of an available server, then select **Next**.



**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 checkmark on the left. Tap the **Show Hidden Shares** slider to view default admin/hidden shares in the share list.
5. Select the desired share and the mount screen will appear. Enter a nickname for this CIFS share (optional) and enter a login username and password (if required).

Choose the **SMB Version** and turn on the SMB 3.0 encryption (if desired), then tap the **Mount** button to login and mount the CIFS share.



**Note:** Due to network security concerns, OpenText TX2 does not support SMB 1.0 as a mounting option for CIFS shares.

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** link under the CIFS drive tile, enable or disable the desired **Username** and **Password** values to be saved, and then tap **SAVE AS BOOKMARK**.

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), and some media utilities.

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 in this procedure 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. OpenText TX2 allows you to pause and resume imaging jobs with the following output file formats: E01, Ex01, DD, and DMG.

#### To pause a running duplication job:

1. Locate the desired job in the **Active Jobs** area of the **Jobs** tab.
2. Tap its **Pause** button , and then confirm the operation.

The job will be moved to the **Recent** area, with a status of **Paused**.

#### To resume a paused job:

- Choose one of the following options.
  - Tap the **Play** button  on the paused job in the **Recent** area of the **Jobs** tab.
  - Tap the **Resume Job** button in the header of the **Job Status** screen.



**Tip:** The **Job Status** screen can be viewed by tapping on the job in the **Jobs** tab.

- Tap 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 appear.

This screen checks for the availability of the original job's source and destination drives before allowing resumption of the paused job. If Verification was not turned on in the original imaging job setup, the **Resume Duplication** screen offers a means to turn it on before resuming the job.

In addition to manually initiated pause and resume, OpenText TX2 also supports power loss situations. For the supported job types (E01, Ex01, DD, and 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 before the original job can be resumed after a power loss event. This procedure assumes the original drives are connected and available to the system after the power was restored.

1. On the **Home** screen, tap the menu icon  to open the side navigation menu.
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 that this is the job that was running when power was lost, and 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.

5. When you are comfortable that this is the log for the job you want to resume, tap **Resume Job** on the **Log Details** screen.

The **Resume Duplication** screen for the paused job will appear.

If the original drives are present and available to OpenText TX2, they will be shown with green checkmarks and the **Resume Duplication** button will be lit up.

If your original job did not have Verification turned on, you can turn it on in this screen (if desired).

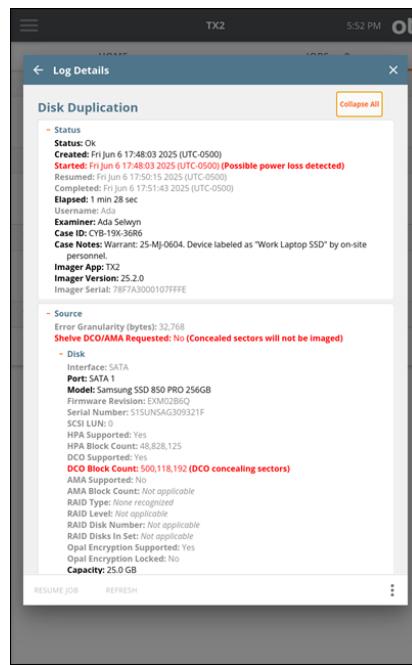
6. Tap **Resume Duplication** to resume the job.

The **Job Status** screen for the resumed job will appear. 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.



Users can resume 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 previous sections. 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.



**Note:** For each of these situations, the exact same source and destination drives are required before the job can be resumed. In the case of a full destination drive, other files (unrelated to the in-progress job) must be deleted from the

drive before the job can be resumed. Alternatively, in the case of a failed job due to a full destination, a second destination drive can be added to resume the job; for more information, see ["Spanning a destination drive" on page 78](#).

#### 4.3.7 Spanning a destination drive

When a physical imaging job (non clone) fails due to a full destination drive, you can add a new destination drive and resume the duplication process. This functionality allows you to image a larger source drive to multiple destination drives, preventing the need to start over such failed jobs from the beginning with a single, larger destination drive. An unlimited number of drives may be used for spanning.

To speed up the imaging process, OpenText recommends leaving the original destination drive attached when connecting the second destination drive and continuing the imaging process. This allows OpenText TX2 to copy any partial image segment files to the second destination drive and then resume the imaging process from where it failed. In this case, any partially processed segment files are removed from the first destination drive; both destination drives will contain only complete segment files at job completion.



**Note:** For simplicity and clarity, this section refers to a first/original drive and a second drive, but destination drive spanning can be done to any number of destination drives.

When disconnecting the first destination drive prior to connecting the second one, OpenText TX2 resumes the imaging process from the end of the last complete image segment file. This method requires more time to complete the overall job, because any partial segment files from the first drive will be discarded and have to be recreated on the second drive. In this scenario, the first drive will retain all segment files (complete and partial) that were created prior to the job failure, and the second drive will contain only the files processed after resuming the job.

**To span a destination drive:**

1. After a **Duplicate (image type)** job fails due to the destination filesystem being full, make sure that at least one other formatted destination drive is connected.



**Note:** To take advantage of this destination drive spanning feature, only one destination may be used in the original job configuration, and the image file size setting must not be set to **Unlimited**.

2. Tap the **Resume Job** button on the failed job's status screen.

The job may also be resumed via the log list, by tapping the **Resume Job** button on the failed job log.

3. In the **Resume Duplication** screen, review the **Destination(s)** area. Tap the **Span** button on the right side of the original destination drive tile.

The **Select a destination** screen appears, displaying the list of connected destination drives.

4. Select a new destination drive to continue the job with, by tapping its drive tile.

The **Resume Duplication** screen appears, showing the original and the new destination drives. A green checkmark indicates the newly selected destination drive.



**Note:** If readback verification was not enabled for the original job, you can enable it here but only if all associated destination drives are still connected.

5. Tap **Resume Duplication** to resume the job.

The **Job Status** screen for the resumed job appears. From this point on, the job will carry on as if it had never been paused, with the exception of logged pause/resume events.

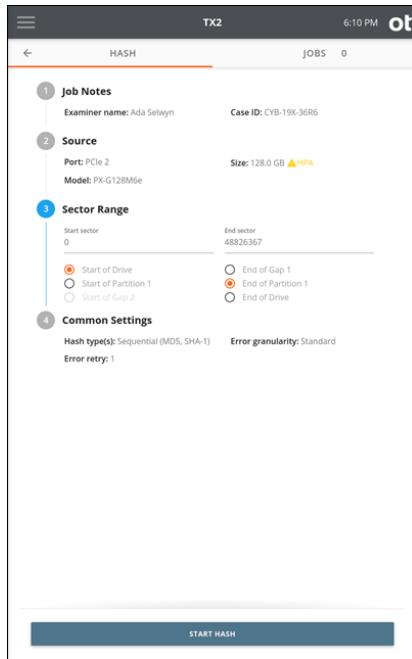
When a source drive is imaged to multiple destination drives, the job log clearly lists every destination drive used during the job. A forensic examiner will need to consolidate the files from all the destination drives into a single folder prior to importing them into a forensic analysis tool (like OpenText Forensic). Forensic analysis tools require all segment files to be placed together in a single folder, ensuring a complete and accurate reconstruction of the source image.

## 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 in either the legacy sequential mode or using OpenText Tree Hashing. 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:
  - a. Use the option 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.
  - b. 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.



4. In the **Common Settings** section, select up to two hash types of MD5, SHA-1, and SHA-256. Then select the hash method:

- **Sequential Hashing:** This is the traditional hashing method that has been used in the digital forensics industry for many years. It hashes each byte of data in a sequential manner which makes this process the bottleneck in most forensic data acquisition jobs.
- **Tree Hashing:** This new hashing method breaks data into fixed size chunks and then calculates the selected hash(es) (MD5, SHA-1, and/or SHA-256) of each of those chunks. These intermediate chunk hashes can be calculated independently which allows for massive parallelization of what used to be a slow, sequential process. Then those stored intermediate hash values are sequentially hashed with the same selected algorithm(s) (MD5, SHA-1, and/or SHA-256) to generate the tree hash value for the entire source data set. The OpenText Forensic TX2 Imager keeps track of the order of each of the chunk hash values which ensures reliable and predictable final tree hash values that will be repeatable and verifiable with the OpenText Forensic TX2 Imager, OpenText Forensic, and any forensic tool that supports this hashing method.

**! Important**

The hash value calculated during a tree hash job and that of a sequential hash job for the same exact source data set will be different. The OpenText Forensic TX2 Imager's tree hash implementation (patent pending) is based on the publicly available Sakura tree hashing framework ([https://link.springer.com/chapter/10.1007/978-3-319-07536-5\\_14#preview](https://link.springer.com/chapter/10.1007/978-3-319-07536-5_14#preview)). OpenText's specific implementation of

this framework, as defined in the *OpenText Tree Hashing Specification*, allows the implementation of this tree hashing method by other forensic tool providers, subject to any applicable license agreements. OpenText Forensic TX2 Imager can verify any tree-hashed forensic file set (E01, Ex01, DD, and DMG) or clone it creates. Additionally, OpenText Forensic (formerly EnCase Forensic) version 25.2 (or later) is able to verify tree-hashed E01 and Ex01 file sets created by an OpenText Forensic TX2 Imager.

OpenText is actively working to proliferate this technology by making it easy for other tool vendors to implement it in their digital forensics tools.

To see more information about *OpenText Tree Hashing* and to view the specification document, tap the help button  available for the Tree Hash setting in either the Common Settings area of a Duplicate job setup screen or in the Job settings area of the Defaults settings screen. If you are using the remote user interface, there is an additional option to download the specification document from the help message window. For more information about *OpenText Tree Hashing*, or to discuss collaboration with other tool vendors, contact your OpenText Forensic Equipment distributor or an OpenText Forensic salesperson.



**Note:** The E01 file format does not support SHA-256 hash values (sequential or tree), and thus OpenText Forensic TX2 Imager will not allow that selection. For all other image file output types (Ex01, DD, DMG) and for clone jobs, any two hashes can be selected from the three supported algorithms (MD5, SHA-1, SHA-256).

5. To start the hash job, tap the **Start Hash** button at the bottom of the screen. The **Job Status** screen appears.
6. To cancel the hash job, 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, 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

OpenText TX2 provides a powerful logical imaging function that allows for file-based evidence acquisition from locally attached and network-based source filesystems. Logical imaging saves valuable time by focusing on specific files of interest rather than acquiring the entire physical drive. This logical imaging understands the structure of the recognized filesystem and acquires the desired source file data and/or metadata.

Logical imaging operations on OpenText TX2 can be targeted, limiting which files are acquired to only a subset of the source filesystem. This device 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 allows for rapid acquisition of source file data, providing 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 OpenText Forensic. 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, the device 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 deactivated and no source file filtering), 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 51 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 is displayed.

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/device, tap the **2** or **Source** heading.

A list of attached drives/devices appears, with a filesystem summary tile shown under each drive for any filesystems recognized by OpenText TX2.



**Note:** Devices such as mobile phones, tablets, and digital cameras that support Picture Transfer Protocol (PTP) or Media Transfer Protocol (MTP) are mountable by OpenText TX2 and will be shown as an option for logical imaging source selection. The rest of the logical imaging job setup process is exactly the same as a standard drive-based filesystem source.

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, as described in “[Duplication over a network](#)” on page 73. Source network shares can be mounted from the **Sources** button on the main screen as well.

Once the source drive/share list is set, 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), 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.

### **Notes**

- Within any screen displaying a list of drives, you can tap the options icon (three vertical dots) located on the drive tile, to see more drive details and access any available media utilities.
- 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. 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. The default setting is to acquire **All files and folders** and is initially selected.

Use the default setting if your job does not benefit from targeted down-selection of source files/folders and continue with step 7.

If no source file or folder filtering is required, 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 filtering of the source dataset is required, there are two different starting points, as follows:

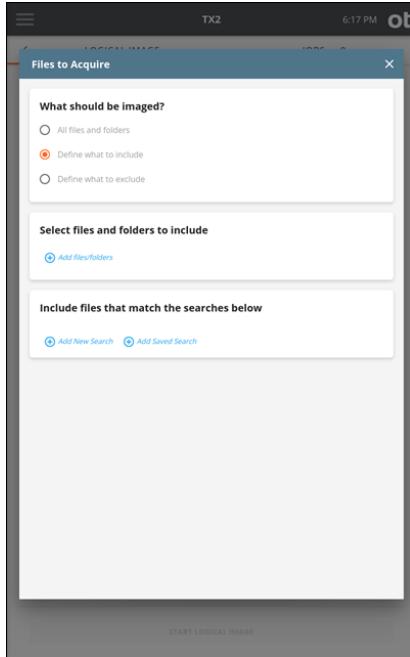
- **Define what to include:** This setting assumes that NO files and folders will be acquired, which requires defining which ones to INCLUDE.
- **Define what to exclude:** This setting assumes that ALL files and folders will be acquired, which requires defining which ones to EXCLUDE.

### **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. For more information about the optional Advanced mode search, see “[Advanced logical imaging setup](#)” [on page 97](#).

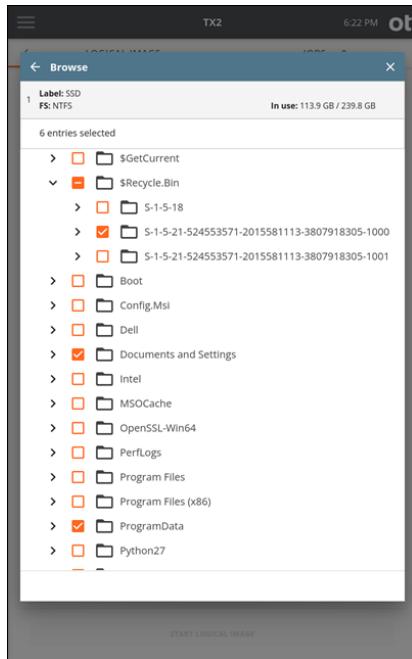
Regardless of whether you are including items in an empty dataset or excluding items from a full dataset, the same setup style is used to limit what is acquired.

6. Select files and folders to include/exclude.



In the example shown here, **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 launches a browse modal window which shows the entire contents of the selected source filesystem.



Individual files and/or folders can be manually selected by clicking the orange box preceding the desired item. In the preceding example, 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.

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 88](#) 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; each drive will have a corresponding tile showing 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 OpenText TX2 is Lx01, which will contain all the data and metadata for every acquired file and folder. Additionally, this device 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**. For more information regarding captured metadata, see “[Source file metadata](#)” on page 100.
- **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) on 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.



**Note:** Tree Hashing is not supported for logical imaging. The main benefit of tree hashing is to improve performance when hashing large amounts of block device data, which does not apply to file-based logical acquisition.

Turn on or turn off **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 allows forensic analysis tools such as OpenText Forensic 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 turned on/turned off. 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 **START LOGICAL IMAGE**.

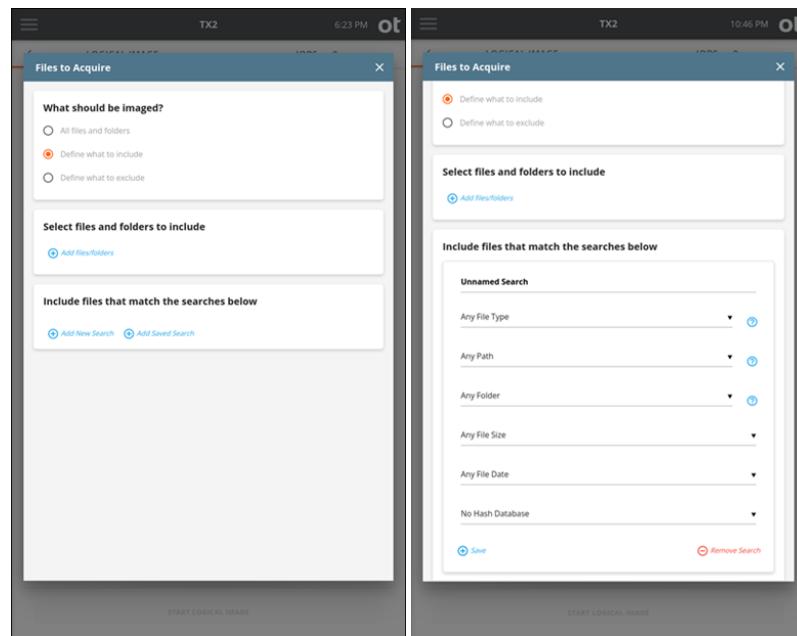
#### 4.5.2 Include/exclude criteria

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 in this section 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. 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, 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 shown here, 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 match 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 that are included in filesystems that this device supports.



**Note:** For characters that are composed of multiple, distinct codepoints (for example the German umlaut over the letter A), OpenText TX2 uses the NFC (Normalization Form Canonical) composition standard to normalize the various encoding methods into a common hex value for matching purposes. For more information, search for “Unicode normalization” on the internet.

The following sections provide detailed information about the various search parameters you can use when defining the include/exclude criteria for a local image acquisition.

#### 4.5.2.1 File type

The **File type** search parameter uses the file’s MIME type to locate files of the specified type(s). 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 MIME types known to be associated with that type of file. The lists can be seen by tapping the help button  on the file type parameter field. All MIME types associated with each file type category are found in “[File MIME types](#)” on page 98.

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

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

#### 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 following table shows the available wildcards and a few examples. 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).	<p><b>Law*</b> Matches: Law, Laws, Lawyer Does not match: NoLaw, La, aw</p> <p><b>*Law*</b> Matches: Law, NoLaw, Lawyer Does not match: La, aw</p>
?	Matches any single character.	<p><b>?it</b> Matches: Bit, bit, Sit, sit Does not match: it, slit, bits</p>
[abc]	Matches one instance of any of the characters between the brackets.	<p><b>[SB]it</b> Matches: Sit, Bit Does not match: sit, bit, it</p> <p> <b>Note:</b> 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]].</p>
[!abc]	Matches one instance of any of the characters that are not what is between the "!" and the "]".	<p><b>[!S]it</b> Matches: sit, Bit, bit Does not match: Sit</p>
-	Used within brackets, hyphens denote a range of characters to be matched.	<p><b>money[1-5]; equivalent to money[12345]</b> Matches: money1, money2, money5 Does not match: money, money6</p> <p> <b>Note:</b> To match a hyphen itself within brackets, it must be the first or last item inside the brackets.</p>

### 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 on 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 99.

The **Custom Folder** setting allows you to browse the source for a folder to match. After selecting **Custom Folder** from the drop-down list, 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 >=** 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 >=** 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 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. OpenText TX2 treats these timestamps as if created with a UTC+0000 timestamp.

#### 4.5.2.6 Hash database

When a hash database is applied as a logical image search, the acquisition process will either include or exclude files matching or not matching a set of known file hashes, depending on the desired behavior.

To use a file hash database to search for source drive files with specific hash values, tap **Add New Search** in the **Files to Acquire** screen, and then tap **No Hash Database** at the bottom of the search type list. The options are as follows:

- **No Hash Database** indicates that no hash database has been selected and thus this search method will not be used. This is the default setting.
- **In Hash Database** lets you select (from an attached accessory or destination drive) a hash database to be used in this search. All source files with hashes that match file hash values in the selected database will either be included or excluded in the logical acquisition, depending on your include/exclude selection at the top of the **Files to Acquire** screen.
- **Not in Hash Database** lets you select (from an attached accessory or destination drive) a hash database to be used in this search. All source files with hashes that do not match file hash values in the database will either be included or excluded in the logical acquisition, depending on your include/exclude selection at the top of the **Files to Acquire** screen.



**Note:** Hash database files must be in the native TX2 format. See “[Appendix A: Logical Imaging Engine hash databases](#)” on page 149 for detailed information about this hash database format and for a supporting Python script, that can create an OpenText TX2 hash database file from other file hash sources, including NSRL RDS, NSRL CAID, and several other custom options.

In order to minimize the file hash database search time and allow for unique identification of specific hash databases (for logging and internal job startup pre-condition check purposes), all externally created database files must be stored in a root folder named `otfe_hash_databases`. OpenText TX2 will automatically find any such stored database files and display in the database file selection screen.

Hash database files are not stored directly on OpenText TX2. A drive with the desired hash database file(s) will need to be connected to an accessory or destination port.

### 4.5.3 About the logical imaging process

This section provides information about the OpenText TX2 logical imaging process.

#### 4.5.3.1 Logical image job status

Once a logical image job has been started, the **Job Status** screen is automatically displayed.

This status screen is similar to other OpenText TX2 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). OpenText TX2 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, and 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 OpenText TX2, 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 in ["Source file metadata" on page 100](#).
- `[image_name].tx2_packed_log` contains a copy of the forensic log that your device can read and use for later standalone verification of the `Lx01` file set.

All these 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 `.tx2_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 OpenText TX2 can be switched to **Advanced Logical Imaging** mode. By turning on this setting in the **Default** settings screen, two additional search setup options are activated. If **Advanced Logical Imaging Setup** is turned on, but none of its additional features are used, OpenText TX2 behaves the same as in **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 near those two settings. If you turn on the invert switch near **Archives**, the rule now matches any file that is not an Archive. If you turn on the switch near the File Sizes range, the rule will now match any file outside the range 1 MB to 100 MB. 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 MIME types

Related to logical imaging, OpenText TX2 can search for file types based on the following MIME types shown in the following table. The names in the **File Type** column match predefined file type searches that can be selected during search setup. Custom MIME type values may also be used. For more information about using custom MIME types, see “[File type](#)” on page 90.

File Type	MIME Type
Archives	application/*gzip, application/java-archive, application/vnd.android.package-archive, application/vnd.rar, application/x-7z-compressed, application/x-archive, application/x-bzip*, application/x-freearc, application/x-rar-compressed, application/x-tar, application/x-xz, application/x-zip-compressed, application/x-zstd, application/zip, multipart/x-zip
Databases	application/*sqlite*, application/vnd.oasis.opendocument.database, application/x-berkeley-db, application/x-dbf, application/x-evtx, application/x-filemaker, application/x-msaccess

File Type	MIME Type
Documents	application/*document*, application/*spreadsheet*, application/epub+zip, application/msword, application/pdf, application/rtf, application/vnd.apple.keynote, application/vnd.apple.numbers, application/vnd.apple.pages, application/vnd.ms-excel*, application/vnd.ms-powerpoint*, application/vnd.ms-visio*, application/vnd.ms-word*, application/vnd.oasis.opendocument.base*, application/vnd.oasis.opendocument.graphics*, application/vnd.oasis.opendocument.presentation*, application/vnd.oasis.opendocument.text*, application/vnd.openxmlformats-officedocument.*, application/vnd.wordperfect, application/x-appleworks*, application/x-latex, application/x-mobipocket-ebook, application/x-ole-storage, application/x-tex, application/xml, text/*
Emails	application/mbox, application/vnd.ms-outlook, message/*
Multimedia	audio/*, video/*
Pictures	image/*

## 4.5.5 Folders

During logical imaging, OpenText TX2 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/
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 OpenText TX2 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 ISO 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 ISO 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 ISO 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 ISO 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.
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.

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. On 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 OpenText TX2 packed log file. Tap the desired file to select it and then close the browse modal.



**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.  
Once started, a **Job Status** modal is displayed, showing the verification job 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, which includes all calculated verification hashes. The log for the completed job can be viewed by tapping on the **View Log** link on 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). 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.



**Note:** Devices such as mobile phones, tablets, and digital cameras that support Picture Transfer Protocol (PTP) or Media Transfer Protocol (MTP) are mountable by OpenText TX2 and will be shown as options in the list of source drives/filesystems to select for browsing.

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

OpenText TX2 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 identify valuable evidence before deciding whether to acquire a given drive or set of files. Files that can be viewed directly on this device 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. Selecting the desired viewable file will activate the **View** button at the bottom of the **Browse** screen. Tap the **View** button to see the text or image file.

OpenText TX2 provides a convenient gallery view feature that allows all images in a selected folder to be viewed as thumbnails. To use gallery view, select a folder from the browse window and then tap the **GALLERY** button at the bottom of the screen. All the images in that folder will be displayed in the new screen, with a maximum of 12 images per page. If there are multiple pages of images, tap the **PREVIOUS PAGE** and **NEXT PAGE** buttons at the bottom of the gallery view screen, to scroll through them all. Tapping an individual thumbnail image will show the full-size image on the screen. If there is more than one viewable image in the selected folder, tapping on the left side or on the right side of the displayed full-size image will display either the previous or next image in full-size, respectively.

Packed log files generated by OpenText TX2 (extension `.tx2_packed_log`) can also be viewed directly on this device. 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 the touchscreen display. For that reason, only the first 256 KB of data is shown when viewing text files directly on this device. 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 OpenText Forensic to see the full contents of any large text files.

OpenText TX2 determines which text and image files are viewable by file extension only. The following file extensions are currently viewable.

<b>Text type</b>	.bat, .c, .conf, .csv, .h, .htm, .html, .ini, .js, .json, .log, .nfo, .py, .readme, .sh, .text, .tsv, .txt, .xml
<b>Image type</b>	.apng, .bmp, .gif, .ico, .cur, .jpg, .jpeg, .jfif, .pjpeg, .ppj, .png, .svg, .webp, .avif
<b>TX2 packed log type</b>	tx2_packed_log

While it is possible to view the text version of a 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, 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 OpenText TX2 forensic 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 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, and DMG). It does not support restoration from a logical image file set (Lx01).

At the beginning of a Restore job, OpenText TX2 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. Your device 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 OpenText TX2.



**Note:** The **Restore** function requires a TX2 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. On 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 `.tx2_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, turn on 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, turn on the 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.



**Tip:** 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 **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 Mobile backup acquisition

OpenText TX2 has the capability to detect iOS- and Android-based mobile devices (phones and tablets) via its USB source port and initiate backup files which can then be acquired. Those backup files contain forensically valuable user data from the devices, and they can be used as input to forensic investigation tools such as OpenText Forensic, and analysis by tools such as OpenText Mobile Investigator.



**Note:** OpenText TX2 can perform backup file acquisitions from most mobile devices running any iOS version and select Android versions. Android has dropped support for Android Debug Bridge (ADB) starting with version 12. OpenText TX2 relies on that interface to detect Android-based mobile devices and acquire their backup files. If OpenText TX2 is not detecting an Android-based device, check its OS version for compatibility.

OpenText TX2 can detect devices that expose a USB Still Image Capture device running either Picture Transfer Protocol (PTP) or Media Transfer Protocol (MTP). These protocols are commonly supported and can serve as an alternative method of acquiring logical (file) based evidence from a mobile device.

The way OpenText TX2 interacts with mobile devices is much different than how it interacts with traditional drives (HDDs and SSDs). That includes the way they are detected, the commands sent, and the information received. This section is focused exclusively on mobile device backup acquisition, including acquisition workflow steps and general information that will maximize the value of your experience with this feature.



**Note:** The access code for a mobile device will be needed to acquire its backup file. Biometric security methods are not typically allowed by mobile devices for the types of host computer interactions needed to acquire their backup files. Newer OS versions will likely require multiple entries of that code at various stages of the mobile backup acquisition workflow.

As a separate mobile device-related feature, OpenText TX2 can display and capture still images and other exposed files from detected mobile devices and cameras that support Picture Transfer Protocol (PTP) or Media Transfer Protocol (MTP). This allows you to directly preview/triage and acquire mobile device evidence without needing to acquire the entire backup file. For more information, see ["Browsing" on page 102](#) and ["Logical imaging" on page 82](#).

## 4.9.1 Connecting and detecting mobile devices

Mobile devices must always be connected to a USB source port on OpenText TX2. The mobile device connector type can be different on different devices, but the communication protocol is always USB. Please be aware that the types of connectors provided or supported by mobile device manufacturers may change over time.

Regardless of the mobile device's connector type, commercial adapter cables are readily available at a reasonable price, and they will be needed to connect different mobile devices to OpenText TX2. Contact OpenText Customer Support if you need assistance finding a suitable USB adapter cable for your mobile device.

After connecting a mobile device to OpenText TX2, it should automatically be detected and shown in the **Sources** list available at the bottom of the **Home** screen. That same **Sources** list is available when selecting a source for a mobile backup acquisition job. However, unlike standard media (HDDs, SSDs), there are some notable nuances with mobile device detection that can cause a connected device to not show as detected/usable. These are covered in the following paragraphs.



**Note:** Mobile devices that support Picture Transfer Protocol (PTP) or Media Transfer Protocol (MTP) will typically show as two separate entries in the Sources list. The main device item (which is what will be used for mobile backup file acquisition) will be shown with a mobile phone icon on the left and the tile will contain the device's name and serial number. The PTP/MTP

portion of the mobile device will be shown with a camera icon on the left and will have an informative message at the top identifying it as a *Picture Transfer Protocol (PTP) device*. For more information about utilizing the PTP/MTP portion of a mobile device to live triage pictures and logically acquire all exposed files, see “[Viewing text and image files](#)” on page 102 and “[Performing a logical image acquisition](#)” on page 83.

For an iOS device, if trust has not been established with OpenText TX2 as the host computer system, the device will still be detected by OpenText TX2, but it will not be selectable as the source for a mobile backup acquisition job, and only limited device information will be displayed. To establish host computer trust on an iOS device after it is attached to OpenText TX2, unlock the screen and acknowledge that you want to trust the attached computer. Once the trust has been established, additional device details are displayed, and the device becomes selectable as the source for a mobile backup acquisition job.

OpenText TX2 can detect an Android device only when the device’s **USB Debugging** setting is turned on. The method to get that setting to appear differs across phone vendors and/or Android OS versions. However, most devices seem to require that they first be put in Developer mode in order to access the **USB Debugging** setting. Check with the vendor of the device to see how to turn on Developer mode. Once **USB Debugging** mode is turned on, trust will be fully established and the device will be selectable for a mobile backup acquisition job.



**Note:** The methods employed by Apple and the various Android phone manufacturers to trust OpenText TX2 as a host computer will likely change over time. Referring to the manufacturer’s documentation for a given phone and OS version will be the best way to establish host computer trust to allow backup file acquisition from that device.

#### 4.9.2 Mobile device details

Once a mobile device is detected and trust has been established between the device and OpenText TX2, additional information about the device will be shown in the **Source Mobile Details** screen on OpenText TX2. That screen can be viewed by tapping on the mobile device tile in either the main Sources list or via the source selection step in the mobile backup acquisition job setup screen.

The following information is available on a mobile device details screen:

- Model
- Serial number
- Name (user entered)
- Vendor
- Operating System
- OS Version
- Product name/number

- Product name (familiar)
- IMEI

This detailed device information will be included in the forensic log for a mobile backup acquisition job.

### 4.9.3 Mobile device media utilities

The available media utilities for a connected mobile device can be viewed at the bottom of the mobile device details screen. They are **Manage Backup Encryption** and **Eject**.

#### Manage Backup Encryption

As indicated in the source selection area of the mobile backup acquisition job stepper, the state of the backup file encryption setting on a mobile device is a key piece of forensic information. In general, if backup file encryption is turned on, more mobile device user information will be included in the backup.

One key difference between iOS and Android devices is how backup encryption is managed. For iOS devices, an informative message will appear in the source selection section of the mobile backup acquisition job setup screen to indicate the encryption state of the mobile backup. If backup encryption is turned off on an iOS device, it can be turned on via the **Manage Backup Encryption** media utility. For Android devices, OpenText TX2 is not able to discern the state of encryption on the device. Care must be taken to ensure that backup encryption is set to the desired state directly on the Android device prior to beginning a mobile backup acquisition job.

Whether backup encryption was previously turned on from the mobile device or turned on as part of the acquisition workflow, the password will need to be noted, as it will be required to access the backup file contents in the upstream forensic investigation software tool.



**Note:** Mobile device backup files can be encrypted (via a setting on the device), which typically results in more user data being included in the backup file, which is forensically desirable. However, the files included in an encrypted backup will typically have different encrypted data from job to job (with the exact same source file data), which makes the encrypted backup file hashes inconsistent between subsequent backup acquisition jobs on the same source device. Keep this in mind as you use mobile backup encryption in your digital forensic investigations.

#### Eject

It is highly recommended to eject all media devices before removing them from OpenText TX2. To eject a mobile device, tap the mobile device tile from the Sources list and the mobile device details screen will appear with the media utilities at the bottom. Tap **Eject** and then confirm the dialog prompt to complete the ejection process. The device may then be physically removed from OpenText TX2.

#### 4.9.4 Performing a mobile backup acquisition

##### To perform a mobile backup acquisition:

1. Using commercially available adapter cables, attach the mobile device to OpenText TX2 USB source port. For the destination, attach a formatted drive to one of the OpenText TX2 destination ports or mount an available iSCSI drive or CIFS/SMB share on your network.



**Note:** Multiple mobile backup acquisition jobs can be run simultaneously, although a USB hub will be needed to provide the required number of USB source ports. Unlike other OpenText TX2 job types, only a single destination can be used with a mobile backup acquisition job. Multiple mobile backup jobs can use a single destination.

2. On the **Home** screen, tap the **Mobile** icon.

The Mobile 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 (with a detected filesystem) are connected, they are automatically selected. If you are satisfied with the default settings and the selected source and destination devices, press the **Start Mobile Backup Acquisition** 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 **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 to make data entry easier.
4. To change or add a source mobile device, tap the **2** or **Source** heading. From the displayed source list modal, select a mobile device from the list. A green check confirms your selection.

If a connected mobile device is not showing in the source list, it is possible that phone interaction/configuration is required. For more information, see ["Connecting and detecting mobile devices" on page 105](#).

As indicated in the source selection area of the mobile backup acquisition job stepper, the state of the backup file encryption setting on a mobile device is a key piece of forensic information. In general, if backup file encryption is turned on, more mobile device user information will be included in the backup.

One key difference between iOS and Android devices is how backup encryption is managed. For iOS devices, an informative message will appear in the source selection section of the mobile backup acquisition job setup screen to indicate the encryption state of the mobile backup. If backup encryption is turned off on an iOS device, it can be turned on via the **Manage Backup Encryption** media utility available in the mobile device details screen. For Android devices,

OpenText TX2 is not able to discern the state of encryption on the device. Care must be taken to ensure that backup encryption is set to the desired state on Android devices prior to beginning a mobile backup acquisition job on OpenText TX2.

Whether backup encryption was previously turned on from the mobile device or turned on as part of the acquisition workflow, the password will need to be noted as it will be required to access the backup file contents in the upstream forensic investigation software tool. See ["Mobile device media utilities" on page 107](#) for more information on mobile device backup encryption for iOS and Android devices.



### Notes

- If there are source drives connected to OpenText TX2 in addition to mobile devices, all of them will show in the sources list, but the non-mobile devices will be grayed out and not selectable for a mobile backup acquisition job.
- To help users identify which source devices have already been acquired, a green checkmark is shown in the bottom right portion of the device tile for any mobile device that has been used in a previous, successful mobile backup acquisition job. For active jobs, a hollow checkmark will appear which will turn solid green once the job has successfully completed.
- Within any screen displaying a list of devices, you can tap the **Options** icon located on the device tile, to see more details and access any available media utilities.

Close the source selection modal by tapping the **X** in the upper right corner or by tapping outside of the modal. If a different source is desired, go back into the **Select a mobile device** screen by tapping on the **2** or **Source** heading from the Mobile job setup stepper.

5. To change or add a destination drive/filesystem tap the **3** or **Destination** heading. From the destination list modal, select one drive/filesystem from the list.



**Note:** Unlike other types of OpenText TX2 jobs, only one destination drive/filesystem can be selected for a mobile backup file acquisition job.

After selecting a filesystem, 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.

To make a network share visible in the destination selection list of a mobile backup acquisition job setup screen, first add and mount the share from the

**Destinations** drive list accessible on the **Home** screen. For more information about mounting iSCSI drives and CIFS/SMB network shares, see “[Duplication over a network](#)” on page 73.

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

The setting options and default setting values are as follows:

- **File output:**
  - **Native** will capture the mobile device’s backup file structure directly to the chosen destination drive/filesystem. This is the factory default setting.
  - **Lx01** will create a logical image file set that encapsulates the **Native** files.
- **Metadata:** OpenText TX2 can create a metadata file in CSV format during a mobile backup acquisition job. This metadata file includes the hash(es) of each acquired backup file, among other things. For more information about the metadata file, see “[Files created during mobile backup acquisition](#)” on page 111. By default, this setting is turned on.
- **Hash type(s):** The available hash types are **MD5** and **SHA-1**. This setting is only available for **Native** file output jobs that have metadata turned on and **Lx01** file output jobs (regardless of the metadata setting). The hash value(s) will be stored in the metadata CSV file (if turned on) for both output types and in the **Lx01** file. The default setting is MD5 and SHA-1 or whatever is set for Hash type(s) in the **Defaults** setting area.



**Note:** Tree Hashing is not supported for mobile backup file imaging. The main benefit of tree hashing is to improve performance when hashing large amounts of block device data, which does not apply to file-based mobile backup file acquisition.

- **Image name:** This will be the base filename for the generated outputs. For iOS sources, this is also the name of the sub-folder that is used to store the actual backup files. The default value is **image** or whatever is set for **Image name** in the **Defaults** setting area. To change the default image name, tap the **Image name** field. For more information about folders and files created during a mobile acquisition job, see “[Files created during mobile backup acquisition](#)” on page 111.
- **File size:** This setting only applies if **Lx01** is selected for the **File output** setting. It sets the size of the **Lx01** segment files that will be written to the destination. Note that the size of the files captured in **Native** format are determined by the mobile device. The **Lx01** segment file size options are: 2 GB, 4 GB, 8 GB, and Unlimited. The default value is 2 GB or whatever is set for **File size** in the **Defaults** setting area.
- **Compression:** This setting only applies if **Lx01** is selected for the **File output** setting. If this setting is turned on, the data will be compressed as it is written out to the **Lx01** segment file(s) to help minimize the overall size of the output file set. By default, this setting is turned on.

- **Readback verification:** This setting only applies if **Lx01** is selected for the **File output** setting. If this setting is turned on, the acquired files in the **Lx01** segments will be read back from the destination and hashed, and then those hashes will be compared to the original acquisition file hashes. The default value is **enabled** or whatever is set for **Readback verification** in the **Defaults** setting area.

7. Once you are satisfied with your settings and device/drive selections, tap **Start Mobile Backup Acquisition**.

A **Job Status** screen is automatically displayed.



**Note:** It may be required to interact with the mobile device after the job has been started on OpenText TX2. It is not always possible to provide informative messaging in the job status screen to indicate that input is required on the mobile device before the actual backup file creation process begins. This can lead to job failures that appear to have no logical explanation. In particular, for Android devices that have backup encryption turned on, the encryption password must be entered prior to accepting the backup request. Accepting the backup request without entering the encryption password will result in a job failure with no indication as to why the job failed. In general and particularly for Android backup jobs, always unlock the device before starting the job and pay close attention to any messaging on the phone to ensure the backup file creation process has actually started.

## 4.9.5 Files created during mobile backup acquisition

The files created by OpenText TX2 during a mobile backup acquisition job are different for each of the supported device OS types (iOS and Android). Even within a given device type, the outputs are different between **Native** and **Lx01** output types. The following sections describe the mobile backup acquisition output files generated in these different scenarios.

### 4.9.5.1 Files created during iOS device backup acquisition

The output files generated by OpenText TX2 during backup acquisition of an iOS based mobile device depend on the **File output** setting. The two options are **Native** and **Lx01**.

#### Native iOS backup acquisition job output files

Currently, the files created by an iOS mobile device during a **Native** file type backup are written to the destination drive according to the following convention:

```
[image base directory]/
[directory name]/
[image name]/
[iOS backup folder name]
[iOS sub-folder name 1]
[iOS filename 1]
:
:
```

```
[iOS filename n]
[iOS sub-folder name 2]
.
.
.
[iOS sub-folder name n]
Info.plist
Manifest.db
Manifest.plist
Status.plist
[image name].csv
[image name].log.html
[image name].log.txt
```

[image base directory] is defined in **Setting Defaults** or when selecting a destination drive/filesystem. The default is /tx2\_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 mobile backup acquisition job setup. The default setting is Date and Time.

[iOS backup folder name] is the base folder as provided by the iOS device which contains all the backup content in its sub-folders and files. OpenText TX2 has no control over the naming of this folder nor any of the folder/file names contained therein, and they are subject to change at Apple's discretion.

[image name].csv is the generated mobile backup acquisition job metadata file. It contains the following information for each file acquired during the backup job:

- Path – The overall path name of the folder/file as written to the destination drive/filesystem.
- Type – Identification of the path type (Directory or File).
- Filesize – Size of any entries of type file (in bytes).
- Date/time stamps – These are the dates/times related to when the files were written to the OpenText TX2 destination during the backup acquisition job, not the dates/times of the original source files from the mobile device's perspective.
- MD5 Hash/SHA1 Hash – Acquired file hash values as calculated by reading back the files that were written to the destination. This is not exactly the same as an acquisition hash (which is created for physical and logical imaging jobs before the data is written to the destination), but it is the best that can be done to mimic an acquisition hash for mobile backup acquisition jobs. The hash values in this CSV file are what is used to readback verify a **Native** format backup acquisition job.



**Note:** Mobile device backup files can be encrypted (via a setting on the device), which typically results in more user data being included in the backup file, which is forensically desirable. However, the files included in an encrypted backup will typically have different encrypted data from job to job (with the exact same source file data), which makes the encrypted backup file hashes inconsistent between subsequent backup acquisition jobs

on the same source device. Keep this in mind as you use mobile backup encryption in your digital forensic investigations.

- File Status – Status of the file as read back from the destination during a mobile backup acquisition job. If the job completed successfully and there were no errors reading back the files from the destination (when creating the hashes for the metadata file), all will show as “OK”. If there was an issue while reading the files back, the job will fail and the offending file will show an “Error” status in the CSV file.
- Matched Rules – For mobile backup acquisition jobs, this will always indicate “Y” for yes. This field is more pertinent to logical imaging jobs, which use it to indicate if a given file was part of a targeted/filtered collection.

[image name].log.html/txt – These are the two forensic log files (HTML and text formats) generated for each job (mobile backup acquisition included). These logs are accessible through the **Logs** list in the side navigation menu. They can also be exported to a destination drive/filesystem (local or network based) for further analysis and/or case documentation purposes.

### Lx01 iOS backup acquisition job output files

When **Lx01** is selected as the file output type for an iOS backup acquisition job, the same **Native** iOS backup files are still acquired. However, instead of the **Native** files being kept on the destination drive, they are read back into OpenText TX2 and then packed into **Lx01** segment files that are then written to the destination. This type of output helps protect the **Native** backup files from unintentional modification, just as a logical image job does with files from a source filesystem. The files created by an iOS mobile device during a **Lx01** file type backup are written to the destination drive according to the following convention:

```
[image base directory]/
  [directory name]/
    [image name].csv
    [image name].log.html
    [image name].log.txt
    [image name].Lx01
    [image name].Lx02
    .
    .
    .
    [image name].Lx99, etc.
```

[image base directory] is defined in Setting Defaults or when selecting a destination drive. The default is /tx2\_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 mobile backup acquisition job setup. The default setting is Date and Time.

[image name].csv is the generated mobile backup acquisition job metadata file. It contains the following information for each **Native** iOS file acquired during the backup job:

- Path – The overall path name of the folder/file as written to the destination drive/filesystem during the backup acquisition job. Note that those folders/files will not

be seen directly on the destination after completion of the Lx01 output job, but they represent the file structure that is encapsulated in the Lx01 file set.

- Type – Identification of the path type (Directory or File).
- Filesize – Size of any entries of type file (in bytes).
- Date/time stamps – These are the dates/times related to when the files were written to the OpenText TX2 destination during the backup acquisition job, not the dates/times of the original source files from the mobile device's perspective.
- MD5 Hash/SHA1 Hash – Acquired file hash values as calculated by reading back the files that were written to the destination (as encapsulated in the Lx01 segment files). This is not exactly the same as an acquisition hash (which is created for physical and logical imaging jobs before the data is written to the destination), but it is the best that can be done to mimic an acquisition hash for mobile backup acquisition jobs. The hash values in this CSV file are what is used to readback verify an Lx01 format backup acquisition job.



**Note:** Mobile device backup files can be encrypted (via a setting on the device), which typically results in more user data being included in the backup file, which is forensically desirable. However, the files included in an encrypted backup will typically have different encrypted data from job to job (with the exact same source file data), which makes the encrypted backup file hashes inconsistent between subsequent backup acquisition jobs on the same source device. Keep this in mind as you use this feature in your digital forensic investigations.

- File Status – Status of the file as read back from the destination during a mobile backup acquisition job. If the job completed successfully and there were no errors reading back the files from the destination (to pack them into the Lx01 segment files), all will show as "OK". If there was an issue while reading the files back, the job will fail and the offending file will show an "Error" status in the CSV file.
- Matched Rules – For mobile backup acquisition jobs, this will always indicate "Y" for yes. This field is more pertinent to logical imaging jobs, which use it to indicate if a given file was part of a targeted/filtered collection.

[image name].log.html/txt – These are the two forensic log files (HTML and text formats) generated for each job (mobile backup acquisition included). These logs are accessible through the **Logs** list in the side navigation menu. They can also be exported to a destination drive/filesystem (local or network based) for further analysis and/or case documentation purposes.

[image\_name].Lx01, [image\_name].Lx02, ... are the forensic evidence files for the operation. They contain all the data and metadata for each acquired mobile backup file and folder.

### 4.9.5.2 Files created during Android device backup acquisition

The output files generated by OpenText TX2 during backup acquisition of an Android based mobile device depend on the **File output** setting. The two options are **Native** and **Lx01**.

#### Native Android backup acquisition job output files

Currently, the files created by an Android mobile device during a **Native** file type backup are written to the destination drive according to the following convention:

```
[image base directory]/
  [directory name]/
    [image name]/
      [image name].ab
    [image name].csv
    [image name].log.html
    [image name].log.txt
```

[image base directory] is defined in **Setting Defaults** or when selecting a destination drive. The default is `/tx2_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 mobile backup acquisition job setup. The default setting is Date and Time.

[image name] is the base folder used to store the Android backup file ([image name].ab).

[image name].csv is the generated mobile backup acquisition job metadata file. It contains the following information for each file acquired during the backup job:

- Path – The overall path name of the folder/file as written to the destination drive/filesystem.
- Type – Identification of the path type (Directory or File).
- Filesize – Size of any entries of type file (in bytes).
- Date/time stamps – These are the dates/times related to when the files were written to the OpenText TX2 destination during the backup acquisition job, not the dates/times of the original source files from the mobile device's perspective.
- MD5 Hash/SHA1 Hash – Acquired file hash values as calculated by reading back the files that were written to the destination. This is not exactly the same as an acquisition hash (which is created for physical and logical imaging jobs before the data is written to the destination), but it is the best that can be done to mimic an acquisition hash for mobile backup acquisition jobs. The hash values in this CSV file are what is used to readback verify a **Native** format backup acquisition job.



**Note:** Mobile device backup files can be encrypted (via a setting on the device), which typically results in more user data being included in the backup file, which is forensically desirable. However, the files included in an encrypted backup will typically have different encrypted data from job

to job (with the exact same source file data), which makes the encrypted backup file hashes inconsistent between subsequent backup acquisition jobs on the same source device. Keep this in mind as you use this feature in your digital forensic investigations.

- File Status – Status of the file as read back from the destination during a mobile backup acquisition job. If the job completed successfully and there were no errors reading back the files from the destination (when creating the hashes for the metadata file), all will show as “OK”. If there was an issue while reading the files back, the job will fail and the offending file will show an “Error” status in the CSV file.
- Matched Rules – For mobile backup acquisition jobs, this will always indicate “Y” for yes. This field is more pertinent to logical imaging jobs, which use it to indicate if a given file was part of a targeted/filtered collection.

[image name].log.html/txt – These are the two forensic log files (HTML and text formats) generated for each job (mobile backup acquisition included). These logs are accessible through the **Logs** list in the side navigation menu. They can also be exported to a destination drive/filesystem (local or network based) for further analysis and/or case documentation purposes.

### **Lx01 Android backup acquisition job output files**

When **Lx01** is selected as the file output type for an Android backup acquisition job, the same **Native** Android backup file is still acquired. However, instead of the **Native** file being kept on the destination drive, it is read back into OpenText TX2 and then packed into **Lx01** segment files that are then written to the destination. This type of output helps protect the **Native** backup file from unintentional modification just as a logical image job does with files from a source filesystem. The files created by an Android mobile device during a **Lx01** file type backup are written to the destination drive according to the following convention:

```
[image base directory]/  
[directory name]/  
  [image name].csv  
  [image name].log.html  
  [image name].log.txt  
  [image name].Lx01  
  [image name].Lx02  
  .  
  .  
  [image name].Lx99, etc.
```

[image base directory] is defined in **Setting Defaults** or when selecting a destination drive. The default is `/tx2_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.

[image name].csv is the generated mobile backup acquisition job metadata file. It contains the following information for each **Native** iOS file acquired during the backup job:

- Path – The overall path name of the folder/file as written to the destination drive/filesystem during the backup acquisition job. Note that those folders/files will not be seen directly on the destination after completion of the Lx01 output job, but they represent the file structure that is encapsulated in the Lx01 file set.
- Type – Identification of the path type (Directory or File).
- Filesize – Size of any entries of type file (in bytes).
- Date/time stamps – These are the dates/times related to when the files were written to the OpenText TX2 destination during the backup acquisition job, not the dates/times of the original source files from the mobile device's perspective.
- MD5 Hash/SHA1 Hash – Acquired file hash values as calculated by reading back the files that were written to the destination (as encapsulated in the Lx01 segment files). This is not exactly the same as an acquisition hash (which is created for physical and logical imaging jobs before the data is written to the destination), but it is the best that can be done to mimic an acquisition hash for mobile backup acquisition jobs. The hash values in this CSV file are what is used to readback verify an Lx01 format backup acquisition job.



**Note:** Mobile device backup files can be encrypted (via a setting on the device), which typically results in more user data being included in the backup file, which is forensically desirable. However, the files included in an encrypted backup will typically have different encrypted data from job to job (with the exact same source file data), which makes the encrypted backup file hashes inconsistent between subsequent backup acquisition jobs on the same source device. Keep this in mind as you use this feature in your digital forensic investigations.

- File Status – Status of the file as read back from the destination during a mobile backup acquisition job. If the job completed successfully and there were no errors reading back the files from the destination (to pack them into the Lx01 segment files), all will show as "OK". If there was an issue while reading the files back, the job will fail and the offending file will show an "Error" status in the CSV file.
- Matched Rules – For mobile backup acquisition jobs, this will always indicate "Y" for yes. This field is more pertinent to logical imaging jobs, which use it to indicate if a given file was part of a targeted/filtered collection.

[image name].log.html/txt – These are the two forensic log files (HTML and text formats) generated for each job (mobile backup acquisition included). These logs are accessible through the **Logs** list in the side navigation menu. They can also be exported to a destination drive/filesystem (local or network based) for further analysis and/or case documentation purposes.

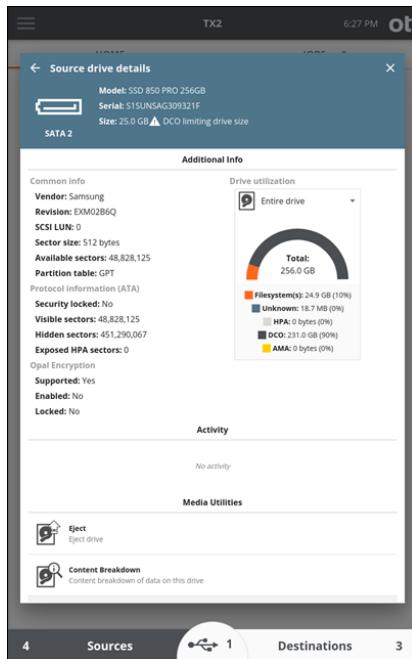
[image name].Lx01, [image\_name].Lx02, ... are the forensic evidence files for the operation. They contain all the data and metadata for each acquired mobile backup file and folder.

## 4.10 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  on 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 OpenText Forensic. Even though OpenText TX2 could not parse or detect partitions and filesystems in this scenario, there is still forensically valuable information that can be carved or extracted using OpenText Forensic.

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 will act on the entire drive.



**Note:** The information in this section 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 OpenText TX2.

Mobile devices (phones, tablets) and digital cameras that support Picture Transfer Protocol (PTP) or Media Transfer Protocol (MTP) and are connected to a USB source port will appear in the Sources list as well. Some mobile devices may show two entries in the list, one for the overall device (for mobile backup acquisition purposes) and one for the embedded PTP/MTP device (for browsing, viewing, and logical imaging of exposed files).

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 solid 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.



**Note:** The green checkmark that indicates a given drive has been successfully acquired will persist in the user interface until your device is power-cycled, at which time no drives will indicate they were acquired (despite the fact that they may have been in a previous session). Also, if an acquired drive is removed and then reinstalled on OpenText TX2 within the same session/power-cycle, that drive will show as having been acquired.

#### 4.10.1 Encryption detection

OpenText TX2 automatically detects certain types of encryption present on attached drives. This detection is possible for software-based encryption types that have known header signatures. OpenText TX2 can also detect the hardware-based Opal encryption method. It 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, OpenText TX2 can also unlock and perform any supported operations on a drive/partition encrypted with BitLocker. For more information, see “[BitLocker encryption](#)” on page 44.

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

#### 4.10.1.1 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.

OpenText TX2 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 **Encryption Unlock** media utility. Once unlocked, an Opal SED can be read from (or written to, in the case of a destination drive). For information related to unlocking Opal SEDs, see “[Encryption Unlock](#)” on page 42.

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

An additional consideration for Opal drives is a unique configuration that exposes a Shadow MBR. This Shadow MBR can be enabled by drive/system manufacturers to

initially identify the drive as a small, non-encrypted volume, which overrides the actual MBR information. A typical use case for this configuration is to enable system manufacturers to request credentials from a user before revealing the actual MBR information on the drive. Regardless of the use case, it is important to be able to identify situations where only the Shadow MBR is revealed, to make it clear that the entire drive contents are not being seen. OpenText TX2 will detect when an Opal Shadow MBR is enabled and clearly inform of its presence. The lock icon will show in the affected drive tile in the Sources list, and the presence of an Opal MBR will be explicitly called out in the drive details screen. Note that the Shadow MBR configuration is essentially a unique form of a locked Opal drive, therefore unlocking the Opal encryption on OpenText TX2 will disable the Shadow MBR (regardless of the underlying encryption state) and make the full, unencrypted drive contents available for triage/acquisition. Also, Opal encryption unlock (including Shadow MBR disablement) is a volatile change, meaning that the drive will revert to its original configuration after it is power cycled.



### Caution

Docking station type devices that contain Opal drives must support ATA command pass-through so that OpenText TX2 properly detects the presence of Opal encryption and allows 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 OpenText TX2 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 user interface, removing the drive from the enclosure and connecting it directly to OpenText TX2 may yield the desired outcome.

#### 4.10.1.2 Apple Core Storage and FileVault 2

OpenText TX2 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.



**Note:** There are many possible ways to configure Apple Core Storage volumes with FileVault 2 encryption. OpenText TX2 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, a warning message will indicate that Core Storage has been detected and FileVault 2 is possible.

#### 4.10.2 RAID detection

In addition to encryption detection, OpenText TX2 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.

OpenText TX2 can 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. Together with 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.*

## 4.11 Logs module

OpenText TX2 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. The following is a summary of the information captured for an image-based duplication job. For specific job log examples, see [sample logs](#).

- **Status** – Overall job status (Incomplete, Ok, Error/Failed, Cancelled) as well as date/time stamps, username, job notes, and OpenText TX2 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 filesystem specific information (if present and supported).
- **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 turned on. 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 this setting was turned on 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). The **Log List** can be filtered to show only specific logs of interest. This is particularly useful when exporting or deleting logs from the system. For more information, see “[Filtering logs](#)” on page 129.

To export all displayed logs (can be the entire list or a filtered list), tap the **Export** 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 OpenText TX2, 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 displayed logs (can be all logs or a filtered list) for OpenText TX2 previously completed (inactive) jobs, tap the **Delete Inactive** button on the bottom of the **Log List** screen. When completed, the only remaining logs will be for any jobs that were active when the delete action was initiated or any logs that were excluded through filtering.



**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.

The following log management options are available at the bottom of the log details screen:

- **Resume Job** – If highlighted (orange), then the job is resumable. For more details, see “[Pausing and resuming a duplication job](#)” on page 75.
- **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 OpenText TX2 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.

### 4.11.1 HTML logs

OpenText TX2 can 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 this device, but this view can easily be switched back to the legacy text view, using the **Options** menu  at the bottom of any log details screen.

A plus sign displayed in the HTML log in a section header 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 pressing the + or - symbol near the section header, respectively), there is a button at the top 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.

### 4.11.2 Sample logs

Two sample logs are shown in this section: one from a successful duplication and one from a failed standalone verification. Note that, while the HTML log view is the default when viewing logs on the OpenText TX2 user interface, the following sample logs are shown in legacy text format, as it provides a better view for showing all the log information at once.

#### Log 1

```
----- Start of TX2 Log Entry -----  
Task: Disk Duplication  
Status: Ok  
Created: Fri Jun  6 13:50:35 2025 (UTC-0500)  
Started: Fri Jun  6 13:50:35 2025 (UTC-0500)  
Closed: Fri Jun  6 13:55:40 2025 (UTC-0500)  
Elapsed: 5 min  
Username: coop
```

```
Examiner: Dale Cooper
Case ID:
  Laura Palmer Murder
Case Notes:
  Second Diary

Imager App: TX2
Imager Ver: 25.2.0
Imager S/N: 78F7A3000100000F

-----Source Disk-----

Interface: NVME
Port: PCIe 2
Model: NVMe 0x126F APPLE SSD SSW256GB
Firmware revision: S0318A0G
Serial number: 19122610375
NVMe namespace: 1
Capacity in bytes: 251,059,544,064 (251.0 GB)
Block Size: 512 bytes
Block Count: 490,350,672
Partition table: GPT
Partition 1
  Start Sector: 2,048
  Sector Count: 490,348,544
  Partition Encryption: None detected
    Filesystem
      Type: exFAT
      Block Size: 131,072 bytes
      Total Blocks: 1,915,296
      Free Blocks: 1,915,284
      Total bytes: 251,041,677,312 (251.0 GB)
      In Use bytes: 1,572,864 (1.5 MB)
  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: 4,000,000,000 (4.0 GB)

-----Image Destination-----

Interface: USB
Port: USB 1
Model: SanDisk Extreme Pro 55AF
Firmware revision: 4060
Serial number: 243621402765
SCSI LUN: 0
USB Serial number: 323433363231343032373635
USB VID: 0x0781
USB PID: 0x55af
USB interface class: 0x08
USB interface subclass: 0x06
USB interface protocol: 0x62
Capacity in bytes: 1,000,171,323,904 (1.0 TB)
Block Size: 512 bytes
Block Count: 1,953,459,617
Partition table: GPT
Partition 1
  Start Sector: 34
  Sector Count: 32,734
  Partition Encryption: None detected
  Filesystems: None Recognized
Partition 2
  Start Sector: 32,768
  Sector Count: 1,953,425,408
  Partition Encryption: None detected
  Filesystem (Target Filesystem)
```

```

Type: exFAT
Block Size: 4,096 bytes
Total Blocks: 243,939,328
Free Blocks: 243,878,893
Total bytes: 999,175,487,488 (999.1 GB)
In Use bytes: 247,541,760 (247.5 MB)
Tableau Encrypted: No
Whole disk encryption: None detected
Folder: /tx2_images/2025_06_06_13_50_35/
File name base: image
Verification Status: Finished OK
    Verification Tree Sha1: da3c 9685 385d a17c e83c a5bb 775a a365 9638 593e
    Verification Tree Md5: 035d 8a95 f1fa 4d0a c7c8 ba22 f90c 7f71

-----Duplication Results-----

LBA Range Duplicated: Entire Source Disk
Total recoverable errors: 0
Total unrecoverable errors: 0
Acquisition Tree Sha1: da3c 9685 385d a17c e83c a5bb 775a a365 9638 593e
Acquisition Tree Md5: 035d 8a95 f1fa 4d0a c7c8 ba22 f90c 7f71

-----End of TX2 Log Entry-----

```

## Log 2

```

-----Start of TX2 Log Entry-----

*** CAUTION: THE OPERATION RECORDED IN THIS LOG DID NOT COMPLETE NORMALLY ***

Task: Image Readback Verification
Status: Error/Failed
Created: Fri Jun  6 14:08:46 2025 (UTC-0500)
Started: Fri Jun  6 14:08:47 2025 (UTC-0500)
Failed: Fri Jun  6 14:09:08 2025 (UTC-0500)
    Destination unreadable - 0xfb1ecc28838903d5
Closed: Fri Jun  6 14:09:08 2025 (UTC-0500)
Elapsed: 21 sec
Username: coop
Examiner: Dale Cooper
Case ID:
    Laura Palmer Murder
Case Notes:
    Second Diary

Imager App: TX2
Imager Ver: 25.2.0
Imager S/N: 78F7A3000100000F

-----Image Destination-----

Interface: USB
Port: USB 1
Model: SanDisk Extreme Pro 55AF
Firmware revision: 4060
Serial number: 243621402765
SCSI LUN: 0
USB Serial number: 323433363231343032373635
USB VID: 0x0781
USB PID: 0x55af
USB interface class: 0x08
USB interface subclass: 0x06
USB interface protocol: 0x62
Capacity in bytes: 1,000,171,323,904 (1.0 TB)
Block Size: 512 bytes
Block Count: 1,953,459,617

```

```
Partition table: GPT
Partition 1
  Start Sector: 34
  Sector Count: 32,734
  Partition Encryption: None detected
  Filesystems: None Recognized
Partition 2
  Start Sector: 32,768
  Sector Count: 1,953,425,408
  Partition Encryption: None detected
    Filesystem (Target Filesystem)
      Type: exFAT
      Block Size: 4,096 bytes
      Total Blocks: 243,939,328
      Free Blocks: 182,556,497
      Total bytes: 999,175,487,488 (999.1 GB)
      In Use bytes: 251,424,075,776 (251.4 GB)
Tableau Encrypted: No
Whole disk encryption: None detected
Folder: /tx2_images/2025_06_06_14_03_21/
File name base: image
File format: Ex01

-----Failure Summary-----
Reason for failure: Destination unreadable
Failure code: 0xfb1ecc28838903d5

-----End of TX2 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 OpenText TX2. The options for this field are: **No**, **Locked**, and **Unlocked**.
- **Whole disk encryption:** This field is populated with the specific type of third-party whole disk encryption that OpenText TX2 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. OpenText TX2 indicates that **FileVault 2** encryption is possible when a **Core Storage** partition is detected.

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

If OpenText TX2 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 32 KB 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 functionality is turned on and your device 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 recommended to export and delete logs from your device after each case. OpenText TX2 can 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.

### 4.11.3 Filtering logs

OpenText TX2 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, tap the log filter icon  at the bottom of the log list screen. The available filter parameters include the following:

- Case information: **Examiner name**, **Case ID**, and **Notes**.
- Drive information: **Model**, **Vendor**, **Serial**, and **Network address** (for CIFS/iSCSI).
- Job type: **Duplicate**, **Logical**, **Mobile Backup**, **Reconfigure**, **Restore**, **Hash**, **Verify**, **SMART**, **HPA/DCO/AMA**, **NVMe Namespace Attach**, and **Blank Check**.
- Job status: **Started**, **Ok**, **Failed**, **Cancelled**, and **Paused**.
- Dates (less than and greater than or equal to an entered job start date).

Each of the filter fields can be populated manually by typing the desired value in the field(s) of interest. Alternatively, for the case and drive information fields, auto-population is possible. To auto-populate these fields, 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.



**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, 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 near 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, 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 re-select 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.

## 4.12 Remote web interface

OpenText TX2 allows you to use a web browser to connect remotely to any units connected to the network. If your unit has a user setup with remote access privileges and a valid IP address, remote access is as simple as opening a web browser and entering the IP address or hostname of OpenText TX2 into the address field.

A login screen will always appear first on the remote browser, and, after credentials are entered and validated, the OpenText TX2 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 unit. However, 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 pull-down 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, 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 **TX2-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. Tap the desired log tile from the main log list. In the **Log Details** screen, click the **Options** icon  and select **Download**.



**Note:** The **Download** button is not displayed when accessing the unit 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 OpenText TX2. However, the logs for any remotely initiated jobs will reflect the time zone as set on OpenText TX2.
- The **Factory Reset** feature is not available to remote users.
- Locking/ unlocking the system (PIN lock) from any location (local or any remote user instance) will lock/unlock all the active screens.

### 4.12.1 SSL certificate setup and installation

OpenText TX2 uses SSL certificates to ensure secure communication during remote sessions. By default, it generates a new self-signed certificate on power up if one does not already exist on the system. A new self-signed certificate will also be generated after a Factory Reset. These self-signed certificates expire after one year, and OpenText TX2 will automatically generate a new one before expiration. Users can manually initiate a new self-signed certificate or install their own SSL certificates.

#### To generate a new self-signed certificate or install a user certificate:

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.

You have two options:

- To create a new self-signed certificate:
  - a. Tap **GENERATE NEW CERT**.  
A warning modal appears asking you to confirm generation of a new self-signed SSL certificate and system reboot.
  - b. Tap **Generate** to confirm the selection and reboot the system.
- To install your own SSL certificate on OpenText TX2:
  - a. Tap **INSTALL USER CERT**.
  - b. Select the desired drive/filesystem from the list.  
A file browser window is displayed.
  - c. 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.
  - d. Select **Install**.  
The system reboots.

Once you install your own certificate, OpenText TX2 will retain it in the event of reboot or power disruption. Manually generating a self-signed

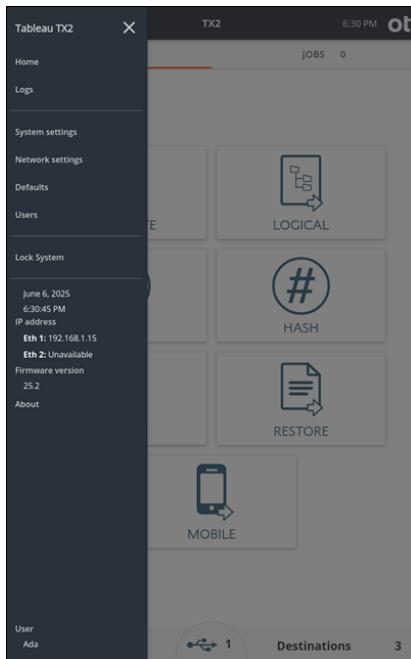
certificate will overwrite your own certificate and return your device to the default state of generating a new self-signed certificate upon annual expiration.

## 4.12.2 Remote access

To access the OpenText TX2 remotely, a valid username and password/Smart Card PIN are required. The user account must also have remote access turned on.

### To access OpenText TX2 remotely:

1. On the **User Management** screen, confirm the user is set up on the system with a password/Smart Card PIN and remote access is turned on. Before turning on the remote access, a password/Smart Card PIN must be added by tapping the **ADD AUTHENTICATION** button. Once a password/Smart Card PIN has been entered, toggle the **Allow Remote** switch.
2. Connect OpenText TX2 to your local area network with an Ethernet cable.
3. Obtain the assigned IP address from the side navigation menu, as shown in the following image.



4. Open a web browser on a computer/device that is connected to the same local area network as OpenText TX2, type the unit'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:** If you have not installed a user certificate on your OpenText TX2 (see “[HTTPS certificate setup](#)” on page 27), your browser will not be able to validate the authenticity of the OpenText TX2 web server, and will display a security warning message that will require exception acknowledgement to proceed. Installing your own HTTPS certificate will establish a trust relationship between the OpenText TX2 web server and your remote web browser.

5. Enter the username and password/Smart Card PIN.

The OpenText TX2 user interface appears in your browser.



**Note:** If any OpenText TX2 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 unit associated with each browser window, which is helpful in a lab environment with multiple units being accessed remotely from the same browser.



# Chapter 5

## Adapters

This chapter describes the add-on drive adapters available for OpenText TX2, which extend its imaging capabilities in an easy to connect and use manner.



**Note:** While OpenText TX2 supports PCIe hot-swap, it is required to eject any PCIe drive prior to removal. While this is generally recommended for all drive types, it is of particular importance for PCIe drives, due to the inherent nature and complexity of the relationship between PCIe peripherals and the host processor system. For more information about OpenText Forensic PCIe Adapters, contact your OpenText Forensic sales representative or go to [opentext.com/products/digital-forensic-hardware](https://www.opentext.com/products/digital-forensic-hardware) (<https://www.opentext.com/products/digital-forensic-hardware>).

The following drive adapters are available individually or as part of an adapter kit for OpenText TX2.

SKU #	Description	Use case
TA7-1	PCIe add-in Card SSD Adapter	PCIe add-in Card SSD to connect to any PCIe-compatible OpenText Forensic devices.
TA7-2	PCIe M.2 SSD Adapter	PCIe M.2 SSD to connect to any PCIe-compatible OpenText Forensic devices.
TA7-3	PCIe Adapter for Apple SSD 2013 and later	Apple PCIe SSD 2013 and later to connect to any PCIe-compatible OpenText Forensic devices.
TA7-4	PCIe U.2 SSD Adapter	Interface type U.2 (SFF-8639) SSD to connect to any PCIe-compatible OpenText Forensic devices.
TA7-5	PCIe-IDE Adapter	IDE devices to connect to any PCIe-compatible OpenText Forensic devices.
TA7-7	PCIe Adapter for Apple SSD 2016 and later	Apple PCIe SSD 2016 and later to connect to any PCIe-compatible OpenText Forensic devices.
TA7-9	PCIe to FireWire Adapter	Connects compatible OpenText Forensic devices to FireWire device or Mac in Target Disk Mode.

SKU #	Description	Use case
TDA3 - 1	3.5 in. SATA - micro SATA Hard Disk Adapter	Connects micro-SATA SSDs to any SATA-compatible OpenText Forensic devices.
TDA3 - 2	SATA - Blade-Type Solid State Drive (SSD) Adapter	Mac Air BLADE Type SSDs to connect to any SATA-compatible OpenText Forensic devices.
TDA3 - 3	mSATA/M.2 SATA SSD Adapter	mSATA or M.2 SATA SSD connection to any SATA compatible OpenText Forensic devices.
TKDA3 - LIF	SATA LIF Hard Disk Adapter with two LIF cables in a Black Nylon Bag	Allows 1.8" Mac Air SATA LIF hard drives to connect to any SATA-compatible OpenText Forensic device.
TDA5 - 18	3.5 in. - 1.8 in. IDE Hard Disk Adapter	Allows 1.8" notebook IDE hard disks to connect to any IDE-compatible OpenText Forensic devices.  Compatible via TA7 - 5 adapter with TC - PCIE4 - 4 / TC - PCIE4 - 8.
TDA5 - 25	3.5 in. - 2.5 in. IDE Hard Disk Adapter	Allows 2.5" notebook IDE hard disks to connect to any IDE-compatible OpenText Forensic devices.  Compatible via TA7 - 5 adapter with TC - PCIE4 - 4 / TC - PCIE4 - 8.
TKDA5 - ZIF	3.5 in. - ZIF IDE Hard Disk Adapter, TC20-BNDL, in a Black Nylon Bag	Allows 1.8" ZIF hard disks to connect to any IDE-compatible OpenText Forensic devices.  Compatible via TA7 - 5 adapter with TC - PCIE4 - 4 / TC - PCIE4 - 8.

## 5.1 PCIe SSD adapters

OpenText Forensic PCIe SSD adapters allow the use of PCIe based SSDs of various types via OpenText TX2 PCIe source and destination ports. The following adapters are available individually or as part of an adapter kit:

- PCIe add-in card SSD adapter – TA7 - 1
- PCIe m.2 SSD adapter – TA7 - 2
- PCIe adapter for Apple SSD 2013-2017 – TA7 - 3
- PCIe u.2 SSD adapter cable – TA7 - 4

- PCIe adapter for Apple SSD 2016-2017 – TA7-7

## 5.2 PCIe IDE adapter (TA7-5)

The OpenText Forensic PCIe IDE Adapter (TA7-5) allows the acquisition of IDE drives via an OpenText TX2 PCIe source port. The two-piece bundle IDE/FireWire adapter kit (TKA-PCIE-2PC) is sold separately and 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 the OpenText TX2 (TC-PCIE4-8) is included in the OpenText TX2 kit.

### 5.2.1 Using the TA7-5

1. With unit powered off, attach the TA7-5 to OpenText TX2 by connecting it to one of the source PCIe ports on the left side of the unit, using a TC-PCIE4-4 or TC-PCIE4-8 cable.
2. Connect an IDE drive using the TC6-2 ribbon cable (blue connector goes to the TA7-5) and TC2-8-R3 power cable. Note that the power cable should plug directly into the TA7-5 adapter.
3. Power on the unit.
4. The **Sources** drive counter will increment by one to let you know your IDE drive is connected. Tap the **Sources** button at the bottom of the **Home** screen to view IDE drive details.

## 5.3 PCIe FireWire adapter (TA7-9)

The OpenText Forensic PCIe FireWire adapter (TA7-9) allows the acquisition of FireWire drives via an OpenText TX2 PCIe source port. The two-piece bundle IDE and FireWire adapter kit (TKA-PCIE-2PC) is sold separately and 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 OpenText TX2 (TC-PCIE4-8) is included in the OpenText TX2 kit.



**Note:** There is a 9-pin 1394b port and a 6-pin 1394a port on TA7-9. If media are attached to both FireWire ports, both will show as available drives in the OpenText TX2 user interface and be simultaneously usable for any available drive actions (imaging, hashing, etc.).

## 5.4 Adapting SAS drives

The source and destination SATA ports on OpenText TX2 are dedicated SATA ports that do not support SAS drives.

To acquire SAS drives, a separate adapter must be used. One option is an OpenText T6u Forensic SAS bridge (sold separately) connected to a USB source port on the imager. Alternatively, there are many high quality and reliable commercially available SAS-USB adapters that can be used.

## 5.5 Apple Target Disk Mode acquisition adapters

OpenText TX2 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. Depending on the Apple computer interface connection, different adapters and/or cables are required to allow for connection to the source side of your OpenText TX2.



**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, OpenText Forensic Equipment devices do not support the special mechanisms to expose Apple memory devices protected by a T2 chip (or any of its follow-on variants), nor a way of providing credentials to allow for the acquisition of the unencrypted data.

### 5.5.1 FireWire adapter cable

To connect an Apple Mac device to OpenText TX2 via the Mac's FireWire interface will require three items. A FireWire 800 cable (9-pin to 9-pin variant) is required to connect between the Mac and an OpenText Forensic PCIe FireWire Adapter (TA7-9). That adapter then connects to OpenText TX2 with an OpenText Forensic PCIe cable (TC-PCIE4-8 or TC-PCIE4-4). Different Macintosh computers have different FireWire connectors, so check the connector type before you begin.



### 5.5.2 Thunderbolt 2 adapter cable

Adapting from a Thunderbolt 2 connector on a Mac to OpenText TX2 requires four separate adapters/cables. A Thunderbolt 2 to FireWire 800 (9-pin) adapter is used along with a FireWire 800 (9-pin to 9-pin) cable to connect between the Thunderbolt 2 port on the Macintosh and an OpenText Forensic PCIe FireWire Adapter (TA7-9). Then connect the TA7-9 adapter to one of the source PCIe ports on OpenText TX2 with an OpenText Forensic PCIe cable (TC-PCIE4-8 or TC-PCIE4-4).





# Chapter 6

## Specifications and troubleshooting

This chapter provides a list of OpenText TX2 specifications and information for troubleshooting the most common issues you may encounter when using this product.

### 6.1 Specifications

Connectors: Source Side	
USB	Two USB 3.2 Gen 2 (10 Gbps) Type-C connectors
PCIe	Two PCIe (Gen 3.0 x4) connectors
SATA	Two SATA (6 Gbps) signal connectors
Drive Power	Two 4-pin Molex Mini-Fit power connectors for SATA drive power
Connectors: Destination Side	
USB	Two USB 3.2 Gen 2 (10 Gbps) Type-C connectors
PCIe	Two PCIe (Gen 3.0 x4) connectors
SATA	Two SATA (6 Gbps) signal connectors
Drive Power	Two 4-pin Molex Mini-Fit power connectors for SATA drive power
TX2-S1	One TX2-S1 signal connector (supporting two SATA 6 Gbps interfaces on the bottom of the unit)
Connectors: Miscellaneous	
Ethernet	Two 10 Gbps Ethernet Connections (Source or Destination); auto-negotiable to 1 Gbps; support for port bonding
USB	Two USB 3.2 Gen 1 (5 Gbps) Type-C Connectors
SD Card	One SD Card Connector for Device Firmware
DC Input	One barrel connector for use with the OpenText Forensic TP8 Power Supply
User Interface	
LCD	8.0 in. graphic LCD (800 x 1280 Resolution) with capacitive touchscreen
Power Button	One on/off power button
Indicators	
Power Indicator	White LED integrated into the power button indicates the unit is powered on; Flashes for power fault

Status Indicator	Multi-color LED located in the lower-right corner of the top panel. Color key: <ul style="list-style-type: none"><li>• White – System startup</li><li>• Blue – Active job</li><li>• Green – Completion of successful job</li><li>• Red – Job failure</li></ul>
Speaker	Audio tones indicate job completion or failure
<b>Physical / Environmental</b>	
Power	60 Watts typical operating power (not including external drive power)  250 Watts maximum allowed power draw (including external drives)
DC Input	24 VDC (Nominal)
DC Output (per drive)	+5V at 3A (USB)  +12V at 4A (PCIe)  +5/12V at 4A (SATA)
Dimensions	9.50 in. (L) x 6.50 in. (W) x 2.63 in. (H)
Weight	44 oz (1,220 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)
EMC Compliance	This product has been evaluated according to the following EMC regulatory standards: <ul style="list-style-type: none"><li>• EN 55032:2015+A1:2020, Class A</li><li>• EN IEC 61000-3-2:2019+A1:2021</li><li>• EN 61000-3-3:2013+A2:2021+AC:2022-01</li><li>• EN 55035:2017+A11:2020</li></ul> This evaluation included ESD immunity testing as per IEC 61000-4-2 standard. For usage precautions and information about interference recovery, see the following <a href="#">Note</a> .
<b>Warranty</b>	
TX2 (main unit)	Three years parts and workmanship from date of purchase
TX2-S1 and TX2 Accessories	One year parts and workmanship from date of purchase



**Note:** OpenText TX2 has passed ESD testing in accordance with the EN 55035 standard, Performance Criterion C.

In environments with strong electrostatic discharge events, this product may exhibit temporary performance degradation, which may require power cycling the unit to restore functionality. This is a temporary condition with no observed or expected product damage.

To restore standard operation after a suspected ESD event, restart the unit. This can be accomplished by pressing and holding the power button for three seconds or by unplugging the power cord, waiting for five seconds, and then plugging the cord back in.

This recovery process is the expected behavior and conforms to the EU EMC regulatory requirements. If any persistent issues are seen with OpenText TX2 after a suspected ESD discharge event, contact OpenText Technical Support.

## 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
- Real-time clock data retention issue

### 6.2.1 Power supply issues

The power supply provided with this device is capable of powering OpenText TX2 and nearly all combinations of drives. OpenText TX2 also employs staggered power sequencing for the source and destination drives, powering up each drive interface in a sequential manner. 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, you may hear the source and destination drives spin up separately.

During power-on initialization and self-test, OpenText TX2 checks its input voltage. If the voltage is below the minimum specification, the unit will power off and the power button LED will blink.

If you have difficulty turning on the device, check the status of the blue DC power LED on the <sup>TP8</sup> power supply connector to ensure that it is on, which indicates that it is connected to power and functioning properly.

## 6.2.2 Thermal issues

OpenText TX2 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 be prompted to immediately shut down the unit to prevent permanent damage. You will have the option to ignore the shutdown 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 the front 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.

In the event of a severe overheating condition that is not covered by the warning conditions described here (for example, a processor or display malfunction that prevents the warnings from being seen), OpenText TX2 has a mechanism by which it will shut down on its own based on low level temperature readings. If this happens to your unit, please contact OpenText Customer Support.

## 6.2.3 Problems with drive detection

When using a product like OpenText TX2, the most common problem you may encounter is a failure to achieve drive detection. Most drive detection problems are the result of worn out cables or faulty/failing drives. The following table lists the most common drive detection problems and possible corrective actions.

Problem	Corrective Action
General drive detection	Check the power and signal cable connections between your unit and the drive to ensure that all connectors are properly seated. If you have a new/different drive cable, use it to see if the problem goes away. Similarly, if you have a different OpenText Forensic Equipment device, try detecting the same drive/cable setup on that device to see if the problem follows the drive/cable or the original unit. If your unit still does not detect the drive, cycle the OpenText TX2 power to attempt to detect the drive during a fresh start-up sequence.
PCIe SSD is not detected	Most OpenText Forensic PCIe adapters have a power LED to indicate that they are receiving power from the host system. Check to make sure that this adapter power LED is on.
IDE drive is not detected	Ensure that the blue end of the IDE signal cable faces the TA7-5 IDE adapter, and that the IDE drive is configured for Master or Single Drive mode. Also, check that the drive power cable is connected directly to the TA7-5 adapter.
SATA drive is not detected	Use only the Unified SATA cables provided by OpenText (TC4-8-R4). 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	Some USB drives have proprietary self-encryption (for example, Kingston's IronKey). These drives expose a small CDFS volume to the host system instead of the main data volume. This CDFS volume typically includes an application that allows for entry of a key/password on the host system. While OpenText TX2 cannot run these unlocking applications, it will detect such a proprietary, self-encrypted drive and report its type to the user in the drive tile.

OpenText TX2 has been tested 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. OpenText issues firmware updates to address most compatibility issues. If a drive is not recognized by your OpenText TX2, go to OpenText My Support <https://support.opentext.com> and ensure you are running the latest firmware. If there are no firmware updates available to resolve your detection issue, contact your reseller or OpenText Customer Support to report your issue.

## 6.2.4 Problems detecting Apple devices in target disk mode

Accessing drives inside Apple devices is accomplished by putting the Apple device into target disk mode (TDM) and then connecting it to the USB source port on your OpenText TX2 using a quality USB adapter cable. Problems with these types of connections are not uncommon. The following table lists the most common issues seen when trying to mount and acquire an Apple computer in target disk mode.

Problem	Corrective Action
No detection on OpenText TX2; Apple computer boots to standard user login screen (not in target disk mode).	<p>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).</p> <p>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).</p> <p>With the advent of the T2 secure enclave interface, it is also possible that the System Preferences on the Apple device are set to block remote booting, which blocks entry into target disk mode. If possible, change the System Preferences to allow remote booting, and then retry the TDM boot sequence.</p>
No detection on OpenText TX2; Apple device is in TDM, but more than the USB logo is shown on the Apple screen.	<p>This condition indicates that the Apple device is not properly connected to a powered-on OpenText TX2 unit. This is most commonly caused by the use of a defective cable between the Apple computer and your unit. For recommended cables, see <a href="#">“Connecting drives” on page 51</a>. 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, or replacing the cable to see if the logo pattern on the Apple screen changes.</p>
<p>No detection on OpenText TX2; Apple device is in TDM with only the USB logo on the screen.</p> <p>When connecting an Apple computer to OpenText TX2, the order of operations matters.</p>	<p>The recommended steps are as follows:</p> <ol style="list-style-type: none"> <li>1. Power-on the Apple computer while pressing the T key until only the USB logo appears on the screen.</li> <li>2. Power on OpenText TX2 and ensure it is fully booted to the <b>Home</b> screen.</li> <li>3. Connect the interface cable to the Apple computer first and then to one of the USB Type-C connectors on the source side of OpenText TX2.</li> </ol>
No detection on OpenText TX2; Apple device is in TDM and all cabling is solid.	<p>As mentioned in <a href="#">“Apple Target Disk Mode acquisition adapters” on page 138</a>, it’s possible that the Apple computer you are attempting to acquire is a version with a T2 (or later) security chip. OpenText TX2 does not currently provide a means of detecting or unlocking drives from a T2 (or later) Apple computer.</p>

Problem	Corrective Action
OpenText TX2 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, OpenText TX2 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 target disk mode, please contact OpenText Customer Support for assistance.

### 6.2.5 Real-time clock data retention issue

Under standard operating conditions, the real-time clock on your OpenText TX2 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 OpenText TX2 for any reason, including battery replacement, as this will void your warranty. If you notice an issue with the time and/or date setting not being retained, contact OpenText Customer Support.



## Chapter 7

# Appendix A: Logical Imaging Engine hash databases

OpenText TX2 allows you to use a file hash database as a logical image search tool, to include or exclude files with matching hashes. The hash database must have a specific format, compatible with the OpenText TX2 Logical Imaging Engine (LIE). This appendix provides the OpenText TX2 hash database format specifications and supporting Python script information.

## 7.1 Hash database format specifications

OpenText TX2 hash databases must be SQLite databases with two tables, the [Header Table](#) and the [Files Table](#).

OpenText TX2 recursively searches the `/otfe_hash_databases/` folder of all accessory and destination drive filesystems for compatible database files. These files will be shown in the hash database selection screen that is displayed during the creation of a logical search. For general information about using this feature, see ["Hash database" on page 94](#).

### Header Table

The header table (header) is formed as:

```
header
CREATE TABLE header (name TEXT NOT NULL, description TEXT NOT NULL, uuid BLOB NOT NULL);
```

The header table must contain the following columns.

<b>name</b>	String name. LIE retains only the first 63 characters of this field.
<b>description</b>	String description. LIE retains only the first 1023 characters of this field.
<b>uuid</b>	The raw 16 bytes of a UUID.  <b>Important</b> This field must be unique. If OpenText TX2 finds duplicate UUIDs between two databases, it rejects both of them.

The header table should have only one row. If more than one row exists, the final one will be assumed to be accurate, as selected by:

```
Getting the correct header row
SELECT name, description, uuid FROM header ORDER BY rowid DESC LIMIT 1;
```



**Note:** rowids are not guaranteed to increase as more rows are added. It is best to ensure that the header only has one row.

## File Table

The file table (`files`) is formed as:

```
files
CREATE TABLE files (size INT NOT NULL, sha1 BLOB NOT NULL, md5 BLOB NOT NULL);
```

The file table must contain at least one of the following columns.

<code>size</code>	The file size, in bytes.
<code>sha1</code>	The sha1 hash, as a binary blob (NOT hexified).
<code>md5</code>	The md5 hash, as a binary blob (NOT hexified)

To maximize performance, it is recommended that your database file contains the `size` and one hash column, preferably the strongest hash available. Multiple hash columns slow down your device when calculating hashes and hinder the search, except for a rare case where the size is not available and one hash happens to have conflicts.

## Recommended Index

If `size` is available, OpenText TX2 verifies that at least one file of that size exists in the database, before it calculates either hash. Therefore, it is strongly recommended that your database has an index on `size` available to it. It is also recommended to have an index on all columns available.

The following example shows an index that achieves both recommendations (because indexes are stored by sorting on the first column first, then the second, then the third).

```
all_index
CREATE INDEX all_index ON files (size, sha1, md5);
```

## Application ID

The `application_id` metadata field should be set to `0x4f544844`, which is `0THD` in ASCII and `1,330,923,588` in base 10.

```
application_id
PRAGMA application_id = 0x4f544844;
```



**Note:** The creation script sets the `application_id`, but OpenText TX2 does not currently check this field.

## User version

The `user_version` metadata field should be set to `1`. This number will be incremented if OpenText changes database schemas.

```
user_version
PRAGMA user_version = 1;
```



**Note:** The creation script sets the `user_version`, but OpenText TX2 does not currently check this field.

## 7.2 Example script 1: create\_othd.py

This Python script creates an OpenText TX2-compatible hash database file from other hash sources, including industry standard and custom formats. The input options and resulting outputs are as follows:

<b>folder</b>	Given a folder on the current system as input, walk the folder recursively and create a database from its contents.
<b>md5_list</b>	Given a text list of hexified MD5 hashes, create a database of those MD5s.
<b>sha1_list</b>	Given a text list of hexified SHA1 hashes, create a database of those SHA1s.
<b>nsrl_rds</b>	Given a NIST NSRL RDS hash set ( <a href="https://www.nist.gov/itl/ssd/software-quality-group/national-software-reference-library-nsrl/nsrl-download/current-rds">https://www.nist.gov/itl/ssd/software-quality-group/national-software-reference-library-nsrl/nsrl-download/current-rds</a> ), create a database from its contents.
<b>nsrl_caid</b>	Given a NIST NSRL CAID hash set ( <a href="https://www.nist.gov/itl/ssd/software-quality-group/national-software-reference-library-nsrl/nsrl-download/non-rds-hash#child-abuse-image-database-caid-sets">https://www.nist.gov/itl/ssd/software-quality-group/national-software-reference-library-nsrl/nsrl-download/non-rds-hash#child-abuse-image-database-caid-sets</a> ), create a database from its contents.
<b>csv</b>	Given a CSV with columns named the same as the OpenText TX2 database format, create a database of its contents.

To download a copy of this script, go to one of the following web pages:

- OpenText My Support create\_othd.zip ([https://support.opentext.com/csm? id=kb\\_article\\_view&sysparm\\_article=KB0858435](https://support.opentext.com/csm? id=kb_article_view&sysparm_article=KB0858435))
- GitHub create\_othd.py ([https://github.com/MicroFocus/OpenText-Forensic-Equipment/blob/main/hash-database/create\\_othd.py](https://github.com/MicroFocus/OpenText-Forensic-Equipment/blob/main/hash-database/create_othd.py))

## 7.3 Example script 2: describe\_othd.py

This Python script describes the contents of a given OpenText TX2 hash database file.

To download a copy of this script, go to one of the following web pages:

- OpenText My Support describe\_othd.zip ([https://support.opentext.com/csm? id=kb\\_article\\_view&sysparm\\_article=KB0858436](https://support.opentext.com/csm? id=kb_article_view&sysparm_article=KB0858436))
- GitHub describe\_othd.py ([https://github.com/MicroFocus/OpenText-Forensic-Equipment/blob/main/hash-database/describe\\_othd.py](https://github.com/MicroFocus/OpenText-Forensic-Equipment/blob/main/hash-database/describe_othd.py))

