

Atola TaskForce Manual

Version: Dec 02 2025

Quickstart

- [Start imaging](#)

Introduction

- [Workflow](#)
- [Update firmware](#)

Unit & extensions

- [Package contents](#)
- [Forensic hardware unit](#)
- [Magnetic cover](#)
- [M.2 SSD extension](#)
- [Apple PCIe SSD extension](#)
- [Thunderbolt extension](#)

Installation & environment setup

- [Extend subscription](#)
- [Connectivity options](#)
- [Multiple user profiles](#)
- [Maximize 10Gb network throughput](#)
- [Use a static IP address](#)
- [Jumbo frames for fast imaging to server](#)
- [Network setup tips](#)
- [Configuring NAS: Synology / QNAP](#)

Working with devices

- [Connecting drives and starting TaskForce](#)
- [Supported drives](#)
- [Drive identification](#)

Interface controls & indicators

- [Main window and controls](#)
- [ATA registers: what they mean](#)

Diagnostics

- [Diagnosing an evidence drive](#)
- [Tracking SMART table changes](#)

Imaging: Basics

- [Imaging to 5 targets](#)
- [Imaging to 2 targets with post-hashing](#)
- [Imaging to an E01 file with dual hash](#)
- [Imaging to E01 on a target drive](#)
- [Imaging to an AFF4 file](#)
- [Imaging a remote iSCSI drive](#)
- [Configure target file templates](#)
- [Imaging to a password-protected server](#)
- [Imaging to a file on an encrypted drive](#)
- [Restore E01, AFF4, RAW image file to drive](#)
- [Clipping target drive to source evidence size](#)
- [Exporting sector lists from imaging sessions](#)
- [Creating, exporting, importing presets](#)

Imaging: Damaged drives & performance

- [Imaging a drive with a damaged head](#)
- [Multipass imaging of damaged drives](#)
- [Imaging freezing drives](#)
- [Imaging a shorted hard drive](#)
- [Express mode: self-launching imaging](#)
- [Logical imaging](#)
- [Imaging only sectors with data](#)
- [Imaging cheat sheet](#)

RAID reassembly and imaging

- [RAID 0 with unknown configuration](#)
- [RAID 5 with unknown configuration](#)
- [RAID 5 with a missing drive](#)
- [RAID 5 imaging with 2 damaged drives](#)
- [RAID 6 with unknown configuration](#)
- [RAID 6 with 2 missing drives](#)
- [RAID 10: reassembly and imaging](#)
- [Synology NAS RAID auto reassembly](#)
- [Instant RAID configuration detection: mdadm](#)
- [Logical imaging of a RAID array](#)
- [Imaging selected partitions of a RAID array](#)
- [Unmounting an assembled RAID](#)
- [RAID types and parameters](#)
- [RAID tags and what they mean](#)
- [RAID cheat sheet](#)

Automation

- [Using Web API in a browser](#)
- [Instantly starting 12 imaging sessions](#)
- [Autostarting image analysis upon completion of imaging](#)

Wipe, hash & special capabilities

- [Multi-launch of wiping, hashing, diagnostics](#)
- [Wiping 18 drives simultaneously](#)
- [Unclip or change HPA, DCO, AMA limitations](#)

- [Calculating hash during imaging](#)
- [Calculating dual hash of an E01 file](#)
- [Segmented hashing](#)
- [Verify segmented hashes](#)

Case management

- [Case management system](#)
- [Add a case](#)
- [Add one device to several cases](#)
- [Add several devices to one case](#)
- [Finding and editing cases](#)
- [Finding reports in a case](#)
- [Printing reports from a case](#)

What else?

- [FAQ](#)
- [Troubleshooting](#)

Quickstart

Atola TaskForce has highly intuitive design. This guide will assist in quickly learning how to use TaskForce and image an evidence drive.

1. Start TaskForce

Switch on the power button on the back side of the system and wait for a message to appear on the small IP screen on the front panel. If TaskForce is connected to the local network with an Ethernet cable, IP addresses will be displayed, which can be entered in a Chrome browser on any device within the same local network. If no Ethernet cable is connected to it, TaskForce will start in Standalone mode, in this case you can use the main screen on top of the device.



TaskForce shows network IP address to connect

2. Plug the source drive into the unit

Atola TaskForce has 18 ports, and each of them can be used as either a source or a target. Before connecting your evidence drive to the appropriate port, make sure the port is in the source mode, thus securing any command that can change the state of the drive from being applied to the evidence drive.

3. Diagnose the source drive

In the left-side menu click the **Diagnose** button and a slide-out Select device menu will appear.

In the appropriate category (SATA, SAS, USB, File, IDE or Extension*) select the source drive and click **START** button. Diagnostics will take a couple of minutes. Should the state of the evidence drive be good, you can proceed to image it.

Atola
TECHNOLOGY

Diagnose
Image
Logical
RAID
Wipe
Other

SAMSUNG HD250HJ S0URJ90P802484
250 GB
SATA 5
Case ID: Not assigned

Creation date: 07/21/2023 10:56 AM
TaskForce serial: 95959595
Location: SATA 5
Device model: SAMSUNG HD250HJ
Size: 250 GB (250,059,350,016 bytes)
Case investigator:

Software version: 2023.4.1
TaskForce IP: 10.0.0.41
Write protection: Off
Device serial: S0URJ90P802484
Device firmware: FH10005
Case description:

Diagnostics completed

Results
No major hardware or firmware issues found
Estimated imaging time: 1 hour 4 minutes

1. Circuit board
Device is powered on. A power cycle is needed...
Powering down the device...
Applying power and watching spin-up currents...

Current oscillogram (12V):

Current oscillogram (5V):

Device has become ready in 5 sec
Registers: Status: 01010000 Error: 00000001
Peak power consumption during spin-up: 5V line = 793.03 mA; 12V line = 1,649.091 mA
Integrity word of device identification data is OK.
Device identified: SAMSUNG HD250HJ SN: S0URJ90P802484
Circuit board check passed
Elapsed: 18.4 sec.

PRINT **GO TO CASE** < >

Diagnostics proves the drive is fine

4. Plug the target drive(s) into the unit

When connecting a target drive (or multiple targets), also make sure that the corresponding ports are set to target mode.

5. Start imaging

In the left-side menu click the **Image** button.

Select the source device you have diagnosed in the menu.

Select the targets in the menu. Once you have selected the source, a right-hand slide-out Select target device menu will appear. Under the same categories you can find and select a single target or multiple targets.

The Imaging initiation page lists the source drive and the case ID at the top of the page, the default settings applied to this imaging session and the list of Targets to be involved in this session. Click Start button at the bottom of the page.




Image acquisition in progress

TaskForce workflow

Atola TaskForce provides a complete feature set for a forensically sound evidence acquisition process. Based on our own decade-long experience of working with data storage devices as well as the experience of our clients in digital forensics market we strongly recommend this workflow:

1. Diagnose the drive

TaskForce is equipped with a fully-automated diagnostics module, which diagnoses all drive systems: printed circuit board (PCB), spindle motor, head stack, firmware, and file systems. Diagnostics will work properly even if the drive has burnt parts or damaged head stack – the routine makes use of the current monitor that is embedded into TaskForce unit.

After diagnostics finishes, the tool prepares a report and lets you know the exact issue with the drive; it also suggests the next step to be able to retrieve the data.

To ensure efficient imaging of both good and damaged drives, TaskForce is equipped with a sophisticated and powerful imaging module that creates a bit-to-bit copy of the evidence. Based on the diagnostics report, image drives with default settings or adjust them, should the media be damaged and require special treatment.

1

CASES

REPORTS

Atola
TECHNOLOGY

DEVICES

Diagnose

Image

Logical

RAID

Wipe

Other

ST9500423AS S2V0EP3P

500 GB

Has issues

SATA 3

Case ID:

Not assigned

1 target

Imaging data...

99%

Pass: 3 of 5


Speed: 89 MB/s

Errors: 59

Sectors imaged: 976,224,496

Sectors scheduled: 976,773,168

Estimated time: 3 s



Log

Target HEX viewer

Signatures : 125892

Date	Message
2023/07/20 16:54:27	Jump to sector 3,002,464
2023/07/20 16:54:31	Can't read sectors 3,997,312 - 4,001,407 (UNC)
2023/07/20 16:54:31	Jump to sector 4,001,408
2023/07/20 16:54:34	Can't read sectors 4,996,256 - 5,000,351 (UNC)
2023/07/20 16:54:34	Jump to sector 5,000,352
2023/07/20 16:54:37	Can't read sectors 5,999,296 - 6,003,391 (UNC)
2023/07/20 16:54:37	Jump to sector 6,003,392

PAUSE


VIEW SETTINGS

Link: PHY

Status: ERR INDX CORR DREQ DREQ TAULT DREQ BUSY

Error: AMAN TONR ABRT - IDNF - UNC ACRC

Imaging damaged drive.



Imaging good SSD. Speed: 550 MB/s.

4. Calculate hash

To ensure forensically sound evidence acquisition process, remember to [calculate hash](#) of the evidence and the image. It is essential way to prove image integrity.

With damaged devices, it is best to calculate hash during imaging (using [segmented hashing](#)). This way data on a fragile device is only read once, and less potential damage to the media is caused.


Linear hash can only be calculated by reading data in sectors consecutively in one pass. When it encounters a bad sector, linear hash calculation is discontinued.

Updating TaskForce firmware

Atola TaskForce firmware is updated on regular basis by our team. You can keep track of the updates we make to the firmware in [TaskForce changelog](#).


Updating TaskForce firmware is easy using a remotely connected computer.

1. Plug TaskForce into to your local Ethernet network.
2. Open the Chrome browser on your PC.
3. Download the [most recent version](#) of the firmware.
4. Enter TaskForce IP address in the Chrome browser.
5. Go to the **Menu** on the top right.
6. Click **Update firmware**.




The Update firmware command in the TaskForce menu.

7. Choose update method and then click **Select file**.




Choosing update method on the Update firmware page.

8. In the file selector, select the firmware file and then click **Open**.



Selecting the firmware file.


9. In the Update firmware dialog, click **Update**.



The Update button in the Update firmware dialog.

Once the update process is completed, TaskForce software switches to the new version. No TaskForce reboot or Chrome restart is required.

To check the current firmware version, go to **Menu > Update firmware**.



Check current firmware version.

Package contents (TaskForce)

The TaskForce package includes the following items:

[TaskForce hardware unit](#)



Power supply




14x SAS/SATA powered cables



Hitachi password extraction adapter



3.5" to 2.5" IDE adapter




IDE power cable




IDE interface cable

Flash card reader



2x Ethernet cat 7 cable





Microfiber cloth



1-page installation instruction

Atola TaskForce Installation & Activation

1. Turn on Power switch on TaskForce's back panel.
2. Power on TaskForce's touchscreen by pressing its own Power button.
3. Enter the Windows password in the touchscreen: **atola** to unlock it.
4. Wait for 2 minutes while TaskForce is booting.
5. Follow instructions on the screen to activate TaskForce.
6. Download the latest firmware version: atola.com
7. Update TaskForce firmware by using the downloaded file: Menu > Update firmware
8. atola.com guide is available at the webpage: atola.com/TF-quickstart

Useful Tips

For remote access and task management:

1. Connect TaskForce to the local network by plugging Ethernet cable into either ETH1 or ETH2 port.
2. Enter IP address displayed on the IP screen in a Chrome browser on another device within the same Ethernet network.

For faster imaging, increase data transfer rate by connecting TaskForce to 10Gb Ethernet network.

All 18 ports can be used in either Source or Target mode. Each is equipped with a Source switch and light indicator. When in Source mode:

- Light indicator is lit
- Port's write protection is enabled
- Device plugged into the port can only be selected as a source for imaging tasks and cannot be selected for wiping tasks

To power down TaskForce, turn Power switch off. The touchscreen will automatically turn off.

Forensic hardware unit

To support multiple simultaneous imaging sessions and other forensic tasks running on its 18 ports at top speeds, TaskForce hardware unit is built on extremely robust, high-capacity components that include a server-grade motherboard, 8-thread Xeon CPU 3.7 GHz and ECC RAM. To ensure high quality and efficiency of our tools, [we test them on hundreds of storage devices](#).

This forensic hardware unit is designed for various types of digital investigations. It can be used in the lab and in standalone mode.

Ports, indicators & switches

Ports:



- 6 SATA
- 6 SATA/SAS
- 4 USB
- IDE
- Extension slot (for Thunderbolt, Apple PCIe SSD and M.2 SSD [extension modules](#))



Source/target switch on each of 18 ports enabling hardware write protection in source mode:

- Source mode: helps safely connect and with the examined evidence drives to TaskForce
- Flexible port configuration: any TaskForce port can serve you as source or target depending on your needs.

LED indicators: two indicators for each of 18 ports. One indicator is located next to the port's source/target switch and lights up when source mode is enabled.

Two 10Gb Ethernet ports

- In the lab: handles 12+ multiple imaging sessions utilizing two 10Gb network connections
- In the field: works fast with modern portable NAS solutions

Other ports: Serial RS232, IDE power, VGA, RS-232, DC IN 24V

Supported hard drive interfaces: SATA I/II/III, SAS3 (12Gb), USB 2.0/3.0, IDE

Power switch



Flash card support via card reader attached to any USB port of TaskForce.

Displays

Main display: Microsoft Surface Pro 4

- Screen: 12.3" PixelSense Display
- Resolution: 2736 x 1824 (267 PPI)
- Touch: 10 point multi-touch
- Detachable from TaskForce if necessary

IP display: OLED (20x2 characters)



Physical / Environmental

- Dimensions (including screen): 12x10x3 inches (30x25x7,5 cm)
- Weight (including screen): 4.5 lbs (3.6 kg)
- TaskForce working temperature range: 0°C – 40°C (32°F – 104°F)
- Screen working temperature range: 0° – 35°C (32° – 95°F)
- Power consumption: 60 Watt average, 280 Watt peak
- Supply Voltage: 100 - 240 VAC, 50 - 60 Hz

Other specs

- Internal OS: Linux running a custom kernel
- Control interface: web-based TaskForce application
- Removable NVMe SSD for internal case management database
- Optional [WiFi 802.11n 150 Mb/s adapter](#) in access point mode for easy connectivity (purchased separately)



[View TaskForce in 3D](#)

Inside TaskForce forensic hardware unit



TaskForce forensic hardware unit is essentially a small server-grade computer running Linux. But because neither BIOS, nor Linux kernel was designed to handle hard disk failures, Atola engineers have invested a significant amount of research and development efforts to build a highly customized and fine-tuned Linux kernel that fully overcomes these issues and handle damaged media properly. Additionally, this kernel features:

- High-speed DMA data transfers, 500+ MB/s
- Full low-level control over SATA, USB and IDE ports
- Full native SATA support
- Reset and SATA PHY control for best handling of severely damaged hard drives
- All BIOS and standard kernel functions are disabled

TaskForce hardware also features Atola's proprietary circuitry for ultimate drive's power control:

- Current sensor for in-depth hard disk diagnosis
- Automatic overcurrent and short-circuit protection
- Overvoltage protection

These features are a must for proper handling of damaged drives.

For instance, low-level control of the SATA, SAS, USB and IDE ports allows TaskForce to handle the devices that do not properly initialize, have many bad sectors, or frequently freeze due to internal (mechanical) failures. SATA PHY control allows resetting a frozen hard drive without a power cycle, thus saving time during imaging, and reducing the chances of further hard disk degradation and failure. Current sensing allows TaskForce to diagnose a failed drive even if it has electronic or mechanical damage.



Overcurrent protection detects when the drive draws abnormal current and stops it to prevent any further damage. Overvoltage protection circuit ensures that in the unlikely event of the forensic hardware malfunction, the attached drives are not damaged in any way.

TaskForce forensic hardware unit is fully controlled by the software via Chrome browser, therefore no Linux experience is required to operate it.

Magnetic screen cover

The cover [protects the TaskForce screen](#) from physical damage during transportation and allows for quick and easy packing of the unit.

Made of 2-mm thick aluminum, it is sturdy to protect the screen from scratching and breaking, yet is lightweight and thin to add no additional bulk to the device.

5 magnetic screws on TaskForce's upper panel help attach the lid quickly and easily. It is as easily removed with a measured pulling motion.

The cover comes in the standard package with TaskForces assembled from September 2020 onwards. All customers with an active TaskForce subscription can receive the complimentary cover by contacting our support and providing their unit's serial number and shipping address.

The only magnetic parts used are the magnetic screws. While our tests have shown that they are not able to damage a drive, we recommend that no storage or other devices are placed near the magnets.

Magnetic screen cover

The cover is made of 2-mm lightweight aluminum to protect the screen from physical damage during transportation and allows for instant packing of the unit.

To start using the cover, screw 5 magnetic screws into TaskForce's upper panel as indicated below:



[Attach the lid](#) by placing the cover against the magnets and remove with a careful pulling motion.

The only magnetic parts used are the magnetic screws. Do not place drives on top of the screws. While our tests have shown that they are not able to damage a drive, we recommend that no storage or other devices are placed near the magnets.


The cover comes in the standard package with TaskForces assembled from September 2020 onwards. All customers with an active TaskForce subscription can receive the complimentary cover by contacting our support and providing their unit's serial number and shipping address.

M.2 SSD extension module



You can connect M.2 PCIe NVMe, M.2 PCIe AHCI, and M.2 SATA solid state drives to Atola TaskForce using the M.2 SSD extension module.

This extension module supports only B & M key and M key interface drives.



M.2 SSD extension module works with Atola DiskSense 2 as well.

TaskForce features supported for the M.2 SSD extension

All standard TaskForce operations and features are supported for the M.2 PCIe NVMe, M.2 PCIe AHCI, and M.2 SATA drives, including:

- Max read/write speed: 1300 MB/s
- Drive hotplug (excluding M.2 PCIe AHCI drives)
- Write protection with the physical source/target switch
- Diagnostics
- Physical and logical imaging
- Hash calculation
- Damaged drive support

Plug and unplug the M.2 SSD extension module

TaskForce hardware unit is equipped with the PCI Express port on its back panel, which is referred to as an Extension port. It is used to plug [Atola hardware extension modules](#) supported by Atola TaskForce software.



Important: Do not plug or unplug the extension module when TaskForce is powered on, this can damage the extension module or TaskForce hardware unit. However, you can hotplug M.2 PCIe NVMe and M.2 SATA drives, when the extension module is already connected and fixed in place.

To connect the M.2 SSD extension module to TaskForce, do the following:

1. Power off TaskForce.
2. Align 3 holes on the extension module and 3 screw holes on the Taskforce back panel. Firmly plug the M.2 SSD extension module into the Extension port and fasten the module with 3 screws.
3. Plug an M.2 PCIe NVMe, M.2 PCIe AHCI, or M.2 SATA drive into the extension and fasten the drive in place with the black plastic slider.
4. **Optional:** If you plan to work with M.2 SATA drives, connect the extension to one of the SATA ports on TaskForce using a SATA cable. For details, see [Connect and identify an M.2 SATA drive](#) section on this page.
5. Power on TaskForce.



To disconnect the M.2 SSD extension module or replace it with another extension module, do the following steps:

1. Power off TaskForce.
2. Disconnect the SATA cable from the extension module.
3. Release the screws, which hold the module, and unplug it from the Extension port.
4. **Optional:** Plug another extension module into the Extension port and fasten the module with a screw.
5. Power on TaskForce.

Distinguish between M.2 NVMe/M.2 AHCI, and M.2 SATA drives

All solid state drives with M.2 form factor look pretty much the same.

But there are different ways to connect M.2 PCIe NVMe and M.2 PCIe AHCI drives on one hand, and M.2 SATA drives on the other hand to TaskForce using the M.2 SSD extension module.

The easiest way to tell apart M.2 NVMe and M.2 PCIe AHCI drives from M.2 SATA drives – when the extension is connected and TaskForce is powered on – is to look at the LED indicator on the extension module, next to the words *"When LED is on, plug SATA source port cable here"*:

- If the LED is off, the drive connected to the M.2 SSD extension module is either M.2 NVMe or M.2 PCIe AHCI. See [Connect and identify an M.2 NVMe or M.2 AHCI drive](#).
- If the LED is on, the drive is M.2 SATA. See [Connect and identify an M.2 SATA drive](#).


To define the type of the particular M.2 drive before connecting it to TaskForce, check the markings on the drive or refer to the manufacturer's specifications.



Connect and identify an M.2 NVMe or M.2 AHCI drive

To identify the M.2 PCIe NVMe or M.2 PCIe AHCI drive, TaskForce must be first booted with that particular type of drive connected to the M.2 SSD extension module:

1. Power off TaskForce.
2. Plug an M.2 PCIe NVMe or M.2 PCIe AHCI drive into the extension and fasten the drive with the black plastic slider.
3. Power on TaskForce and wait for the booting to be completed.
4. In the TaskForce main window, click **Devices**.
5. Click the port in the **EXT** section of the **Select device** panel.



NVMe M.2 drive in the EXT section of the Select device panel.

Work consecutively with several M.2 NVMe drives

Drive hotplug is supported for M.2 PCIe NVMe drives, but not for M.2 PCIe AHCI drives. You can plug several M.2 NVMe drives one after another, without turning TaskForce off and on.

To identify the M.2 PCIe NVMe drive, TaskForce must be first booted with that particular type of drive connected to the M.2 SSD extension module. After you boot TaskForce with an M.2 NVMe drive connected to the extension module, the system identifies all other M.2 NVMe drives without rebooting the hardware unit.

To define the type of the particular drive, check the markings on it or refer to the manufacturer's specifications.

The hotplug for M.2 PCIe NVMe drives was added in TaskForce firmware update 2021.4.

If you have an earlier firmware version, to replace an NVMe drive, TaskForce needs to be powered off:

1. Power off TaskForce.
2. Remove the power cable and wait for 1 minute.
3. Unplug M.2 extension with the M.2 drive installed and carefully plug another M.2 drive.
4. Power on TaskForce.

Connect and identify an M.2 SATA drive

When you plug an M.2 SATA drive into the M.2 SSD extension module, TaskForce lights up the LED indicator on the module, next to the words "*When LED is on, plug SATA source port cable here*".

M.2 SATA drives appear in the **SATA** section in the **Select device** panel, not in the **EXT** section, as M.2 NVMe or M.2 AHCI drives do. For example, if you connect M.2 SSD extension with the SATA cable to SATA 1 port on the TaskForce hardware unit, the drive will appear in the corresponding SATA 1 port in the TaskForce software.

To connect an M.2 SATA drive to TaskForce, do the following steps:

1. Using the SATA cable, connect the M.2 SSD extension module to one of the SATA ports on TaskForce.
2. In Atola TaskForce software, power off the SATA port.
3. Plug an M.2 SATA drive into the extension and fasten it in place with the black plastic slider.
4. In Atola TaskForce software, power on the SATA port.



Work consecutively with several M.2 SATA drives

For M.2 SATA, drive hotplug is supported. It allows connecting and replacing M.2 SATA drives without turning TaskForce off and on again.

To determine if a particular M.2 drive is SATA type or not, check the markings on it or refer to the manufacturer's specifications.

Also, if your drive belongs to the SATA type, TaskForce will light up the LED indicator on the extension module, next to the words *"When LED is on, plug SATA source port cable here"*.

Connect a U.2 NVMe drive using adapter


To use a drive with U.2 interface, attach the drive to TaskForce with the help of a U.2-to-M.2 adapter and a cable (not included in the package).

To connect a U.2 drive to TaskForce, do the following:

1. Plug the U.2 drive into the U.2-to-M.2 adapter using the cable.
2. Plug the U.2-to-M.2 adapter into the Atola M.2 SSD extension.
3. Plug the extension into TaskForce while the unit is powered off.

U.2 SSD adapters are available to combine them with the Atola M.2 SSD extension. Please contact your Atola dealer for more information.

Apple PCIe SSD extension module



Apple PCIe SSD extension lets you connect TaskForce to the PCIe SSDs with the custom proprietary M.2 interface within Apple laptops:

- MacBook Pro, Late 2013-2015
- MacBook Air, 2013-2015

TaskForce Apple PCIe SSD extension module works with Atola DiskSense 2 as well.

TaskForce features supported for Apple PCIe SSD extension

The following TaskForce operations and features are supported for the Apple drives:

- Write protection with the physical source/target switch
- Diagnostics
- Physical and logical imaging
- Hash calculation
- Damaged drive support

Plug and unplug the Apple PCIe SSD extension module

TaskForce hardware unit is equipped with the PCI Express port on its back panel, which is referred to as Extension port. It is used to plug [Atola hardware extension modules](#) supported by Atola TaskForce software.



Important: Do not plug or unplug the extension module when TaskForce is turned on, it can damage the extension module or TaskForce hardware unit.

To connect the Apple PCIe SSD extension module to TaskForce, do the following:

1. Power off TaskForce.
2. Align the screw on the extension module and the top screw hole on the Taskforce back panel. Plug the Apple PCIe SSD extension module into the Extension port and fix the module with a screw.
3. Plug an Apple PCIe SSD drive into the extension and fasten the drive with the black plastic latch.
4. Power on TaskForce.

Important: Drive hotplug is not supported yet. TaskForce must be powered off before installing or replacing Apple PCIe SSDs.



To disconnect the Apple PCIe SSD extension module or replace it with another extension module, do the following steps:

1. Power off TaskForce.
2. Release the screw, which holds the module, and unplug it from the Extension port.
3. **Optional:** Plug another extension module into the Extension port and fix the module with a screw.
4. Power on TaskForce.

Thunderbolt extension module



Thunderbolt extension enables TaskForce to work on MacBooks with the following interfaces:

1. FireWire
2. Thunderbolt 2
3. Thunderbolt 3, 2016-2017 models

No SSD removal is necessary, the extension allows connecting TaskForce directly to a MacBook.

The extension module comes with:

- FireWire cable (comes in white or black color)
- Thunderbolt 2 to FireWire adapter (by Apple)
- Thunderbolt 3 to Thunderbolt 2 adapter (by Apple)



TaskForce Thunderbolt extension module works with Atola DiskSense 2 as well.

TaskForce features supported for Thunderbolt extension

TaskForce supports the following operations and features on MacBooks when connected through Thunderbolt extension:

- Write protection with the physical source/target switch
- Physical and logical imaging
- Hash calculation
- Damaged drive support

Plug and unplug the Thunderbolt extension module

TaskForce hardware unit is equipped with the PCI Express port on its back panel, which is referred to as an Extension port. It is used to plug [Atola hardware extension modules](#) supported by Atola TaskForce software.



Important: Do not plug or unplug the extension module when TaskForce is turned on, this can damage the extension module or TaskForce hardware unit.

To **connect** the Thunderbolt extension module to TaskForce, do the following:

1. Power off TaskForce.
2. Align the screw on the extension module and the top screw hole on the Taskforce back panel. Plug the Thunderbolt extension module into the Extension port and fasten the module with a screw.
3. Power on TaskForce.

To **disconnect** the Thunderbolt extension module or replace it with another extension module, do the following steps:

1. Power off TaskForce.
2. Release the screw, which holds the module, and unplug it from the Extension port.
3. **Optional:** Plug another extension module into the Extension port and fix the module with a screw.
4. Power on TaskForce.

Connect MacBook using Thunderbolt extension module

First, write down or take a photo of the serial number located on the bottom side of the MacBook. It will be needed later.




Then do the following steps:

1. Turn off both the MacBook and TaskForce.
2. Plug the Thunderbolt extension module into the Extension port and fasten the module with a screw.
3. Connect the MacBook to TaskForce unit with the help of the Thunderbolt extension and the FireWire cable. Use adapters to connect to the MacBooks with Thunderbolt 2 or Thunderbolt 3 interface.
4. Boot the MacBook in the Target Disk Mode. To do that, start it up while holding down the T key. You should see a Firewire or Thunderbolt icon displayed on screen, signifying that Target Disk Mode is detected and working.




5. Power on TaskForce and wait for the booting to be completed.
6. In the TaskForce main window, open the **Device** panel.
7. Click the port in the **EXT** section of the **Device** panel.




MacBook device in the EXT section of the Device panel.

8. If it is the first time this MacBook is identified by TaskForce, in the **Enter MacBook serial** window, enter the serial number located on the bottom side of the MacBook, and then click **OK**.



Enter MacBook's serial number.

During the subsequent identifications of the MacBook connected to TaskForce, its serial number can be selected from the drop-down menu in the **Enter MacBook serial** window. TaskForce will look up its case management system and offer the choice of MacBooks with the same drive size.



Select the MacBook with the serial number you have connected to TaskForce.

Extend subscription

Atola TaskForce comes with a complimentary 1-year subscription. It covers regular software updates, includes training and technical support from our in-house team of developers, and secures a [lifetime warranty](#).

Buy subscription

To extend your subscription for another period, you first need to place a purchase order [with the reseller](#) that sold you the unit or [on the Atola website](#).

If your subscription has not yet expired, you can still purchase and activate a new one. The new subscription period will commence the day following the current subscription's expiration date.

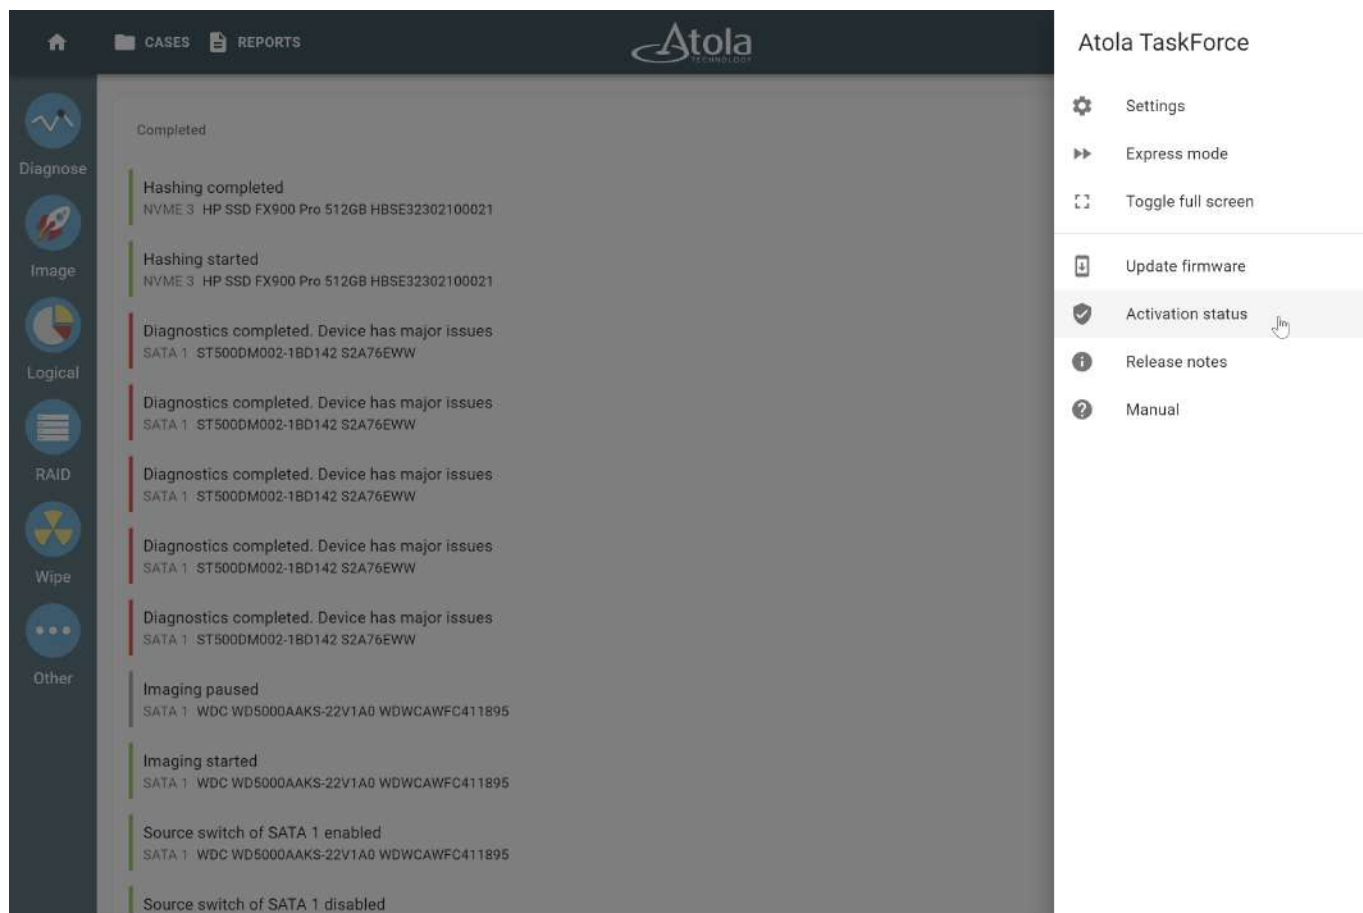
Extend subscription

You will need:

- A device with an internet connection (PC or mobile).
- The serial number located on the bottom side of the TaskForce unit.


After you have purchased a subscription, do the following:

1. On a device with an internet connection, open the TaskForce user interface in your web browser.
2. At the top right, click **Menu > Activation status**.



The Activation status option in the TaskForce Menu.

3. Click Extend subscription.



The Extend subscription button on the Activation status page.

4. Enter the **serial number** of the TaskForce unit.
5. Select **I already extended my subscription through a distributor**, and then click **Next**.

The screenshot shows the 'Extend subscription' page in the Atola Technology web interface. The top navigation bar includes a home icon, 'CASES', 'REPORTS', the Atola Technology logo, 'DEVICES', and a menu icon. A left sidebar contains icons for 'Diagnose', 'Image', 'Logical', 'RAID', 'Wipe', and 'Other'. The main content area is titled 'Extend subscription' and features a two-step process. Step 1, 'Serial number', includes a text box for the serial number (containing '3' and '8') and two radio button options: 'I already extended my subscription through a distributor' (selected) and 'I want to extend subscription by key'. A 'NEXT' button is below these options. Step 2, 'Activation code', is partially visible at the bottom.

The Extend subscription page.

6. Visit the Atola licensing page: a.atola.com or scan QR code on the screen.
7. On the Atola licensing page, enter the license key shown in TaskForce web interface.
8. Fill out all other fields on the Atola licensing webpage, including:
 - o Customer name
 - o E-mail address for activation code
 - o Telephone
 - o Organization, country and city

Atola Licensing

Software Activation Request

License key

Customer name

E-mail address for activation code

Telephone

Organization


Country

City

Note: all fields are required.

Submit

9. Click **Submit**. The Atola licensing webpage generates an activation code.
10. Go back to TaskForce user interface and enter the **activation code** in the respective field.



Entering the activation code on the Extend subscription page.

11. Click **Activate**. TaskForce confirms that reactivation has been successfully completed.

Connectivity and multi-user access

Atola TaskForce has three connectivity options:

1. 10Gb Ethernet network
2. Standalone mode
3. WiFi access point (optional)

10Gb Ethernet network

TaskForce is equipped with two 10Gb Ethernet ports. Whenever the system is connected to a local network via one of its Ethernet ports, an IP address will be displayed on the IP screen on the front panel of the system.

Supported network features:

- Two 10Gb Ethernet ports
- Dynamic (DHCP) IP / Static IP address
- Settings: DNS nameserver, Default gateway, Jumbo frames
- Secure connection with password-protected network folders
- Support of HTTPS with external organizational and self-signed certificates




IP address of Atola TaskForce

If the system is connected via both Ethernet ports, two IP addresses will be displayed on the screen. These IP addresses are assigned to TaskForce by your DHCP server.

Multi-user access

With the help of these IP addresses, TaskForce can be operated by multiple users simultaneously from their workstations or mobile devices:

- enter either of the IP addresses as shown on the IP screen in Chrome browser on another device within the same local network.



Enter IP address in Chrome browser

Through Chrome browser one can remotely track and manage tasks, power devices on and off, open, edit and print cases etc.

Types of devices that can be used to access TaskForce simultaneously include:

- Desktop PC
- Laptop
- Tablet
- Smartphone
- Built-in TaskForce touch screen

TaskForce software can be open in Chrome browser within any OS.

This functionality enables a group of users to work on different assignments using the same tool. This helps utilize TaskForce's multitasking capabilities to the maximum and track operation progress remotely. The number of users accessing TaskForce simultaneously is unlimited.

Standalone mode

TaskForce is equipped with highly responsive HD screen (see [Hardware specs](#)), which allows the system's use in standalone mode.

Whenever the system is not connected to a network via its 10Gb Ethernet ports, *Standalone mode* status will be displayed on the IP screen.

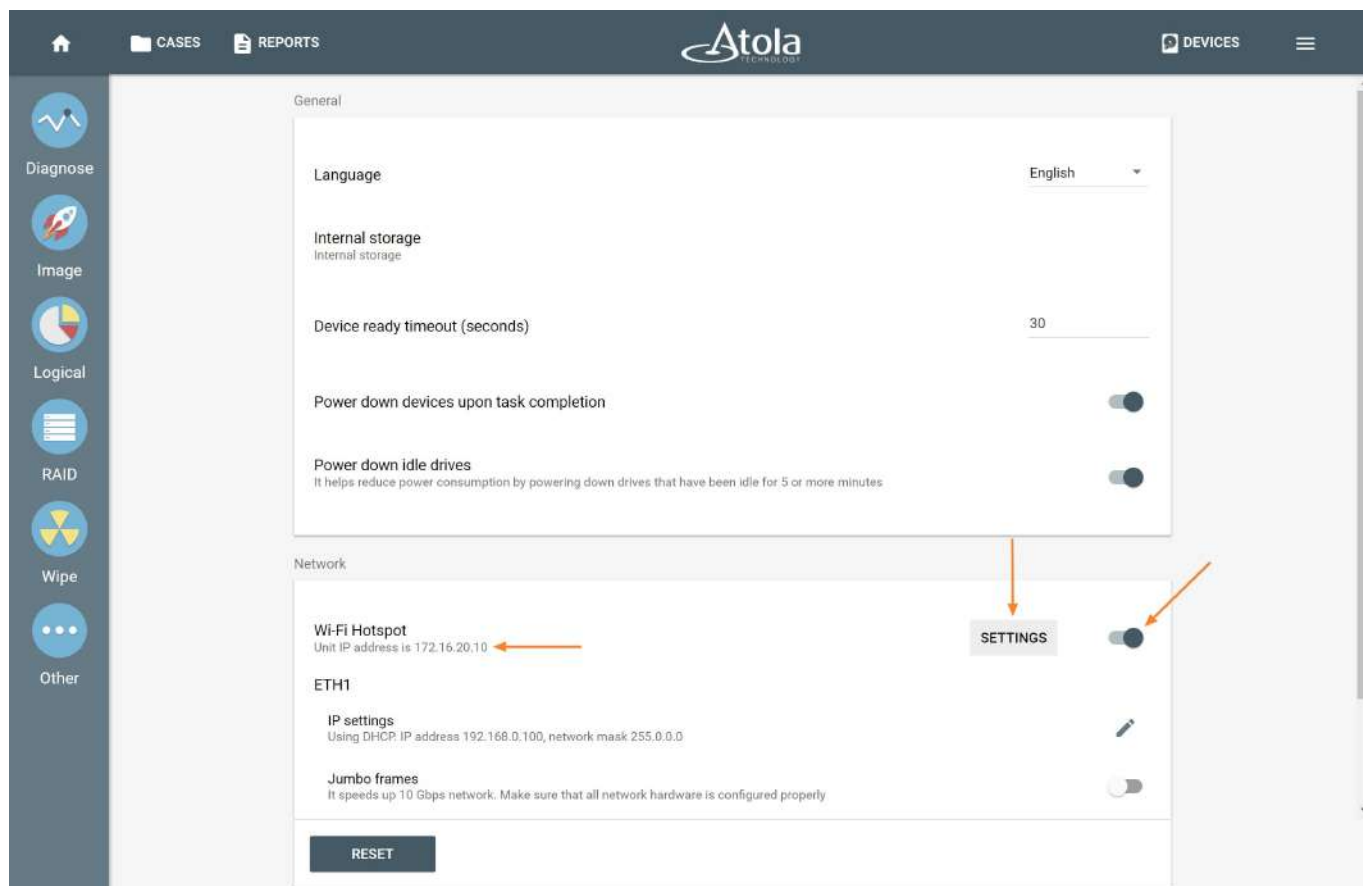
In this mode, you can image data from multiple source drives to target drives in parallel. Thanks to its compact size TaskForce can be easily used in the field.



Wi-Fi access point


The third way to access TaskForce's user interface is via Wi-Fi 802.11n 150 Mb/s adapter. The adapter is optional and can be purchased from your preferred retailer. Here is a link to the [supported adapter](#) sold on Amazon. To enable the adapter, follow these steps:

1. Connect the adapter to one of your TaskForce's USB ports.
2. Go to **Menu** in the top right corner of the TaskForce web page.
3. Click **Settings**.
4. Enable **Wi-Fi Hotspot**. An IP address will appear underneath the **Wi-Fi Hotspot** category.
5. Click **Settings**.



Enable Wi-Fi Hotspot

5. Enable **Wi-Fi Hotspot**.
6. Set **SSID** and **Password**. To make the network invisible to other devices, select **Hidden mode**. Click **Save**.
7. Use these details to connect to the Hotspot from another device.



Adjust Wi-Fi settings

8. To open the TaskForce interface, enter the IP address (indicated under **Wi-Fi Hotspot** category in the **Settings** window, see Step 3) in the Chrome browser of the device you have connected to the Hotspot.

The vast connectivity options make TaskForce a great tool for using both in the lab and in the field.

Use TaskForce with multiple user profiles

For security reasons, you can allow access to Atola TaskForce only for authorized users, protect each user profile with a password, and keep processes, reports, and cases separate and confidential for each user.

Also, you can set up Atola TaskForce to automatically lock its screen after a certain time of inactivity and to prompt a user to log in.

Setting up user roles and profiles is an optimal solution when you want to:

- Share one Atola TaskForce unit with multiple colleagues at the same time.
- Prevent other users from interrupting the processes you started on your TaskForce.
- Prevent your colleagues from seeing each other's cases.
- Temporarily pass your TaskForce to another forensic lab, department, or agency without giving access to your cases.

What others can see and do when you share one TaskForce


After you enable User management, you become an administrator and can create up to 20 profiles with either *Admin* or *User* roles. These roles have different permissions.

What an Admin can see and do

The first profile that you create after enabling User management has an Admin role. You can have more than one profile with the Admin role on your TaskForce.

As an *Admin*, you can:

- See and edit all cases of other users.
- See and stop all processes run by other users.
- See and print all reports created by other users.
- Access and change all TaskForce settings.
- Enable or disable User management.
- Add, edit, or delete other user profiles with either *Admin* or *User* role.
- Change other user's password.
- Set TaskForce to automatically lock the screen after a certain time of inactivity.



Admin sees all processes and cases of other users.

What a User can see and do


A profile with a *User* role is created by *Admin* and has limited access to TaskForce.

As a *User*, you can:


- Log in only to your user profile.
- See and edit only your cases.
- See and print only your reports.
- See, run, and stop only your processes.
- Change only your password.

As a *User*, you can't:

- Access somebody else's user profile.
- See and edit other users' cases.
- See and print other users' reports.
- See or interrupt other users' processes.
- Access a device that is being used by another user
- Access and change TaskForce settings.
- Enable or disable User management.
- Add, edit, or delete other user profiles.
- Change other user's password.
- Disable or configure Automatically lock screen feature set by Admin.



Profiles with a User role see only their cases and processes.



Profiles with a User role don't see each other's cases and processes.

Enable User management


The first profile, what you create after enabling User management, has an *Admin* role.

To enable the multi-user mode in Atola TaskForce, do the following steps:

1. In the Atola TaskForce window, go to **Menu > Settings**.
2. In the **Users** section, toggle **User management**.
3. Enter the username and password for the administrator's profile, and then click **Create**.

Now you can add, edit, or delete other user profiles.

After you enable User management, Express mode is disabled due to security reasons.



User management toggle in the Users section of the Settings menu.


Add, edit, or delete users

Only *Admin* can add, edit, or delete other user profiles with either *Admin* or *User* role.

Add a user

To add a user, do the next steps:

1. Log in to a profile that has the *Admin* role.
2. In the Atola TaskForce window, go to **Menu > Settings**.
3. In the **Users** section, click **Manage**.
4. On the **Users** page, click **Create user**.
5. Enter the username and password for a user.
Optional: To grant this user an *Admin* role, select **Admin** checkbox.
6. Click **Create**.




The Create user button and dialog.

Edit a user

To edit a user, do the following:

1. Log in to a profile that has the *Admin* role.
2. In the Atola TaskForce window, go to **Menu > Settings**.
3. In the **Users** section, click **Manage**.
4. On the **Users** page, select a user you want to edit.
5. In the **Edit user** dialog, edit the username.
Optional: To grant this user an *Admin* role, select **Admin** checkbox.
6. Click **Save**.




The Edit user dialog.

Delete a user

After deleting a user, you can still access their cases and reports under a profile with the *Admin* role.

To delete a user, do the following steps:

1. Log in to a profile that has the *Admin* role.
2. In the Atola TaskForce window, go to **Menu > Settings**.
3. In the **Users** section, click **Manage**.
4. On the **Users** page, select a user you want to delete.
5. In the **Edit user** dialog, click the **Delete** icon.
6. In the confirmation dialog, enter *YES*, and then click **Delete**.




The Delete icon in the Edit user dialog.

Log in and log out

To log out, click **Menu > Log out**. Atola TaskForce locks itself and prompts you to enter a username and password. All the processes you started before logging out are still running in the background.

To log in, enter your username and password, and then click **Log in**. Atola TaskForce unlocks itself and shows its Home screen.




Atola TaskForce is locked and prompts you to enter a username and password.

Change the password

As a *User*, you can change only your password. As an *Admin*, you can change other users' passwords as well.

Change the password for your profile

To change the password for your profile, go to **Menu > Change your password**, enter your current and new passwords and click **Save**.




Change your password command in the Menu.

Change the password for another user as an administrator

To change the password for another user, do the following:

1. Log in to a profile that has the *Admin* role.
2. In the Atola TaskForce window, go to **Menu > Settings**.
3. In the **Users** section, click **Manage**.
4. On the **Users** page, select a user you want to change the password for.
5. In the **Edit user** dialog, enter a new password, confirm the password, and then click **Save**.

If you forgot or don't know the password for an administrator profile on your TaskForce, contact [Atola Support](#).

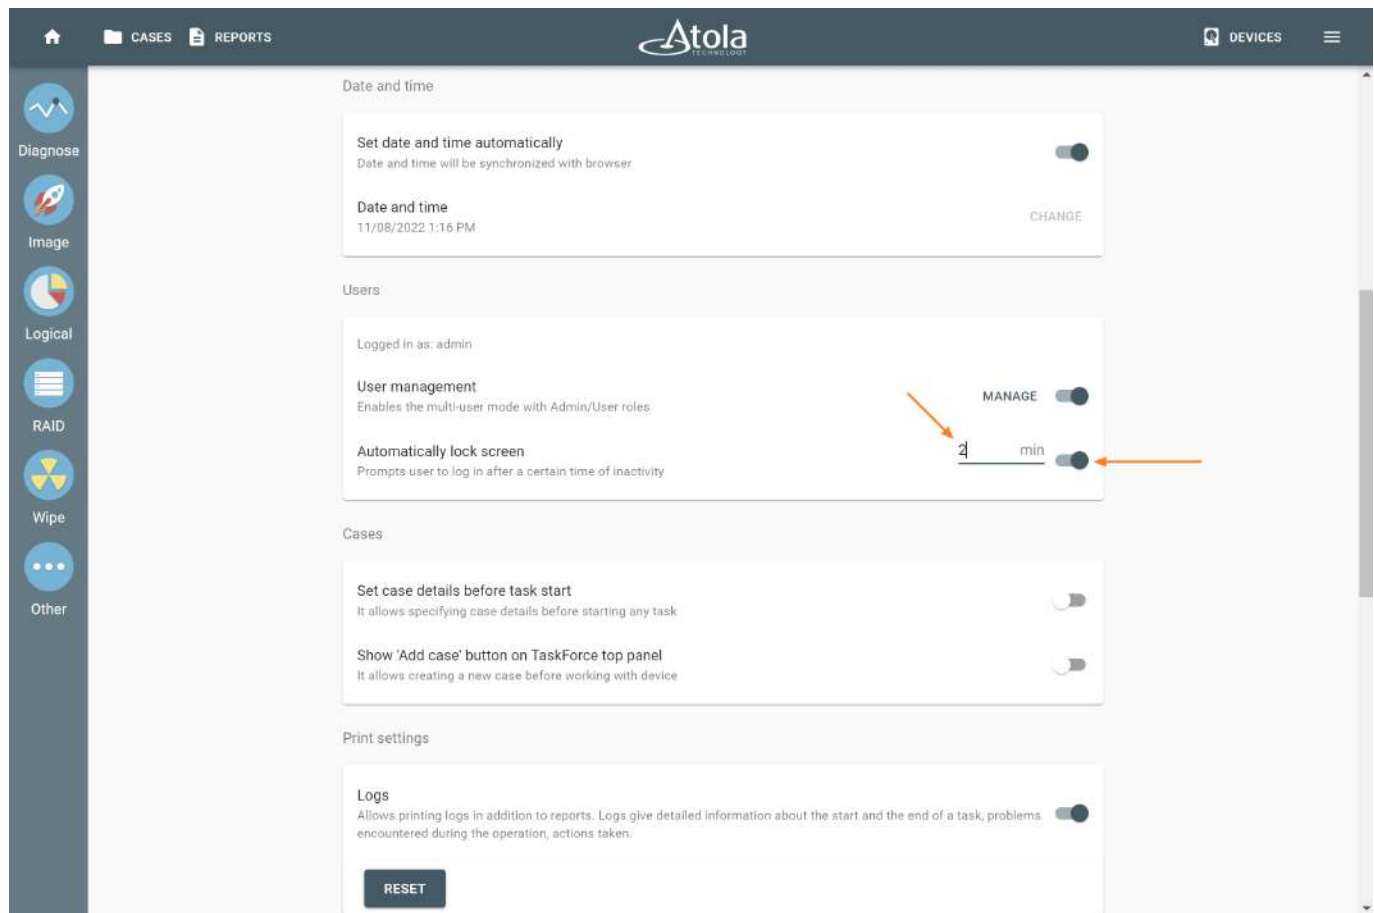


Changing a user password in the Edit user dialog on the Users page.

Automatically lock screen

As an Admin, you can set up Atola TaskForce to automatically lock its screen after a certain time of inactivity. To enable this feature, log in to a profile with the Admin role and do the next steps:

1. Go to **Menu > Settings**.
2. In the **Users** section, toggle **Automatically lock screen**.
3. Enter a time interval in minutes you want TaskForce to wait before automatically locking in.



The Automatically lock screen toggle in the Users section of the Settings menu.


Disable User management and password protection

As an *Admin*, you can disable user management and password protection for your Atola TaskForce.

If you disable User management, cases and reports will become available to all users of that TaskForce unit.

To disable user management and password protection, do the following steps:

1. Log in to a profile that has the Admin role.
2. Go to **Menu > Settings**.
3. In the **Users** section, toggle off **User management**.



Toggling off User management.

Maximize 10Gb network throughput

If you image evidence drives to a file on a server or a computer using a local network and want to achieve the best imaging speeds, maximize your network throughput by following these tips.

Set up 10Gb Ethernet network

Atola TaskForce is equipped with two 10Gb Ethernet ports. To fully use this potential and achieve the best imaging speeds, every node in your local network between the TaskForce unit and your computer or server need to support 10Gb connection as well.


Use 10Gb network adapter

Make sure that a network adapter on your server or your computer supports a 10Gb connection:

- Check vendor specifications for your network adapter.
- Check the network adapter speed in your OS settings.

How to check Ethernet connection speed on Windows 10 and Windows 11

1. Go to **Control panel > Network & Internet > Network and Sharing Center**.
2. In the **View your active networks** section, click your Ethernet connection.
3. In the **Status** window, check the connection speed. It must be equal to 10 Gbps.



Use 10Gb network cable

Make sure that the network cable you use to connect your TaskForce unit and your local network nodes supports 10Gb connection. It must have Cat6 (with the length of less than 55 meters), Cat6a, Cat7, Cat7a, or Cat8 marking on it.

Use 10Gb router or switch

If your TaskForce unit is connected through a network router or switch, make sure that every such device in your local network in between TaskForce and the target computer or server supports a 10Gb connection:

- Check vendor specifications for your network router or switch.
- Check that the 10Gb network cable is connected to a 10Gb port on your router or switch.

Enable jumbo frames

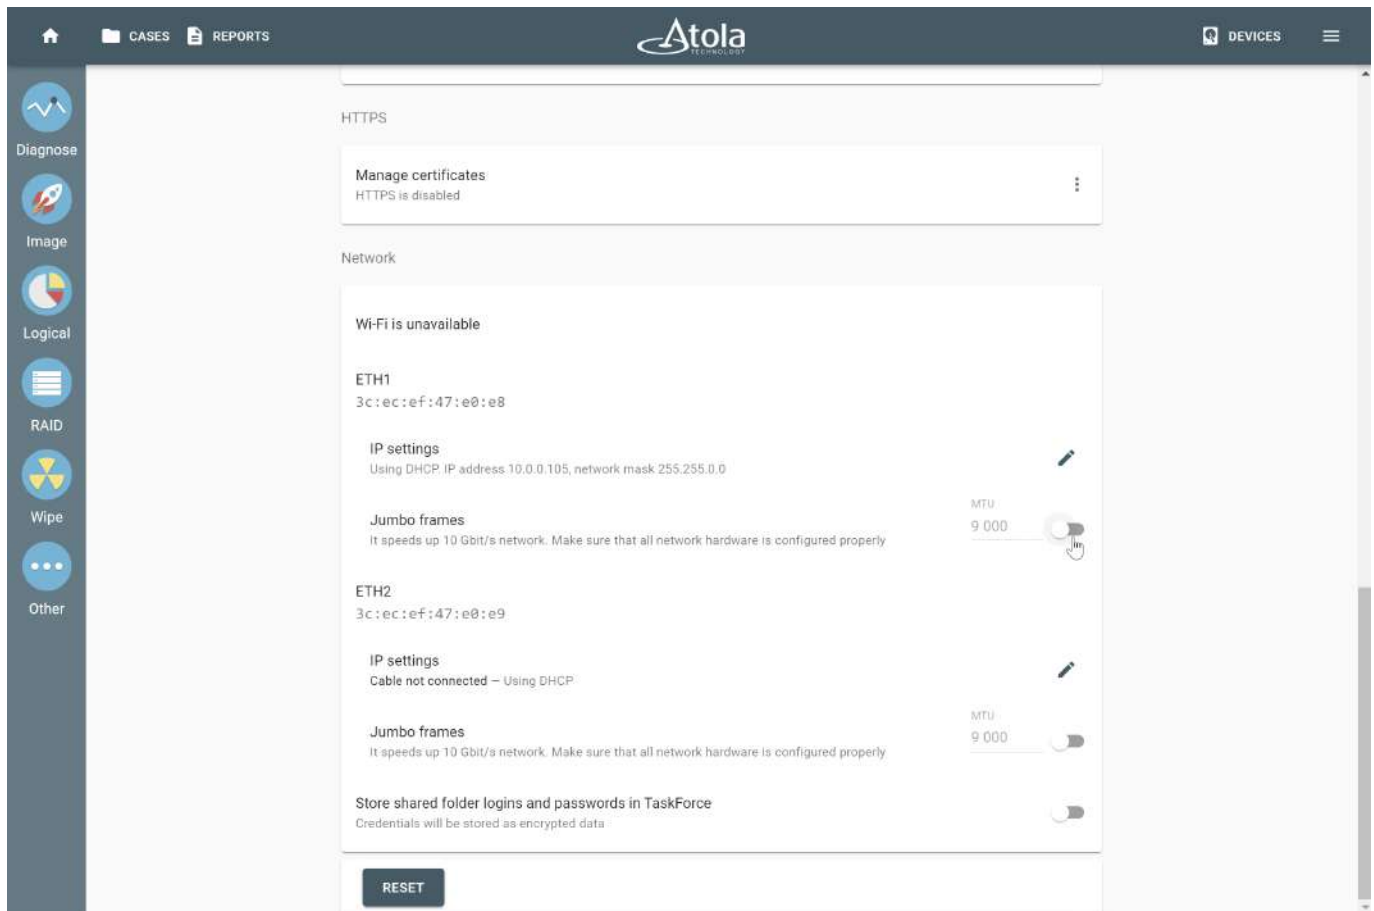
To speed up your 10Gb network, enable jumbo frames on your TaskForce, your target device, and every network switch or router in between them.

Jumbo frames increase the efficiency of broadband Ethernet processing because they carry 9000 bytes of payload, which is more than the standard limit of payload.

Enable jumbo frames on TaskForce

To enable jumbo frames on your TaskForce unit, do the following:

1. In the Atola TaskForce interface, go to **Menu > Settings**.
2. In the **Network** section, toggle Jumbo frames for respective Ethernet port (ETH1 or ETH2). Default MTU (maximum transmission unit) value is set to 9000 bytes, don't change it.




Jumbo frames toggle on the Settings page in Atola TaskForce.

Enable jumbo frames on a target computer

To enable jumbo frames for 10Gb network adapter on a target computer running Windows 10 or Windows 11, do the next steps:

1. Go to **Control panel > Network & Internet > Network and Sharing Center**.

2. In the **View your active networks** section, click on your Ethernet connection.
3. In the **Status** window, click **Properties**.
4. In the **Properties** window, click **Configure** and go to the **Advanced** tab.
5. On the **Advanced** tab, select **Jumbo Packet** and change its **Value** to 9014 Bytes.
6. Click **OK**.



Enable jumbo frames on a network router or switch

Jumbo frames need to be enabled on network routers or switches, which connect your TaskForce unit with your target computer or server.

Many 10Gb routers and switches support jumbo frames by default.

For specifications and instructions, refer to a router or switch user guide provided by its manufacturer.

Use Samba version 3.0 or later

Samba is a free software that provides fast shared access to files and printers for all network clients using the SMB/CIFS protocol.

We recommend using [Samba version 3.0 or later](#) to speed up the imaging to the files on network-connected drives.

For guidance, refer to the [Samba installation guide](#).

Update OS and drivers

Install the latest OS and driver updates.

Due to various software issues, the local network speed can be lower than expected if the operating system and network drivers are outdated.

Check server use quota

Sometimes, a network administrator can create quotas, which limit server space or throughput for end users.

Check if there is an opportunity to change or lift these quotas to achieve maximum data throughput during imaging.

Use faster target drives

To achieve higher imaging speeds, we recommend using faster destination drives with reading and writing speeds of no less than 600 MB per second.

If you're using a RAID array as a storage, consider choosing RAID types with faster reading and writing speed, such as RAID 0 and RAID 5.


When storing very big files, opt for a large cluster size on your server. The bigger files you want to store, the larger cluster size should be.

Use a static IP address

If you plan to connect TaskForce directly to your computer, we recommend configuring your computer and TaskForce to use a static IP address. Here's how to do it.


Configure TaskForce for using a static IP address

1. On your computer, open a Chrome browser and in its address bar enter the IP address, displayed on the TaskForce front panel. You can use TaskForce touchscreen as well.
2. In the TaskForce interface, go to **Menu > Settings**.



TaskForce system menu.

3. In the **Network** section, find the **IP settings** for the ETH1 or ETH2 port and click the **Edit** icon next to it.




The Edit icon next to IP settings for the ETH 1 port.

4. Enter the following network settings:

- **IP address:** 10.0.0.XXX, where XXX can be any number from 1 to 254.
Default static IP addresses of the TaskForce unit are 10.0.0.215 and 192.168.0.215.
The IP address of your PC's Ethernet card must be different from that of the TaskForce unit.
- **Network mask:** 255.0.0.0.
If your PC and the TaskForce unit belong to different subnets, the connection can't be established.
- **Default gateway** and **DNS** can be left empty or set to any value.

5. Click **Save**. The new IP address appears on the TaskForce front panel.



Network settings for the ETH1 port.

Connect TaskForce directly to your computer

To use a static IP address, connect your computer to the ETH1 or ETH2 port of TaskForce with an Ethernet network cable.

If your computer has no Ethernet port, you can use the USB-to-Ethernet adapter (not included in the package). Connect it to your computer before configuring a network.

Configure your computer for using a static IP address


1. In Windows, open **Network Connections**: press *Win+R*, enter *ncpa.cpl* and click **OK**.
2. Right-click your Ethernet adapter, and then select **Properties**.
3. Select **Internet Protocol Version 4**, and then click **Properties**.
4. Enter the following network settings:
 - **IP address**: 10.0.0.XXX, where XXX can be any number from 1 to 254 except for 215.
Default static IP addresses of the TaskForce unit are 10.0.0.215 and 192.168.0.215.
The IP address of your PC's Ethernet card must be different from that of the TaskForce unit.
 - **Network mask**: 255.0.0.0.
If your PC and the TaskForce unit belong to different subnets, the connection can't be established.
 - **Gateway** and **DNS server** can be left empty or set to any value.
5. Click **OK**.

Jumbo frames for fast imaging to server

In TaskForce, Jumbo frames are activated by default to ensure maximum data transfer rates when imaging to a file on your server.


However, if Jumbo frames have been disabled, it is easy to enable them again and experience substantial boost to the speed of imaging!

First, create a file on the server, to which you will be imaging.



Creating a target file.


When you start to image to your server with Jumbo frame disabled, the data transfer speed will not exceed 500 MB/s. The actual speed will also depend on the configuration and current traffic in the network.




Speed of imaging without Jumbo frames.

To boost the speed:

1. Pause the imaging session.
2. Go to the **Menu** at the top right.
3. Click **Settings**.



4. In the **Network** section, enable **Jumbo frames** of the Ethernet port you are using and set **MTU** to 9000.




Enabling Jumbo frames.

For fast imaging of files to your server via 10Gb network, you need to activate Jumbo framework in the settings of the server's network adapter as well as in the settings of the network switch, should it be necessary (10Gb switches normally have Jumbo framework activated by default).

Then you can return to your imaging session:

1. Go to **Image**.
2. Select the source.
3. Click **Resume**.

This time, the speed is way higher!



Imaging speed with Jumbo frames enabled.

Network setup tips

- [Configuring 10Gb network with DHCP-enabled switch](#)
- [Getting maximum performance of Ubiquiti network](#)
- [Configuring a dynamic IP for TaskForce in a network without router or DHCP-enabled switch](#)
- [Accessing Windows Server 2012 shared folder](#)
- [Setting up Synology DS218 as storage server](#)

Configuring 10Gb network with DHCP-enabled switch

You need to create or extend network with DHCP-enabled switch with 10Gb connection.

Example. [Ubiquiti EdgeSwitch 16 XG](#): four 10Gb Ethernet ports, twelve 10Gb SFP ports. Approximate price: \$600.

This kind of switch supports static IP setup via simple web admin. So you could set the IP addresses you need for each current network device.

How to configure Ubiquiti DHCP server:

1. Connect PC and TaskForce to Ubiquiti switch.
2. Set static IP address of PC to 192.168.1.4.
3. Open a browser and enter 192.168.1.2 (default Ubiquiti switch IP).
4. Log in with default credentials: ubnt (both in name and password fields).
5. Go to **System > Advanced Configuration > DHCP server > Global**.
6. Activate **Admin mode** by checking a necessary checkbox and pressing **Submit** button.

7. Go to **Pool Summary** and press **Add** to make a new address pool.

8. Enter your:

- pool name
- network base address (for example, 192.168.1.0)
- network mask (in most cases, it should be 255.255.255.0)
- put Default Router Address and DNS

After creating your pool, you can change it via **Pool configuration** tab.

9. Click **Save configuration** button in the upper right corner of the window and click **Save**.

You can check this [Youtube guide](#) for alternative instructions on network setup using Ubiquiti switch.

Getting maximum performance of Ubiquiti EdgeSwitch 16 XG network

To optimize performance using Ubiquiti EdgeSwitch 16 XG, you need to enable 10Gb with jumbo frames:

1. Go to **Basic > Port summary**.
2. Select ports 0/13, 0/14, 0/15, 0/16 and click **Edit**.
3. Change **Maximum Frame Size** to 9014 in **Edit Port** configuration window.

Configuring a dynamic IP for TaskForce in a network without router or DHCP-enabled switch

If there is no hardware in the network that assigns IP address, or if you want to keep a small network with TaskForce and your server/PC connected directly, it is possible to install and setup software DHCP server. The good news is, it does not require any investment. All you need is some time to set it up on any computer in the server network. Follow the instructions from these guides:

- Windows: [Free DHCP server](#).
- Linux: [Included DHCP service](#).

Accessing Windows Server 2012 shared folder

If you want to store a target image file in a Window Server 2012 network folder but it appears missing, please follow these steps:

1. Go to **Control panel**.
2. Enable **Guest** account (**Administrative tools > AD users and computers > Users**).
3. **Network and sharing center > Change advanced sharing settings > Turn On network discovery + Turn on sharing (file and printers + public folders)**.
4. In the shared folder access options, add **Guest** or **Everyone**.

If the shared folder demands restricted access, please follow [this guide](#).

Setting up Synology DS218 as storage server

To set up Synology DS218:

1. Go to **Control panel > File services > SMB > Advanced settings**.
2. Set **Maximum protocol** to SMB3.
3. Go to **Control panel > Shared folder**.
4. Click **Create** button and specify network folder details.

If you need to get a guest account working, run the following actions:

1. Go to **Control panel > User**.
2. Edit for **Guest** user.
3. Clear the **Disable this account** checkbox.

For more instructions and information about check our [Troubleshooting](#) guide and [FAQ](#) page.

Configuring TaskForce use with Synology or QNAP NAS as a LAN

Here is a guide to configure your NAS to be used with TaskForce for storage of image files.

This manual is indicative and may vary between models and manufacturers. We have tested this setup on Synology DS218 and QNAP TS-431K.

Table of contents

[Setting up Synology](#)

[Setting up QNAP](#)

[Direct connection with a Static IP:](#)

1. [Configuring TaskForce](#)
2. [Configuring Static IP in Synology](#)
3. [Configuring Static IP in QNAP](#)

[Setting up LAN using a router with DHCP \(Dynamic IP\):](#)


1. [Configuring the router](#)
2. [Configuring DHCP in TaskForce](#)
3. [Configuring DHCP in Synology](#)
4. [Configuring DHCP in QNAP](#)

Setting up a Synology NAS

1. Install drives in the Synology NAS.
2. Use a LAN cable to connect the DiskStation to your switch, router, or hub.
3. Press the power button to turn on your Synology NAS.
4. Install DiskStation Manager (DSM) – Synology’s browser-based operating system – on your DiskStation.
5. Follow the instructions to create a RAID volume and a partition with a file system on it.


Enabling SMB3 is highly recommended for the best performance:

- a. Go to **Control panel > File services > SMB > Advanced settings**.
- b. Set **Maximum protocol** to *SMB3*.
- c. Go to **Control panel > Shared folder**.
- d. Click **Create** and specify network folder details.



6. Create a Shared folder in the new volume:

Control panel > File sharing > Shared folder > Create > Create Shared Folder




7. In the folder creation menu, set the folder name, e.g. *Image Files* and set access options.


Setting up a QNAP NAS

1. Install drives in the QNAP NAS.
2. Use a LAN cable to connect the QNAP to your switch, router, or hub.


3. Press the power button to turn on your QNAP.
4. Install Qfinder Pro for quick find and easy access to QNAP NAS on the same LAN.
5. Find your QNAP using Qfinder Pro and go to its page in the browser.
6. Follow instructions to complete Smart Installation.




7. Create Storage Pool (RAID):
 - a. Storage & Snapshots.
 - b. Storage (left menu).
 - c. Storage/Snapshots.
 - d. New Storage Pool (in the the top right corner).
 - e. Follow the instructions to set up the Storage Pool.




8. In the final stages of setting up the Storage Pool, create a New Volume.



9. In the File Station utility, you can find the new volume: QNAP. In it, there is a Shared Folder created by default called *Public*.



10. For a Shared Folder with specific access or encoding settings, create a new one by clicking the icon showing a folder with a + and adjust the configuration.



Ethernet



Direct connection with a Static IP:

Since we are not connecting TaskForce and Synology/QNAP NAS to a wider network, it is enough to set up the same Subnet Mask for the NAS and the TaskForce, and configure Static IP addresses that will indicate the same network.

TaskForce configuration

1. Go to **Menu > Settings > Network** section.
2. Select the Ethernet port through which Synology is connected (e.g. ETH1).
3. Click IP settings's **Edit** icon.

Network

Wi-Fi Hotspot

SETTINGS

ETH1

3c:ec:ef:40:33:b0

IP settings

Cable not connected — Using DHCP

Jumbo frames

It speeds up 10 Gbit/s network. Make sure that all network hardware is configured properly

MTU

9,000

ETH2

3c:ec:ef:40:33:b1

IP settings

Cable not connected — Using DHCP

Jumbo frames

It speeds up 10 Gbit/s network. Make sure that all network hardware is configured properly

MTU


9,000

Store shared folder logins and passwords in TaskForce

Credentials will be stored as encrypted data

RESET


4. In the **ETH settings** dialog, enter the address (e.g. *255.255.0.0*) in the **Subnet Mask** field that is identical to that of Synology.
5. Assign an IP address using the Subnet Mask: the first two bytes of TaskForce's IP must be identical to those of IP Synology (e.g. *10.0.0.15*). This will ensure that both TaskForce and Synology belong to the same network. (e.g. *10.0.0.8*).




6. Click **Save**.

Setting up a Static IP in Synology

1. Go to the **Control panel**.
2. Open **Network interface**.
3. Select the current connection to Synology.
4. Click **Edit**.




5. On IPv4 tab, select **Use manual configuration**.
6. Assign **Subnet Mask** (e.g. *255.255.0.0*), which is identical to that of TaskForce.
7. Assign **IP address** (e.g.: *10.0.0.15*).
8. Click **OK**.




9. After Synology has changed the settings, it can be disconnected from the current network.

Setting up Static IP in QNAP

1. In the Control panel, go to Network & Virtual Switch.
2. Select the current connection to QNAP.
3. Open the Menu.
4. Click **Configure**.




5. On the IPv4 tab, select **Use static IP address**.
6. Assign a **Subnet Mask** (e.g. 255.255.0.0), which is identical to that of TaskForce.
7. Assign an **IP address** (e.g.: 10.0.0.16).
8. **Gateway** field should be left blank because you are configuring a LAN and connecting to other networks is unnecessary.
9. **Jumbo Frame** and **Network speed** should be left with the default values.
10. Click **Apply**.



11. After the QNAP has changed the settings you can disconnect it from the current network.


Connecting to the Shared Folder in TaskForce

1. Connect the Synology/QNAP NAS to TaskForce's ETH1 port.
2. Click the **Devices** icon in the upper right corner. Go to the **File** section and click **Select File**.
3. Click **Connect**.
4. Enter the Server name, Username, Password.
If you do not remember the server name, look it up:
Synology > Control Panel > Info Center > Server name
QNAP > Control Panel > General Setting > Server name
5. Click **Connect**.




6. TaskForce has connected to Synology and you can open the previously created Shared Folder.

Setting up a LAN using router with DHCP (Dynamic IP)




Setting up router

1. Connect the router directly to your PC or laptop via the Ethernet port.
2. Log in to the web configurator.
3. Turn on the DHCP server in the local network settings.
4. Apply changes.




Setting up DHCP in TaskForce

1. Open **Menu > Settings**, go to the **Network** section.
2. Select the Ethernet port to which Synology will be connected. (e.g. ETH1).
3. Click the **Edit** icon to adjust the IP settings.
4. Clear the **Use static IP address** checkbox.
5. Click **Save**.




Setting up DHCP in Synology

1. Go to Control panel > Network > Network interface.
2. Select the current connection to Synology and click **Edit**.
3. On the IPv4 tab, select the **Get network configuration automatically (DHCP)** option.




Setting up DHCP in QNAP

1. In the Control panel, go to the Network & Virtual Switch.
2. Select the current connection to QNAP, open the Menu and select Configure.



3. On the IPv4 tab, select Obtain IP address settings automatically via DHCP.



Connecting the LAN devices

After all the reconfigurations, connect TaskForce, NAS and your PC or laptop to the router using Ethernet ports and connect to the Shared Folder [as described above](#).

Connecting drives & starting Atola TaskForce

This page provides information about Atola TaskForce start up procedure to ensure safe and effective operation of the unit.

Powering on TaskForce

The power switch is located on the back panel of the unit. To start TaskForce, turn the power switch on.



Booting

The booting process takes up to 3 minutes.

Once booting is completed, the IP screen on the front panel will display either "Standalone mode" message or the IP address if the unit be connected to the Ethernet. At this point, the unit is ready for operation.



TaskForce screen

The Microsoft Surface Pro tablet that serves as TaskForce's screen is switched on and off independently from the unit, by pressing the button in the tablet's top panel.



Connecting drives

TaskForce supports SAS, SATA, USB, and IDE drives via its 17 ports, as well as other storage devices via Thunderbolt, Apple PCIe, and M.2 SSD extension modules.

To ensure both TaskForce and the devices connected to it are used properly and safely, read the instructions below.

Connecting USB devices

TaskForce system must be **powered on** before a USB device is plugged in. As soon as booting is finished and the IP address is displayed you can plug the USB device into any of 4 USB ports available.

If you connect a USB device before starting TaskForce, the imager will not be able to boot correctly.



Connecting Extensions

Before connecting an extension module, make sure TaskForce is **powered off**. Plug the extension module into the extension slot located on the back panel of TaskForce and power the unit on.



Connecting evidence drives

Each port is equipped with an individual Source switch enabling hard write protection on the port. To make sure data on the drive is not overwritten, make sure the port is in the source before you connect the evidence drive.



Connecting SATA & SAS drives

TaskForce has 6 SATA and 6 SATA/SAS ports. Before connecting evidence and target devices to the imager, make sure the ports are switched to the right mode.

When a drive is connected to a running TaskForce imager, the port is by default powered off. To identify the device plugged into the port, click the **Devices** button in the top panel of the TaskForce interface, and TaskForce will start to identify all connected drives. The system ensures sustainable overall power consumption in the situations when many drives are plugged in.

Powering off TaskForce

If TaskForce is not running any processes, it is safe to power off the unit by turning off the power switch. Any sessions that were active at the moment the unit was powered off (whether it is due to an outage or the power switch being turned off) will not be stopped correctly and cannot be resumed later.

Supported Drives

Atola TaskForce supports all 1.8-inch, 2.5-inch, 3.5-inch IDE, SATA, SAS and USB hard drives, USB Flash media as well as SD, Compactflash, and Memory Stick cards via a generic USB Card Reader.

TaskForce can also work with the following drive types using proprietary Atola extension modules:

- M.2 NVMe/PCIe/SATA SSDs
- latest Apple SSDs via Thunderbolt extension
- the newest PCIe SSDs from Apple MacBooks (2013 - 2015)

Remote image acquisition can be performed using the iSCSI protocol. To do this, you need to mount a drive on another PC as an iSCSI target.

Most functions of Atola TaskForce will work with any hard drive or flash card with either IDE, SATA-1/2/3 or USB-1/2/3 interface (including those attached via adapters).

To ensure high quality and efficiency of our tools, [we test them on hundreds of storage devices](#).

Head selection works for the following HDD models

- SATA and IDE Seagate hard drives (including F3 series)
- SATA and IDE Western Digital, HGST hard drives with exception of some models released since 2018
- SATA and IDE Hitachi hard drives
- SATA and IDE Toshiba hard drives: MG, MK, MQ, DT, HD families with exception of some models released since 2018

RAID support: unknown configuration detection, reassembly, and imaging

- RAID 0, 1, 5, 6, 10, JBOD
- Hardware RAID controllers: Adaptec, Areca, HP, Dell, LSI, Intel RST, IBM ServeRAID, other generic controllers
- Software RAID: mdadm
- Synology NAS RAID
- File systems: NTFS, exFAT, ext4/3/2, XFS, ZFS, Btrfs, APFS, HFS/HFS+, FAT32/16

TaskForce drive identification

Atola TaskForce is designed to perform multiple processes simultaneously and provide its users with unprecedented flexibility when it comes to a variety of devices and configurations in which they can be used. TaskForce also efficiently communicates how a device is being used and helps a user to handle drives correctly.

Switch between the Source and Target mode

When connecting a drive to the system, make sure the right mode is set on the port: in source mode, an evidence drive is automatically write-protected. It can only be changed with Source hardware switches.



Source hardware switches

Initial drive identification


As soon as you choose a particular task or click **Devices** button in the top bar, TaskForce starts sequentially supplying power and sending commands to identify all connected devices.

After a connected drive receives power supply and identification commands from the unit, it responds with device info including:

- device model and serial number;
- device capacity;
- limitations of the drive.

Detection of password protection and HPA/DCO/AMA restrictions

TaskForce software also immediately detects whether the drive is locked by ATA password or the drive's max readable address is limited via [HPA/DCO/AMA](#). The unit indicates these restrictions and notifies a user about those with red color indication in the device menu.




TaskForce detects ATA, HPA/DCO/AMA limitations

These indicators allow a user to make informed decisions on how to proceed with the device, whether unlocking is required to get access to the whole drive space before starting an imaging session.

The 'Device not detected' notification

Notification **Device not detected** may point to one of these issues:

- there is no device on the port;
- the cable is not properly plugged in;
- the device is connected to another port;
- the device is heavily damaged.



Device not detected

Re-identify device

You might want to re-identify a device in one of the following cases:

- A device hasn't been detected from the first try.
- You have connected another device to the drive port.

To re-identify a single device from the **Select device** panel:


1. In the port with your device, click the **More** icon (the one with three dots).
2. Select Re-identify.

To re-identify all connected devices, use the **Re-identify** button in the top right corner of the **Select device** panel. It affects only the devices in unfolded sections of the panel.

You can also re-identify a single device from its device page by clicking the **Re-identify** button at the bottom.


When the drive port is unavailable

If a drive is busy with a running operation, the port will be temporarily unavailable for selection when launching other tasks. In such case, the fonts in the respective box will be a lighter shade of grey, making the port unclickable.



Busy target drive

When selecting a target device for wiping or imaging, source drives are also unavailable to ensure that data on an evidence drive doesn't get overwritten by mistake.



Selecting target device

Add a device to a case

Once a drive is connected to a drive port and identified, you can add it to a new or existing case right from the **Select device** panel:

1. In the port with your device, click the **More** icon (the one with three dots), and then select **Add case**.
2. Choose between **Create new case** or **Select existing case** and click **Next**.
3. Enter details for a new case or select an existing case from a table.

Power off and power on a drive

From the **Select device** panel, you can power off and power on any connected device. To do that, in the port with your device, click the **More** icon (the one with three dots), and then select the **Power on** or **Power off** option.

Power on all devices: if this checkbox at the top right is selected, TaskForce automatically powers on all plugged drives after the sliding panel appears.


You can also power a device on and off from a device page, by clicking the respective button at the bottom of the page.

When a device is idle for 5 minutes, TaskForce automatically powers it off. This is the setting that you can disable.

However, we recommend powering off a device before disconnecting it from TaskForce.

Atola TaskForce: Main window

This article helps in understanding TaskForce's main window, its controls and buttons and how to use them.



An overall view of the TaskForce main window: 1 - Home icon, 2 - Taskbar, 3 - Cases button, 4 - Reports button, 5 - Current overall performance, 6 - Devices button, 7 - Menu.

1. Home icon

The **Home** button brings you back to the **Home** screen. This is where you can check the active and recently completed tasks in the respective sections of the screen.

The number of current active processes is indicated next to the **Home** icon in a small orange circle.


2. Taskbar

The **Taskbar** on the left shows main TaskForce operations:

1. Diagnose
2. Image
3. Logical
4. RAID
5. Wipe

Click the **Other** button in the **Taskbar** to access 4 more important TaskForce operations:


1. Hash: Calculates hash of a device: MD5, SHA-1, SHA-256, SHA-512.
2. View SMART: Shows all SMART attributes as a table.
3. Browse files: Shows the content of partitions on a device.
4. Hidden drive areas: Lets you to unclip or change HPA, DCO, AMA limitations.



Section with other tasks: Hash, View SMART, Browse files, Hidden drive areas.

3. Cases

By clicking the **Cases** button in the top panel you get to the **Cases** page with a list of the latest cases. With the help of the **Search bar** you can [find a specific case](#). The cases are available for import and export between different TaskForce units.




Current overall performance shows after you click Atola logo.

6. Devices

Click **Devices** button in the top panel to see all the drives connected to TaskForce to obtain maximum information about each by simply clicking it.


The **Devices** panel provides additional options for working with the drives: [add a case](#) for a device, re-identify or power off any device.



The Devices panel is expanded and shows all the drives connected to TaskForce.

7. Menu

The Menu contains device settings and features that regulate your use of the TaskForce unit.



The Menu gives access to Atola TaskForce settings, express mode, full screen mode, release notes, and manual.

In **Settings** you can adjust the general, database, network, print, and other settings.

Express mode enables automatic launch of multiple imaging sessions on all ports that are set to source. For more details, see [Express mode: self-launching imaging](#).

Toggle fullscreen option is handy when working with other programs or files.

In **Update firmware** you can check the current TaskForce firmware, choose update method and perform the firmware update by selecting and downloading the firmware file.

Activation status lets you to look up, reactivate the status, or extend the subscription.

In **Release notes** you can read the information about the most recent Atola TaskForce firmware release and track all updates and enhancements by clicking corresponding links.

ATA registers: what they mean



Link Register

It's only enabled when port powered on, device presence detected and PHY communication established.

Status Register

This register contains hard drive status information. It is updated after every single command sent to the drive.

ERR: means last command failed to execute. In this case the Error register contains more details on the specific error.

INDX: obsolete, used to trigger after each spindle revolution

CORR: obsolete, used to trigger after a bad sector was automatically corrected by ECC

DREQ (Data Request): is asserted when hard drive wants to exchange data with the host controller (in either direction)

DRSC (Device Seek Complete): is obsolete; always asserted on modern hard drives

FAULT (Write Fault): is obsolete

RDY (Device Ready): is obsolete; always asserted on modern hard drives

BUSY: indicates that the hard drive is busy executing a command OR initializing (after power on or reset)

Error Register

Error register provides more details if the last command failed. This register is only valid when **ERR** bit of the Status Register is asserted.

AMNF: means Address Mark Not Found (usually occurs on failed read attempt)

TONF (Track 0 Not Found): obsolete

ABRT: command aborted (unsupported command or other failure)

IDNF: sector ID not found (usually occurs on failed read attempt)

UNC: uncorrectable read error; the hard drive was unable to read data even after applying ECC recovery algorithms

ICRC (Interface CRC error): there was CRC error while transferring data between host and the hard drive (usually indicates bad interface cable)

Diagnosing a drive with Atola TaskForce


When an evidence drive lands on investigator's table for the first time, there is always an uncertainty when it comes to the drive's condition. A broken head or scratched surface of the media require different imaging tactics. That's why it is strongly suggested that before imaging, each drive should first be diagnosed.

TaskForce has Atola's unique diagnostics module which checks all systems of the drive:


- Hard drive's motor and electronics (PCB)
- Head stack
- Media surface
- All firmware/system areas
- Partitions and file systems

At the end the system produces a report which sums up all issues. The process will take only 2-5 minutes.

To start, go to **Diagnose**, select the drive and then click **START**.



First, TaskForce checks the drive's printed circuit board. The system applies power to the device and records and analyzes spin-up current curve. This helps detect most issues with the PCB and the motor. Next, TaskForce analyzes the contents of the hard drive's ATA registers and device identification sector.



Circuit board check.

After that, the head stack is tested. Several factors are taken into consideration when diagnosing heads: media access time for each head, power consumption curves, and internal drive's error reporting systems.

If the head stack looks good, the system performs a short media scan. The purpose of this scan is to verify if there are any bad sectors in the starting, middle and ending sectors of the drive pointing to a damage to the media surface or logical errors.

1

CASES

REPORTS

Atola
TECHNOLOGY

DEVICES

Diagnose

Image

Logical

RAID

Wipe

Other

Samsung SSD 860 PRO 256GB S418NF0KB07067E

256 GB

Healthy

SATA 5

Case ID:
MI-r2343451

Verifying ending sectors (LBA: 499,118,192 - 500,118,191)

100%

Check list
Circuit board
Heads
Media surface
Firmware
File system

Solid state drives (SSD) don't have heads.

Elapsed: 0.0 sec.

3. Media surface

Media scan is in progress...

Verifying starting sectors (LBA: 0 - 999,999)

0 error(s) found

Speed: 542 MB/s

Verifying middle sectors (LBA: 250,059,095 - 251,059,094)

0 error(s) found

Speed: 541 MB/s

Verifying ending sectors (LBA: 499,118,192 - 500,118,191)

0 error(s) found

Speed: 542 MB/s

Average speed: 541 MB/s

Media surface check passed

Elapsed: 2.9 sec.

STOP

Link: PHY Status: ERR INDX CORR DREQ DASC TAULT DDDV BUSY Error: AMAN TONR ABIT - IDNF - UNC ICRC ?

Heads and media surface check.

Next, several firmware tests are performed:

1

CASES

REPORTS

Atola
TECHNOLOGY

DEVICES

Diagnose

Image

Logical

RAID

Wipe

Other

Samsung SSD 860 PRO 256GB S418NF0KB07067E

256 GB

Healthy

SATA 5

Case ID:
MI-r2343451

Firmware

100%

Check list
Circuit board
Heads
Media surface
Firmware
File system

4. Firmware

Device is not locked.
Device identification data is valid.
Max address according to device ID: 500,118,191
Native max address (ext): 500,118,191
Max Address from DCO: 500,118,191.
Reported capacity appears logically correct.
Performing SMART checks...
Power cycles: 1477 Smart status: Good Powered on: 18 days 20 hours


#	Attribute Name	Value	Worst	Threshold	RAW	Status
5	Reallocated Sector Ct	100	100	10	0	OK
9	Power On Hours	99	99	0	452	OK
12	Power Cycle Count	98	98	0	1477	OK
177	Wear Leveling Count	93	93	0	149	OK
179	Used Rsvd Blk Cnt Tot	100	100	10	0	OK

STOP

Link: PHY Status: ERR INDX CORR DREQ DREQ TAULT DREQ BUSY Error: AMAN TONF ABRT - IDNF - LINC ICRC ?

Firmware check.

If TaskForce detected no issues by this point, it performs a file system checkup:



File system check.

After this final stage of diagnostics, TaskForce displays the full report. The **Diagnostics** result message box contains a short summary of all tests. It also provides estimated imaging time for this drive.



Diagnose



Image



Logical



RAID



Wipe



Other

Samsung SSD 860 PRO 256GB S418NF0KB07067E
256 GB

SATA 5

Case ID:

Not assigned

Creation date: 07/25/2023 7:39 PM

TaskForce serial: 23166733

Location: SATA 5

Device model: Samsung SSD 860 PRO 256GB

Size: 256 GB (256,060,514,304 bytes)

Case investigator:

Software version: 2023.7

TaskForce IP: 192.168.1.113

Write protection: Off

Device serial: S418NF0KB07067E

Device firmware: RVM01B6Q

Case description:

Diagnostics completed

Results

No major hardware or firmware issues found

75% of the disk is not associated with any partition.

Estimated imaging time: 7 minutes

1. Circuit board

Device is powered on. A power cycle is needed...

Powering down the device...

Applying power and watching spin-up currents...

Current oscillogram (12V):



Current oscillogram (5V):



Device has become ready in 3 sec

Registers. Status: 01010000 Error: 00000001

Peak power consumption during spin-up: 5V line = 322.30 mA; 12V line = 0.00 mA

Integrity word of device identification data is OK.

Device identified: Samsung SSD 860 PRO 256GB SN: S418NF0KB07067E

Logical sector size: 512 bytes. Physical sector size: 512 bytes

Circuit board check passed

Elapsed: 12.3 sec.

2. Heads

Solid state drives (SSD) don't have heads.

Elapsed: 0.0 sec.

3. Media surface

Media scan is in progress...

Verifying starting sectors (LBA: 0 - 999,999)

0 error(s) found

Speed: 542 MB/s

Verifying middle sectors (LBA: 250,059,095 - 251,059,094)

0 error(s) found

Speed: 541 MB/s

Verifying ending sectors (LBA: 499,118,192 - 500,118,191)

0 error(s) found

Speed: 542 MB/s

Average speed: 541 MB/s

Media surface check passed

Elapsed: 2.9 sec.

4. Firmware

Device is not locked.

Max address according to device ID: 500,118,191
Native max address (ext): 500,118,191
Max Address from DCO: 500,118,191.

Performing SMART checks...

#	Attribute Name	Value	Worst	Threshold	RAW	Status
5	Reallocated Sector Cnt	100	100	10	0	OK
9	Power On Hours	99	99	0	452	OK
12	Power Cycle Count	98	98	0	1476	OK
177	Wear Leveling Count	93	93	0	149	OK
179	Used Rsvd Blk Cnt Tot	100	100	10	0	OK
181	Program Fail Cnt Total	100	100	10	0	OK
182	Erase Fail Count Total	100	100	10	0	OK
183	Runtime Bad Block	100	100	10	0	OK
187	Reported Uncorrect	100	100	0	0	OK
190	Airflow Temperature Cel	69	41	0	31	OK
195	Hardware ECC Recovered	200	200	0	0	OK
199	UDMA CRC Error Count	98	98	0	1071	OK
235	POR Recovery Count	99	99	0	334	OK
241	Total LBAs Written	99	99	0	72409915009	OK

The diagram shows the device temperature history during the recent work time intervals between power cycles.



Last significant work time interval before power off or standby (in minutes): 20

5. File system

Elapsed: 0.1 sec.

⏪ ⏩ ⏴ ⏵

Making a definitive conclusion based on the indices in SMART table is not easy: not all parameters are critical, it is usually a combination of bad values of a few parameters that point to a trouble, time factor plays a role too (how fast has the state of the drive been deteriorating).

View SMART table

SMART table is included in the [Diagnostics report](#). If you want to have a look at the current indices:

1. Go to **Other > View SMART**.
2. Select the drive.
3. Click **Start**.

The screenshot shows the TaskForce software interface. The top navigation bar includes 'CASES', 'REPORTS', and 'DEVICES'. The left sidebar has icons for 'Diagnose', 'Image', 'Logical', 'RAID', 'Wipe', and 'Other'. The main content area displays the 'View SMART completed' report for a Kingston SKC600256G 50026B778374DEAC drive (256 GB, SATA 6). The report includes metadata such as creation date, TaskForce serial, location, device model, size, case investigator, software version, TaskForce IP, write protection, device serial, device firmware, and case description. Below this, the SMART table is shown with columns for Attribute Name, Value, Worst, Threshold, RAW, and Status. The table lists 12 attributes, all with a status of 'OK'. At the bottom, there are 'PRINT' and 'GO TO CASE' buttons, and navigation arrows.

View SMART completed

Power cycles: 1639 Smart status: Good Powered on: 27 days 13 hours

#	Attribute Name	Value	Worst	Threshold	RAW	Status
1	Raw Read Error Rate	100	100	0	0	OK
5	Reallocated Sector Ct	100	100	0	0	OK
9	Power On Hours	100	100	0	661	OK
12	Power Cycle Count	100	100	0	1639	OK
148	Total SLC Erase Ct	100	100	0	368	OK
149	Max SLC Erase Ct	100	100	0	42	OK
150	Min SLC Erase Ct	100	100	0	14	OK
151	Average SLC Erase Ct	100	100	0	24	OK
159	DRAM 1 Bit Error Count	100	100	0	0	OK
160	Uncorrectable Error Cnt	100	100	0	0	OK
161	Valid Spare Block Cnt	100	100	0	39	OK

PRINT GO TO CASE

SMART table report.

SMART table attributes may differ depending on the drive manufacturer. The most critical attributes are:


- Reallocated sectors count.
- Current pending sector count.
- Uncorrectable sector count.

When RAW value of any of these attributes is greater than zero, TaskForce highlights it in yellow.

The worse the values, especially in these critical attributes, the more carefully the drive needs to be treated.

Track changes of the SMART table attributes

To keep track of the changes occurring to the attributes of the SMART table, the imaging settings can be easily adjusted to records SMART table indices prior and after each imaging session.



Adjust the imaging settings to keep the record of SMART prior and after the imaging session.

By comparing the two tables, user can evaluate whether the health of a drive has been deteriorating throughout the imaging session and thus assess how quickly its health has been getting worse. Any discrepancies between the two SMART tables will be highlighted in yellow.



How SMART table state changed after image acquisition

Whenever you need to evaluate how the state of the drive has been changing long-term, go to previous imaging sessions and look up SMART table. TaskForce stores this information in its [case management system](#).

Imaging an evidence drive to 5 targets


Atola TaskForce allows imaging to up to 5 targets at a time.

The targets may include:

- E01, AFF4, RAW file on a network server.
- target drive plugged into one of 18 TaskForce ports.


To start an imaging session that includes 5 targets:

1. Go to **Image**.
2. Select a source device.
3. On the **Select target devices** panel:
 - a. Switch to the **Drives** tab and select your target drives.



Selecting target devices.


b. Switch to the **Files** tab and click **Create File**.



Creating target image file.


4. In the **Select image file** window, open the folder on the server where you want the file to be created and click **Create file**.

5. In the **Create image file** dialog, enter a name for the file, select its type (E01, aff4, dd, img, or raw), and then click **Create**.



Entering the name of the file.


6. On **Select target devices** panel, click **Continue**.
7. On the imaging settings page, double-check all settings and the targets selected for the imaging session, and then click **Start**.



The imaging settings page.


In the imaging page, there are two diagrams that show the progress of imaging. The upper one is called *imaging map bar* and shows imaging progress throughout the whole drive space (all successfully imaged sectors on the source drive are marked green, all damaged ones are marked red). The lower diagram is called *read speed graph* and shows the time TaskForce spent reading sectors on the source drive.

Overall imaging speed is always limited by the slowest device: either by the read speed of the source or the write speed of the slowest target.




Imaging in progress.

When imaging is completed, you are redirected to the imaging summary page, where you can review the details of the session including source and target drive details, imaging settings, hash values and the time when imaging session started and when it was completed.



Enabling post-hashing of targets.

Hashing of source drive during imaging is a preferred option because it only requires the data on the evidence drive to be read once, for both imaging and hash calculation. This ensures both a forensically sound process and minimal impact to potentially unstable media. Hashing during imaging does not slow down imaging process.



Imaging progress.

Once imaging is completed, post-hashing begins immediately on both target devices:

1

CASES

REPORTS

Atola
TECHNOLOGY

DEVICES

Diagnose

Image

Logical

RAID

Wipe

Other

Post-hashing

76%

Speed:520 MB/s

Last calculated sector:471 728 127

Estimated time:12 m

Targets:

ST320LT014-9YK142 W0Q326G4320 GBSATA 4

ST320LT009-9WC142 W0Q4NBT5320 GBSATA 5


Date ↑	Message
2023/07/27 16:52:29	Imaging started
2023/07/27 18:24:18	Pass #1 completed
2023/07/27 18:24:18	Imaging completed
2023/07/27 18:24:18	Target ST320LT009-9WC142 W0Q4NBT5: Post-hashing started
2023/07/27 18:24:18	Target ST320LT014-9YK142 W0Q326G4: Post-hashing started

PAUSE

Post-hashing in progress.

In the end, TaskForce produces a report that documents hashes of both source and target devices:


3. On the **Select target devices** panel, switch to the **Files** tab and click **Create file**.
4. In the file selector, find the folder to store the image and click **Create file**.
5. In the **Create image file** dialog, select the E01 file type.
6. Fill in E01 file information, and then click **Create**.
7. On the **Select target devices** panel, click **Continue**.



Creating an E01 file.

Enable dual-hash and start imaging

1. Once you have selected the source drive and created the target file, you end up on the **Settings** summary page. To adjust the imaging settings, click **Change**.
2. On the **Hashes** tab, make sure that **Hash source during imaging** is selected, also select both MD5 and SHA1 hash types.
3. To proceed with imaging, click **Start**.



Adjusting imaging settings.

The report and the E01 file

Upon completion of imaging, you can see both MD5 and SHA1 hash values indicated in the **Imaging completed** report.

CASES

REPORTS

DEVICES

Diagnose

Image

Logical

RAID

Wipe

Other

Samsung SSD 860 EVO 250GB S3YKNS0N625541K
250 GB

SATA 6

Case ID:
957

Creation date: 09/22/2023 7:59 PM

TaskForce serial: 33688108

Location: SATA 6

Device model: Samsung SSD 860 EVO 250GB

Size: 250 GB (250,059,350,016 bytes)

Case investigator: Andy F. Mikelson

Software version: 2023.4.2

TaskForce IP: 10.0.0.105

Write protection: On

Device serial: S3YKNS0N625541K

Device firmware: RVT04B6Q

Case description:

Imaging completed

Targets

Samsung SSD 860 EVO 250GB 250 GB S/shared/
S3YKNS0N625541K.E01
Evidence number: (empty)
Investigator: Andy F. Mikelson
Description: (empty)

Passes:

1 passes with user customized settings

Scheduled sectors:

488,397,168

Imaged sectors:

488,397,168

Errors:

0

Source sector size:

512

Hash method:

Linear

MD5

0407dc0777da8f11d3e789fae2777da

SHA1

f71c7455dde72b41152edb5fb3732ac4a27c3c6

Calculated range:

0 - 488,397,167

Log


PRINT

GO TO CASE

GO TO IMAGE

< >

It is also possible to look up the information of the created E01 file. To do that, perform the following actions:



E01 file with calculated MD5 and SHA1 hashes.

Imaging to a compressed E01 file on a target drive

In TaskForce, you can image source drives to target files on a storage device connected directly to one of the imager's drive ports. E01, AFF4, and RAW file formats are supported. E01 and AFF4 files can be compressed.

To create multiple target images on a directly connected drive, first, you need to [set this device to the Storage mode](#).


If a drive is already in the Storage mode, plug it into one of the TaskForce drive ports and [use it to store target image files](#) without additional formatting. You can image several evidence devices to target files on storage **simultaneously**.

The drive in the Storage mode is marked with a special blue icon on the **Select target devices** panel.

Set a device to the Storage mode


To set a target device to the Storage mode:

1. Go to **Image**.
2. Select the source evidence drive.
3. On the **Select target devices** panel, switch to the **Files** tab and click **Create file**.




Selecting a target file.

4. In the Select image file window, click **Add storage**.




Adding a storage drive.

5. On the **Select device** panel, choose the drive you want to use in the Storage mode. TaskForce uses a lighter shade of blue to indicate that a storage drive is being configured.




Selecting a storage drive.

6. If TaskForce cannot find the appropriate exFAT partition on the selected drive, it offers you to format the device accordingly. In that case, select **Format device to exFAT** and click **Next**.



Formatting the target to exFAT.


7. To launch target device formatting to exFat with a large cluster size (32 MB), click **Format** and enter **YES** for confirmation. This cluster size enables faster imaging to this drive.



Formatting the target to exFAT.

Once the target device is formatted, TaskForce perceives it as a Storage target and you can [proceed to create a compressed E01 image file](#) on the storage device.

The drive in the Storage mode is marked with a special blue icon on the **Select target devices** panel.




A drive in the Storage mode.

Image to a Storage drive


If you already have a drive configured as a Storage, you can plug it into one of the TaskForce drive ports and use it to store target image files without additional formatting:

1. Go to **Image**.
2. Select the source evidence drive.
3. On the **Select target devices** panel, switch to the **Files** tab and click **Create file**.



Selecting a target file.

4. On the bottom left in the **Select image file** window, select your Storage drive. Then [proceed to create a target image file](#).




Selecting a target file.

Create a compressed E01 image file


To proceed with creating a compressed E01 image file on the storage device, do the following:

1. In the **Select image file** window, select the storage drive and then click + **Create file**.




Adding a new file.

2. In the **Create image file** dialog, enter the file name and select the E01 file type.



Creating an image file.


3. Select the **Compress E01** option.
4. Fill out other file details, and then click **Create**.



Configuring the compressed E01 file.

5. Check your imaging settings and click **Start** to proceed with imaging.

When you select the **Compress E01** option in the imaging settings, the multipass imaging system or reverse imaging option cannot be applied to such an imaging session. However, other fine-tuning options remain available including advanced hashing options (pre-hash, post-hash, segmented hashing, etc.) and selective imaging.



The imaging settings screen.

Image several devices to a Storage drive in parallel

You can run multiple imaging sessions to E01/AFF4/RAW files on a Storage drive simultaneously.

To do that, follow the steps described in the [Image to a Storage drive](#) section and then [Create a compressed E01 image file](#) for each evidence device.

Imaging completed report

The **Imaging completed** report provides all the time stamps, hash values, and hash verification result. To look up the settings of the imaging session, you can also see the **Imaging started** report in the [case management system](#).




- Is an open-source format: you can describe it in a court.
- Supports multipass imaging.
- Offers fast compression methods: Snappy and LZ4.
- Supports block hashes.
- Stores binary zeroes as spans similar to sparse files.
- Is vendor-neutral.

AFF4's block hashes are calculated for small segments of data on the drive and are stored in a table inside AFF4 metadata. There is a Block map hash that represents a single SHA-512 hash value for all the individual block hashes based on Merkle tree model. This is great for imaging of damaged drives to a file using TaskForce's [multipass imaging algorithms](#).


Create an AFF4 file as an imaging target

1. In the TaskForce main window, click **Image**.
2. Select a source device.
3. Optional: On the **Imaging sessions** page, click **Start new**.
4. On the **Select target devices** panel, switch to the **Files** tab and click **Create file**.



Creating a file as an imaging target.

5. Choose a folder for your file. You can save a file:
 - on a computer or server in your local network or
 - on a storage device connected to your TaskForce.
6. After selecting a folder, click **Create file**.
7. In the **Create image file** dialog, change file **Type** to AFF4.
8. Define hashing type and compression algorithm for your target AFF4 file:
 - Block hashes: choose from MD5, SHA1, MD5 + SHA1, SHA256, or SHA512
 - Compression: choose from LZ4, Snappy, or Not compressed.
9. Enter other case details and then click **Create**.




The settings for AFF4 target image file.

10. Click **Continue**.

11. Check imaging settings, change them if needed, and then click **Start**.

TaskForce starts imaging your evidence device to an AFF4 target file. After imaging is finished, the system shows an **Imaging completed** report.



The Imaging completed report with an AFF4 file as a target.

Image a remote drive using the iSCSI protocol

Some devices can't be plugged directly into the TaskForce hardware unit for imaging. These could be drives soldered into a motherboard, servers that can't be turned off, or devices you have legal access to but lack the right to seize.

To image such devices remotely with TaskForce, use the iSCSI network protocol.

To set up an iSCSI target correctly and expose a physical or logical drive via iSCSI on a network, you can utilize a Python script provided by Atola that automatically creates iSCSI targets for all drives except for a boot device.

Automatically create iSCSI targets

To expose a physical or logical drive via iSCSI on a network, you need first set up the iSCSI target correctly. To help you with that, Atola engineers created a Python script named **iscsi-targets**, which automatically creates iSCSI targets for all drives except the boot device.

[Download iscsi-targets from GitHub →](#)

Features of the 'iscsi-targets' script

- Automatically creates iSCSI targets for all drives except the boot device.
- Ensures that the iSCSI Qualified Name (IQN) of every iSCSI target includes the drive model and serial number. When you add such an iSCSI target in Atola imagers as a source drive, the imager's software automatically extracts the model and serial number from the IQN into the case details.
- Lets you specify a block device as a script argument to create an iSCSI target only for that device.

What you need to run the 'iscsi-targets' script

The script runs on **Linux only**. It was tested on various distributions such as Ubuntu, Fedora, CentOS, and RHEL, including DFIR boot environments like Paladin, Caine, and Tsurugi.

1. **Python 3.6 or higher** must be installed.
2. On the first run, the script automatically checks for and installs two required dependencies:
 - *targetcli*
 - *python3-rtslib*

How to use the 'iscsi-targets' script

Here are some examples of using the 'iscsi-targets' script.

1. Create iSCSI targets for **all drives** except the boot device:
``sudo python3 iscsi-targets.py``
2. Create a **single iSCSI target** for the specified /dev/sdb1 partition:
``sudo python3 iscsi-targets.py /dev/sdb1``

The example below shows the first run of iscsi-targets.py on Paladin. It has added 3 iSCSI targets for SATA and USB drives.


```
paladin@PALADIN:~/Downloads$ sudo python3 iscsi-targets.py
targetcli is not installed. Do you want to install it? (Y/N): y
Boot device: /dev/sdc
The following block devices will be unmounted and used for iSCSI target creation:
/dev/sda      256G-NVMe-T      0123456789AC
/dev/sdb      HyperX-Fury-3D-240G  50026B77840022F
/dev/sdd      Samsung-Portable-SSD-T  S3SWNV0M400263
Do you want to continue? (Y/N): y
iSCSI target created for /dev/sda with IQN: iqn.2024-05.com.atola:256g-nvme-t-0123456789ac
iSCSI target created for /dev/sdb with IQN: iqn.2024-05.com.atola:hyperx-fury-3d-240g-50026b77840022f
iSCSI target created for /dev/sdd with IQN: iqn.2024-05.com.atola:samsung-portable-ssd-t-s3swnv0m400263
paladin@PALADIN:~/Downloads$
```

Automated iSCSI target creation in Paladin.

Image remote drives in parallel using iSCSI


To image a remote drive in TaskForce using the iSCSI protocol, follow these steps:

1. [Expose](#) a physical or logical drive via iSCSI on a network.
2. In the TaskForce window, click **Image**.
3. Scroll down to the **Remote** section and click **Connect iSCSI target**.
4. Enter **Target IP or DNS** and **Target port**. If needed, also enter a username and password for remote authentication.
5. Click **Discover**.



Discovering the iSCSI target.

6. TaskForce searches for and lists all iSCSI devices available at the IP address and port you provided. Select your iSCSI target, enter a username and password if required and click **Login**.



Selecting the iSCSI target.

7. **Optional:** Select the LUN if multiple devices are exposed.
8. Enter or correct the device model and serial number.
9. Click **Connect**.
10. TaskForce opens the selected iSCSI device as a source. You can now image this device as usual.

Templates for target image files

If you image evidence devices to target files on a network server or storage drive, you can benefit from highly customizable and flexible target file templates.

Thanks to configurable variables, such as the current date, case ID, device serial number, and more, a template lets you create complex subfolder paths and file names with a couple of clicks, according to your internal naming conventions.

Additionally, you can set up other target parameters, such as file type, compression, hashing method, and autofill repeating case details, such as case number, investigator name, or location.


Once created, a template for a target image file can be applied with a single click.

You can use templates when performing physical or [logical imaging](#), as well as in [Express mode](#).

Create template


To configure a template for target image files, do the following:

1. On the Taskbar, click **Image** and select your Source device.
2. On the **Select target devices** panel, switch to the **Files** tab and in the **Templates** section click **Add Template**.



Adding a template from the 'Select target device' panel.

3. The **Add template** dialog opens.




The 'Add template' dialog.

4. To edit the template's name, click the **Pencil** icon next to it.
5. Specify a **Target folder** by clicking **Select folder** and choosing the desired location.
6. Set up a **Subfolder path** and **Name** template for your target files using constants (such as department name) and available variables:
 - a. Year: **YY** (two-digit format) or **YYYY** (four-digit format).
 - b. Month: **MM** (two-digit format).
 - c. Day: **DD** (two-digit format).
 - d. Current time: **hh-mm** (24-hour format).
 - e. Case ID: **CASE_ID** (as specified in the TaskForce case system).
 - f. Evidence ID: **EVIDENCE_ID**.
 - g. Drive model: **MODEL** (acquired automatically by TaskForce during identification).
 - h. Drive serial number: **SERIAL** (acquired automatically by TaskForce during identification).
 - i. Ordinal number variable: **{NUM+}** (increments from the number you specify in curly brackets if the file or folder with the same name already exists).

Variables are highlighted with **pale green** for better visibility.

To create **nested folders** in the specified Target folder, use the forward slash symbol (/) in the **Subfolder path** field.

7. Select the target file **Type**: E01, AFF4, IMG, DD, or RAW.




Selecting target file type for a template.

8. Optional: Configure additional parameters of a target file:

- For E01 files: **Segment size** and **Compression**.
- For AFF4 files: method for **Block hashes** (MD5, SHA1, MD5 + SHA1, SHA256, SHA512) and **Compression** algorithm (LZ4 or Snappy).


9. Choose how to fill in file metadata: automatically (default) or manually.

If you choose **Automatically**, specify the values for Case ID, Investigator's name, Evidence number, Description, or Notes.



Option to fill in file metadata automatically.

If you choose **Manually**, TaskForce will ask you to enter file metadata whenever you create a target file using the template.



The 'Set file metadata' dialog.

10. When all the parameters of the template are set, click **Create**.


Use template

To create a target image file using an existing template, after selecting a Source device, simply click your template on the **Select target devices** panel.

After that, TaskForce creates the required subfolder(s) and adds a file with the specified parameters as an imaging target.

The target file appears on the **Files** tab of the **Select target devices** panel.


Then you can proceed to the Imaging settings screen by clicking **Continue**.



A template for target image files on the 'Select target device' panel.

Edit or delete template

To edit or delete a template for target image files, use the **three-dots** icon next to the template's name on the **Select target devices** panel.




Editing or deleting a template for target image files.

Accessing password-protected servers

Accessing password-protected servers allows saving image files on such servers, imaging or calculating hash of the files located there, etc.

To create an image file on a password-protected server:


1. In the TaskForce main window, click **Image**.
2. Select a source device.
3. **Optional:** On the **Imaging sessions** page, click **Start new**.
4. On the **Select target devices** panel, switch to the **Files** tab and click **Create file**.
5. In the **Select image file** window, choose the server from the list. If the server does not appear in the list, click **Refresh** icon to search for all available directories. If the server still does not appear in the list, click **Connect**.



Selecting a server from the list.


6. The **Connect to server** dialog opens. Enter the name of the server and fill out login details, including domain or workgroup, and then click **Connect**. To learn these details, contact your network administrator.

TaskForce supports NTLM protocol, which is enabled by default in Active Directory and Windows.




Entering server and login details.

7. Go to the folder on the server where you need to store your image and click the **Plus** icon in the right bottom corner of the window.




Creating a new file in the directory.

8. Enter the name of the new image file and select the format. Click Create.




Entering the name and selecting the format of the new image file.

9. On the **Select target devices** panel, click **Continue**.




The Continue button on the Select target devices panel.


10. Check your imaging settings and click **Start** to proceed with imaging.



6. On the **Select device** panel, choose the drive connected to a port in the *Target* mode.
7. Select **Create an encrypted VeraCrypt container (exFAT)** and click **Next**.




8. Enter and confirm the password for the encrypted volume on the drive. Then click **Create**.




9. Confirm the formatting of the device by entering **YES** and clicking **OK**. After this step, the formatting takes a few seconds.

10. Click + **Create file**.

11. Enter the name of the image file and choose the file format (E01, raw, img or dd). Then click **Create**.




12. Once you have created the file, you may add more image files in the same or a different folder.




After you click **Continue**, TaskForce images the evidence into the file on your encrypted target.

Upon completion of the imaging session, check the **Imaging completed** report.




After you enter the password, the volume is mounted and you can access it from Windows Explorer and use the image for subsequent operations.




Restore E01, AFF4, RAW image file to drive

To restore data from an image file to a drive using Atola TaskForce, follow these steps:


1. In the TaskForce main window, click **Image**.




2. On the **Select source device** panel, expand the **File** section and then click **Select file**.



3. Select the E01, AFF4 or RAW file you are planning to restore.




4. Select the target drive and click **Continue**.




If your target drive is larger than the size of the image, the hash values of the target drive will not be identical to that of the image file. With TaskForce, you can limit a SATA target drive's capacity via Host Protected Area (HPA) or Accessible Max Address (AMA). Applying the HPA or AMA will set a new max address in the drive's firmware, which will make the sectors beyond the required capacity inaccessible. This will make the hashing process straightforward and the hashes will match.

To [limit SATA target drive's capacity via HPA or AMA](#), go to the imaging **Settings** and click the **Change** button.



Go to the Miscellaneous tab and enable the Limit target disk size to source size using HPA/AMA option.



5. Click **Start** for the imaging to begin.

1

Home

CASES

REPORTS

Atola
TECHNOLOGY

DEVICES

≡

Diagnose

Image

Logical

RAID

Wipe

Other

LongFileNamesInNTFS.E01

512 GB

//S:/Images/NTFS

Case ID:

Not assigned

1 target

▼

Imaging data...

77%

Pass: 1 of 1

Speed: 109 MB/s

Errors: 0

Sectors imaged: 779 132 928

Sectors scheduled: 1 000 236 032

Estimated time: 17 m

ms

4

2

778 501 632

777 388 288

777 934 592

778 499 072

Log

Target HEX viewer

Signatures : 729 858

Date ↑

Message

2023/08/30 14:14:09

Max address for target device ST3000DM001-9YN166 Z1F168R5 was successfully set

2023/08/30 14:14:23


Imaging started

PAUSE

VIEW SETTINGS

TaskForce automatically creates detailed reports for every session. The imaging report lists all the details of the source, the target, the imaging settings and timestamps including the setting of the new max address of the target drive. This makes such data extraction from a file transparent and forensically sound.

2. Select a source and target device.
3. On the **Imaging sessions** page, click **Start new**.
4. On the **Settings** page, click **Change**.
5. On the **Miscellaneous** tab, toggle the **Limit target disk size to source size using HPA/AMA (SATA target ports only)** option.



Enabling HPA/AMA restriction for target.

You can now proceed with the imaging process by clicking the **Start** button.

Before the imaging starts, TaskForce looks up the size of the evidence drive and limits the space of the target using HPA/AMA to make its capacity identical to that of the evidence drive.

When imaging is complete, the report will contain information about the time when HPA/AMA was enabled.

HOME

CASES

REPORTS

Atola
TECHNOLOGY

DEVICES

Diagnose

Image

Logical

RAID

Wipe

Other

ADATA USB Flash Drive 21B03215801700B0

62 GB

USB 3

Case ID: Not assigned

Creation date: 08/30/2023 4:50 PM

TaskForce serial: 33688108

Location: USB 3

Device model: ADATA USB Flash Drive

Size: 62 GB (62,075,699,200 bytes)

Case investigator:

Software version: 2023.7

TaskForce IP: 10.0.0.246

Write protection: On

Device serial: 21B03215801700B0

Device firmware: 1100

Case description:

Imaging completed

Targets

KINGSTON SUV400S37120G 50026B726A041902 120 GB SATA 5

Passes:

1 pass with default settings for healthy drives

Scheduled sectors:

121,241,600

Imaged sectors:

121,241,600

Errors:

0

Source sector size:

512

Hash method:

Linear

MD5

f3ec16a870e564f7926d6c4fc2511be7

Calculated range:

0 - 121,241,599

Signatures:

4,152

Top 10 found signatures

File descriptor	Bytes	Count
NTFS File Record	46 49 4C 45	1708
NTFS Index Record	49 4E 44 58	1076
Unicode XML/HTML Document	3C 3F 78 6D 6C	951
Perl Module	70 61 63 6B 61 67 65	233
Adobe Extension Package	03 00 00 00 01 00 00 00	85
Adobe Photoshop Actions Settings File	00 00 00 10 00 00 00	31
Portable Network Graphic	89 50 4E 47 0D 0A 1A 0A 00 00 00 0D 49 48 44	17
Unicode XML/HTML Document	EF BB BF	13
ClickOnce Deployment Binary Manifest File	50 63 6D 48 01 00 00 00 00 00 00 00	6
MP3 Audio File	49 44 33	6

Found signatures.csv (Internal storage)

Log

Date ↑

Message

2023/08/30 14:33:01

Max address for target device KINGSTON SUV400S37120G 50026B726A041902 was successfully set

2023/08/30 14:33:08

Imaging started

2023/08/30 15:50:11

Pass #1 completed

2023/08/30 15:50:11

Imaging completed

PRINT

GO TO CASE

GO TO IMAGE

⏪

⏴

⏵

⏩

The Imaging report indicates the change to the target drive capacity.

The target disk's port in Devices menu now contains an HPA/AMA indicator, thus informing you that HPA/AMA has been enabled on this drive.

Home

CASE

Select device

Power on all devices

Diagnose

Image

Logical

RAID

Wipe

Other

SATA

1

WD2500AAKS-00F0A0
WCAT1F625038
62 GB - limited via HPA

Case: Not assigned

target

3

ST320LT014-9YK142
W0Q326G4
320 GB

target

5

KINGSTON SUV400S37120G
50026B726A041902
62 GB - limited via HPA

Case: Not assigned

target

2

WDC WD10JPVT-75A1YT0
WXF1C12Y0097
1 TB

target

4

HGST HTSS45050A7E680
RB697AMBG0ULNN
500 GB

target

6

ST16000NM001G-2KK103
WL2013X5
16 TB

target

SAS

1

SEAGATE ST3300657SS-H
3SJ372DS
295 GB

target

3

Device not detected

target

5

Device not detected

target

2

Device not detected

target

4

SEAGATE ST373455SS
3LQ1GY2C
73 GB

target

6

Device not detected

target

USB

IDE

EXT

FILE

RAID

The HPA indicator in the port of the Device menu.

There will also be a report created in the case management system, which indicates the old (native) and the new (as set by HPA/AMA) max address.

The screenshot shows the Atola Technology interface. On the left is a sidebar with icons for Diagnose, Image, Logical, RAID, Wipe, and Other. The main area displays a report for a Kingston SUV400S37120G 62 GB SATA 5 drive. The report includes the following details:

Creation date:	08/30/2023 4:33 PM	Software version:	2023.7
TaskForce serial:	33688108	TaskForce IP:	10.0.0.246
Location:	SATA 5	Write protection:	
Device model:	KINGSTON SUV400S37120G	Device serial:	50026B726A041902
Size:	62 GB (62,075,699,200 bytes)	Device firmware:	0C3J96R9
Case investigator:		Case description:	

Device size was reduced by HPA

This action was initiated by an Imaging setting (automatically clip target device to the size of source).
Current device was limited to the source size by HPA.

Old max address: 234,441,647
New max address: 121,241,599

At the bottom of the report, there are buttons for 'PRINT' and 'GO TO CASE', and navigation arrows.

The report about HPA activation.

Now you can calculate hash on both drives to make sure the hash values are identical.

Enabling HPA/AMA is an option available only for SATA target drives.


To learn how to unclip HPA/AMA, read [Unclip or change HPA, DCO, AMA limitations](#) in our manual.

Exporting sector lists from an imaging session

When an imaging session is completed or paused, it is possible to see its summary in the **Imaging sessions summary** page. Now there is also a possibility to export lists which would clearly indicate which of the sectors on the source drive have been successfully imaged, which have not (if any), and which of the sectors contained errors.

To export such list:

1. Go to **Image** and select the source drive.
2. In the session summary, click the **Export** icon.



The Export icon in the session summary.

3. Select the sectors you are interested in (for example, Imaged sectors).

Atola TECHNOLOGY

Diagnose
Image
Logical
RAID
Wipe
Other

ST9500423AS S2V0EP3P
500 GB

Not diagnosed
SATA 3

Case ID:
Not assigned

Recent imaging sessions

ST9500423AS S2V0EP3PE01	39%	Scheduled: 976,773,168 out of 976,773,168 Imaged: 386,172,235 Errors: 248
-------------------------	-----	---

Export to file
Imaged sectors
Non-imaged sectors
Sectors with errors

START NEW

Link: PHY Status: ERR INDX CORR DREQ DSCP TAULT DDDV BUSY Error: AMAN TONF ABIT - IDNF - LINC ICRC

Selecting the type of sectors.

4. Save the downloaded .csv file. This file shows the ranges of imaged sectors:


	A	B	C	D	
332	412934912	412959487			
333	413959488	413984063			
334	414984064	415008639			
335	416008640	416033215			
336	417033216	417057791			
337	418057792	418086463			
338	419086464	419106943			
339	420106944	420123327			
340	421123328	421139711			
341	422139712	422156095			
342	423156096	423168383			
343	424168384	424176575			
344	425176576	425180671			

Should an imaging session be completed, the list of non-imaged sectors will be blank.

Imaging presets

For examiners who need to ensure they are using specific imaging settings for certain types of drives or cases, TaskForce allows creating different presets for easy, one-click switching to a specific imaging routine. Presets also allow these specific imaging settings to be shared with colleagues who use another TaskForce by exporting presets from one device and importing them onto another one.

TaskForce has two presets called *Default* and *Damaged* recommended for healthy and faulty drives respectively.




Default presets.

Manage custom presets

To create a custom preset:

1. Go to **Image** and select source and target devices.
2. On the **Imaging settings** page, click **Change**.
3. Change the imaging settings.
4. Click the three-dot icon in the bottom right corner and choose **Save to**.
5. In the dialog window, enter the name of the preset and click **Save**.




Saving a preset.

Once you have saved a preset, they are stored in TaskForce's work folder and can be used by other operators using the same imager. They can find the presets you created under the **Custom** button.

The Chrome browser of the user who created a preset, saves a copy of the preset locally. When this user opens the **Custom** menu, both presets are displayed:

- one is the preset saved in TaskForce's work folder,
- another one is the local copy.

If one becomes redundant, it can be easily deleted.



All custom presets.

To delete a preset:

1. Select the preset in the **Custom** menu.
2. Go to the three-dot menu in the bottom-right corner and click **Delete**.


Share presets with another TaskForce

Once you have created your custom presets, you can share them with colleagues who use another TaskForce imager.

To export a preset:

1. Select the preset from the **Custom** menu.
2. Click the three-dot icon.
3. Click **Export**.


The preset will be downloaded in .json format.



Exporting a preset.

To import a preset:

1. Click the three-dot icon.
2. Click **Import**.
3. In the **Import settings** window, click **Select file**.
4. Find the file in the file selector and click **Open**.
5. Double-check the preset and click **Import**.



Importing a preset.

After this import, the preset can be found in the **Custom** menu.

Imaging a drive with a damaged head

The diagnostics module, selective head imaging and multipass imaging algorithm allow Atola TaskForce to handle a drive with damaged heads gently and effectively. All these techniques help minimize the risk of losing more data on the working part of the head stack.

Diagnose first

The built-in diagnostics module of TaskForce automatically checks all major subsystems of the evidence drive: circuit board, heads, media surface, firmware and file system.

A diagnostics report provides detailed information about the heads. In addition, it offers recommendations for the optimal imaging strategy for your damaged hard drive.

[Diagnostics completed](#)

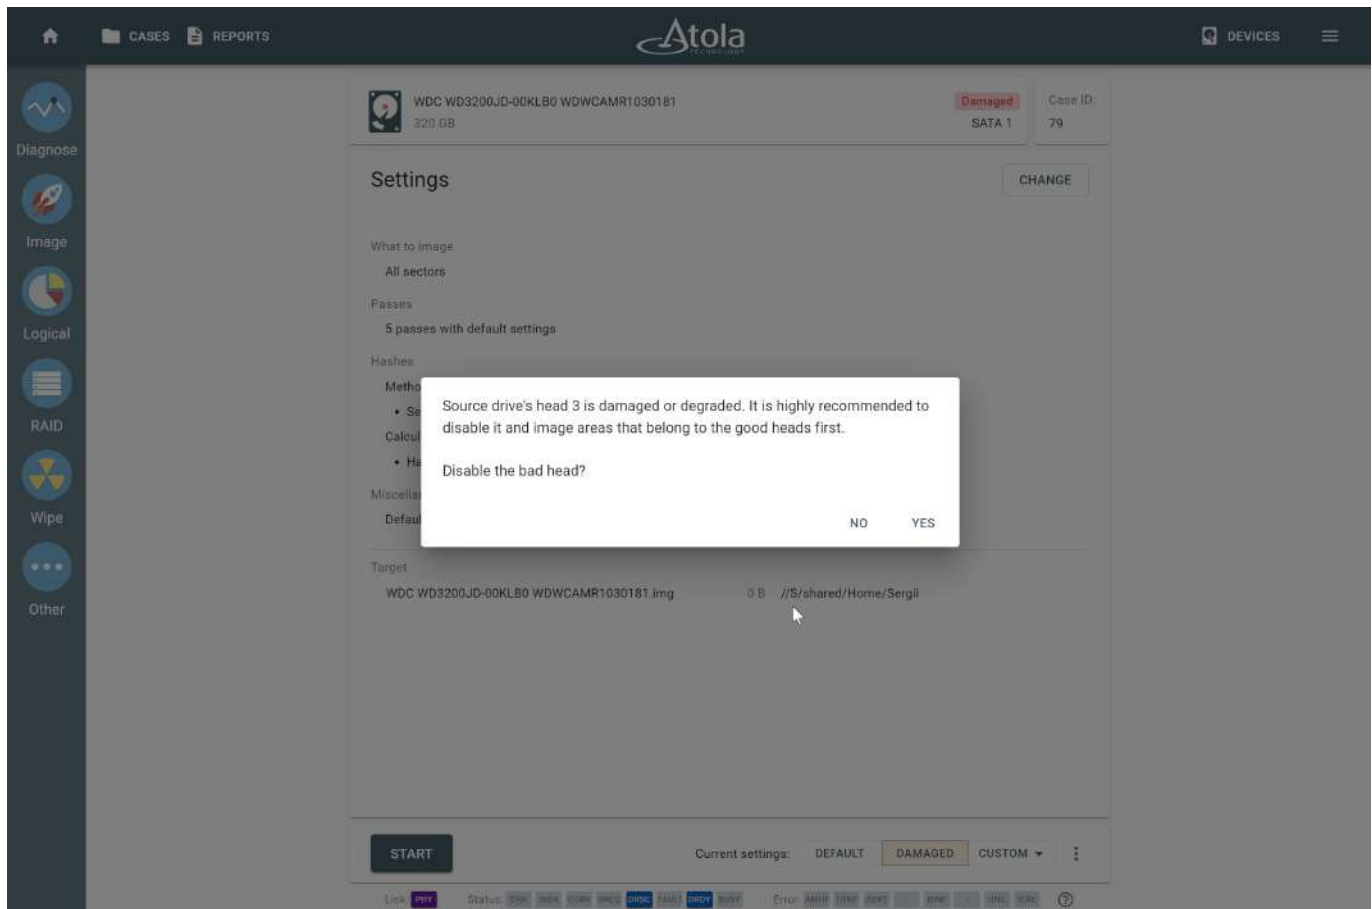


The above diagnostics report informs the operator that the drive's hardware has major issues and points to defects in the media and a damaged head (Head#3). The report contains a recommendation to disable the damaged head in the imaging settings.

Selective head imaging


Atola engineers recommend that the good heads are imaged first. To do that:

1. Go to **Image**.
2. Select your source and target devices.
3. Click **Continue**.
4. If a head was identified as damaged during the diagnostics, at this stage TaskForce shows a dialog prompting you to disable the damaged head. To confirm that the head should be automatically disabled for the subsequent imaging session, click **YES**.




Alternatively, the damaged head can be disabled in the imaging settings:

1. On the **Imaging settings** page, click **Change** to adjust the settings for your imaging session.
2. In **What to image** section, click on **All sectors** to configure the selective imaging.




3. Unselect the damaged head and click **Save**.
4. To launch your imaging session, click **Start**.



Multipass imaging algorithm

As you can see in the screenshot below, some errors were found in the course of imaging on the space of the drive that is read with the Head#4. It is common for a drive with a bad head to also contain errors on the platters that are read with other heads.



TaskForce uses its [multipass imaging algorithm](#) when encountering a bad sector that belongs to a good head. It allows handling errors and retrieving data from some of the bad sectors. For as long as it is possible to read data from the sector or block of sectors within the specified pass timeout, TaskForce will be able to image this data.

Having completed imaging from the good heads, the system pauses the session and produces a detailed imaging report with a log of all actions performed during the imaging session.

TaskForce allows editing settings of all unstarted imaging passes, adding or removing passes, etc. So if later you think you may be able to retrieve important data with Head#4, you can add another pass and configure the settings of the new pass accordingly.

TaskForce automatically inserts and launches an additional imaging pass after you click **Resume** on the pause session. The new pass will include all non-imaged sectors.

- To image the unselected bad head, simply click **Resume**.

Atola
TECHNOLOGY

CASES

REPORTS

DEVICES

Diagnose

Image

Logical

RAID

Wipe

Other

WDC WD3200JD-00KLB0 WDWCAMR1030181
320 GB

SATA 1

Case ID:
Not assigned

Recent imaging sessions

ST500DM002-1BD142 S2AB7G2C

83%

Scheduled: 625,142,448 out of 625,142,448

Imaged: 520,961,167

Errors: 131

RESUME

START NEW

Link: **PHY**

Status: **ERR** **INDX** **CORR** **DRSC** **FAULT** **DRDY** **BUSY**

Error: **AMNF** **TONF** **ABRT** **DRNF** **UNCG** **ICRC**

TaskForce resumes the imaging session to focus on the area that belongs to the damaged head.

1

HOME

CASES

REPORTS

Atola
TECHNOLOGY

DEVICES

≡

Diagnose

Image

Logical

RAID

Wipe

Other

WDC WD3200JD-00KLB0 WDWCAMR1030181

320 GB

SATA 1

Case ID:
Not assigned

1 target

Imaging data...

LEGEND

83%

Pass: 6 of 6

Speed: 220 B/s

Errors: 382

Sectors imaged: 520,961,167

Sectors scheduled: 625,142,448

Estimated time: 2808 d 12 h

Head number: 3

ms

4

2

112,437

112,687

Date ↑

Message

2023/09/01 11:40:59

Can't read sector 112,682 (AMNF). Head 3

2023/09/01 11:41:01

Can't read sector 112,683 (UNC). Head 3

2023/09/01 11:41:03

Can't read sector 112,684 (UNC). Head 3

2023/09/01 11:41:06

Can't read sector 112,685 (UNC). Head 3

2023/09/01 11:41:08

Can't read sector 112,686 (UNC). Head 3

2023/09/01 11:41:10

Can't read sector 112,687 (AMNF). Head 3

PAUSE

VIEW SETTINGS

Link: PHY

Status: ERR

INDEX

CORR

DRSC

FAULT

DRDY

BUSY

Error: AMNF

TORQ

ABRT


DRNF

UNC

ICRC


?

If the number of errors keeps growing, while the number of the imaged sectors remains unchanged, pause your imaging session and power down the drive because the head seems to be completely inoperable.



In the **Imaging report** above, you can see that TaskForce imaged 520,961,167 sectors out of 625,142,448, having extracted as much data from good heads as was possible with the default settings.

For more details scroll the report down to check the **Log**:



Multipass imaging of damaged drives

Damaged media require a sophisticated imaging approach to balance out thorough data extraction with forensic units' need for expediency and careful treatment of damaged media.

Atola TaskForce has a complex imaging functionality that lets you image even physically damaged drives while avoiding causing further drive damage. You can [use a predefined setting](#) for imaging a damaged drive, or [fine-tune all imaging parameters](#) yourself.

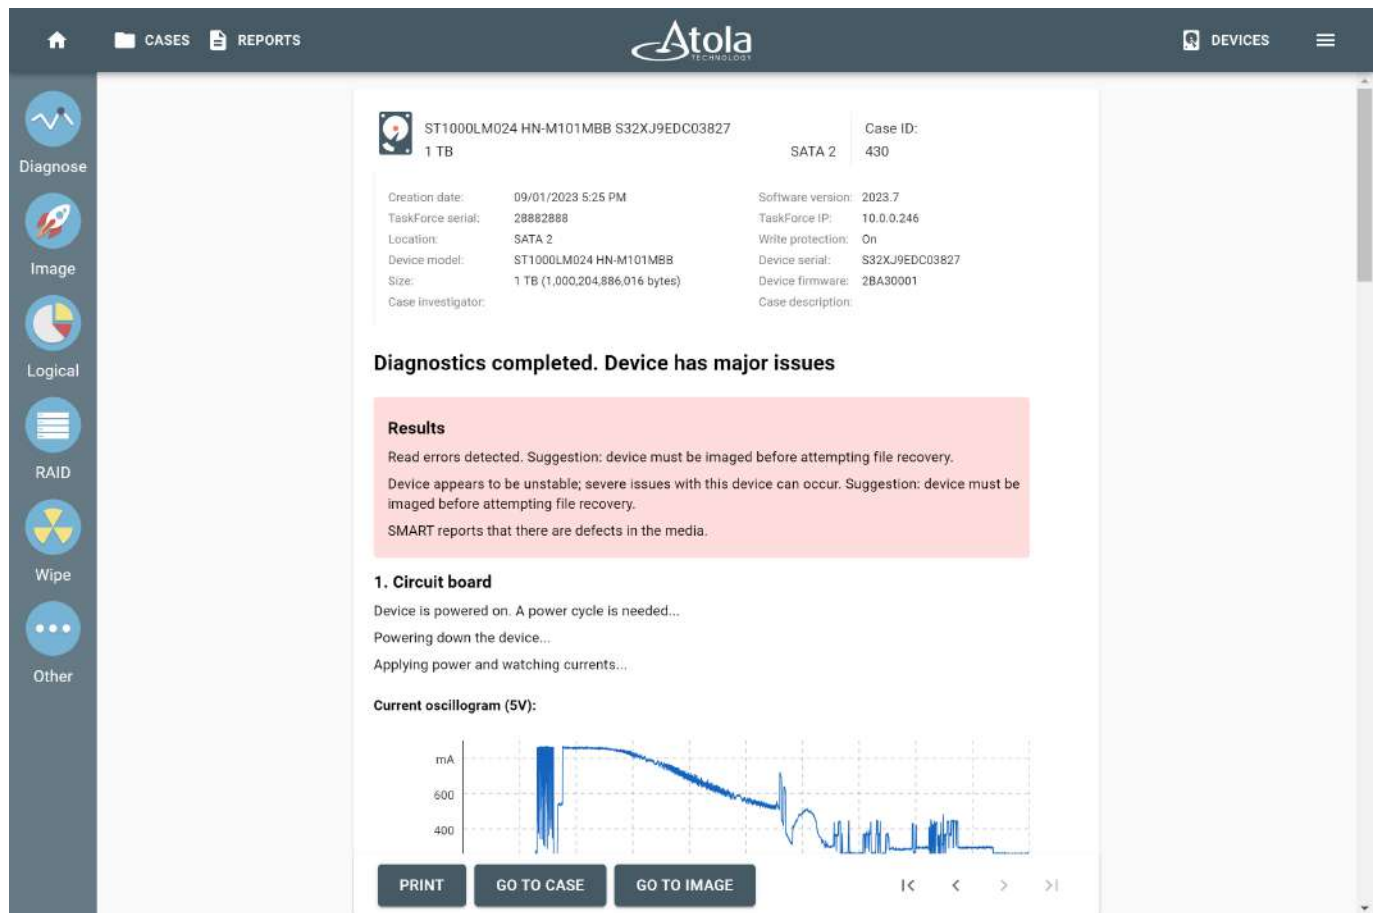
Diagnose first

Before imaging any evidence device, we strongly recommend [diagnosing](#) it.

To image an evidence drive without causing further damage to it, TaskForce is equipped with an automatic diagnostics module that evaluates the state of the drive, identifies specific errors and their location, and recommends the best approach for data acquisition.

The detailed diagnostics report contains detailed recommendations on how to acquire data from this particular drive, based on its condition.

If Automatic Checkup detects a damaged or degraded head, disable the head in the imaging settings for the initial imaging session. For details, see [Imaging drives with damaged heads](#).




The diagnostics report for a damaged device.

Preset for imaging damaged drives

If a device has issues or is damaged, TaskForce shows the respective tag next to it and suggests using the *Damaged* preset, designed specifically for imaging faulty drives.

In this case, the *Damaged* preset is automatically selected on the **Imaging settings** screen.

The default preset for bad drives is based on our decades-long experience in the data recovery market to overcome most types of damage to the drives. That's why we suggest that you use the *Damaged* preset for bad drives unless a particular drive requires a specific imaging approach.




The Damaged preset is selected on the Imaging settings screen.

To thoroughly retrieve maximum data from an unstable drive in a forensically sound way, the preset for damaged hard drives has several major differences from default imaging settings:

- **Number of passes:** 5 passes for damaged drives instead of 1 pass for good drives.
- **Different timeouts for each pass:** 1 second on the first pass, 5 seconds on the second, third, and fourth passes, and 60 seconds on the last fifth pass.
- **Jump on errors:** from 1,000,000 sectors on the first pass to only 1 sector jump on the last pass.
- **Different read block sizes for each pass:** 4,096 on the first, second, and third passes, 256 for the fourth and fifth passes.
- **Segmented hashing** with 4 GB segments instead of linear hashing used by default for good drives.

If you use an E01 file as a target along with a linear hashing method, multiple passes aren't available. In such a case, we recommend using an AFF4 or RAW file format instead.



The Damaged preset for imaging faulty drives.

Pass is a single complete cycle of reading blocks from a source device and writing them to a target device, beginning from a start sector and finishing at an end sector (as specified in the **What to image** field).

Timeout is a max time for a single read block attempt during this pass.

Jump on errors is a number of consecutive sectors that TaskForce will skip if it can't read a block from a source device.

Max read block size is a maximum number of sectors that TaskForce reads from a source device at a time.

Segmented hashing is a process of calculating hashes for a series of corresponding LBA ranges of the image. For more details, see [Segmented hashing for data verification](#).

Multipass algorithm for imaging damaged drives

To approach bad drives in the most gentle way possible, TaskForce uses its special multipass imaging system.

Most forensic imagers can only do linear imaging, which dramatically slows down the imaging process whenever a bad sector is encountered, and, as a result, the drive may freeze. To speed up the imaging of damaged media and maximize the amount of successfully retrieved data, TaskForce has a special imaging algorithm that includes a deliberate timeout and block size control.

Timeouts and block size control

Using a small block size pays off when you need to thoroughly retrieve maximum data from an unstable drive, but it also significantly slows down the imaging process. What's worse, such an imaging approach may cause further damage to the media.

That's why TaskForce's multipass imaging engine uses **large blocks with short timeouts** on the **first few passes**, scheduling reads inside slow areas for later and then using the **smallest block size on the last pass** when very few sectors are left to be read. TaskForce handles block size automatically, to provide the best possible results in the shortest time.


This technique helps achieve imaging max speeds in good areas of the drive. At the same time, it lets you approach bad areas in the most delicate way possible and retrieve as much data as possible.

First pass

On the first pass, TaskForce allows a 1-second **Timeout** per block, and the **Max read block size** is set to 4096 sectors. This allows smooth sequential imaging of all healthy modern drives.

But when imaging damaged media, these settings let TaskForce skip any areas that slow down the process and perform **Jump on error** by 1,000,000 sectors at a time.


This way all the good areas of the drive are imaged at top speed, while forcing TaskForce to return to the problematic areas on the following passes, narrowing down the bad areas and allowing more time to retrieve the data within them.



Imaging the first pass. Empty areas where errors were encountered and jumps were performed.

Second and third passes


While the **Max read block size** remains the same during the second and the third passes, the **Jump on error** is set to 20,000 sectors and 4,096 sectors respectively, and slightly longer, 5-second **Timeouts** are allowed for attempted reading of the blocks.



Empty areas start filling up with data, as the jumps become smaller.

Fourth pass

On the fourth pass, both Jump on error and Max read block size are reduced to 256 sectors to try reading problematic zones in a more granular way.



The amount of data retrieved is already 99%.


Fifth pass

On the fifth pass, TaskForce allocates 60-second Timeouts to read the **Maximum block size** of 256 with just 1-sector **Jump on error**. It is the last and the most thorough attempt to retrieve data from the remaining bad areas of the drive.

The last pass has a unique feature that is not used during previous passes: an internal sector-by-sector auto-reread procedure for an error block. It is defined by an unchangeable Jump size = 1 sector.

How the imaging engine works on the last pass:

- It reads a block using the Max Block Size pass setting (256 by default).
- If the reading is successful, it proceeds to the next non-imaged block.
- If a read error occurs, the engine re-reads the whole error block sector by sector.



On the fifth pass TaskForce attempts to read the data for the last time.

After the final pass, the Imaging Results report will indicate the eventual number of errors on the drive and other detailed statistics.

The screenshot displays the Atola software interface. The top navigation bar includes a home icon, 'CASES', 'REPORTS', the Atola logo, 'DEVICES', and a menu icon. A left sidebar contains icons for 'Diagnose', 'Image', 'Logical', 'RAID', 'Wipe', and 'Other'. The main content area shows the 'Imaging completed' report for a device identified as 'ST3160811AS 5PT0DR8A' (160 GB) on 'SATA 1'. The report includes a table of metadata:

Device Information		Case Information	
Creation date:	09/01/2023 6:32 PM	Software version:	2023.7
TaskForce serial:	28882888	TaskForce IP:	10.0.0.246
Location:	SATA 1	Write protection:	On
Device model:	ST3160811AS	Device serial:	5PT0DR8A
Size:	160 GB (160,041,885,696 bytes)	Device firmware:	3AAE
Case investigator:		Case description:	

Below the metadata, the 'Imaging completed' section shows the following details:

- Targets:** ST3160811AS 5PT0DR8A.img 160 GB S:/shared/Home/
- Passes:** 5 passes with default settings for damaged drives
- Scheduled sectors:** 312,581,808
- Imaged sectors:** 312,581,711
- Errors:** 97
- Source sector size:** 512
- Pattern for unreadable sectors (HEX):** 00
- Hash method:** Segmented
- Segment size:** 4 GB
- Hash type:** MD5
- Segmented hashes:** MD5_HashSegments-ST3160811AS_5PT0DR8A_3AAE.csv (Kingston DataTraveler Max 204479A6A72071A5080C)
- Signatures:** 144,731

At the bottom of the report, there are buttons for 'PRINT', 'GO TO CASE', and 'GO TO IMAGE', along with navigation arrows.

The Imaging Results report shows the number of errors on the drive.

Customize imaging settings for each pass


To cope with a severely damaged drive, you can adjust the following parameters of any imaging pass:

- Timeout
- Jump on error
- Max read block address
- Start and end LBA
- Image in reverse direction
- Disable read look-ahead

Also, you can [add](#) or [delete](#) an imaging pass.

To customize settings for a certain imaging pass, do the following:

1. On the **Imaging settings** page, click **Change**.



The Change button on the Imaging settings page.

2. Click the pass you want to change.

Atola
Technologies

Home CASES REPORTS DEVICES

Diagnose Image Logical RAID Wipe Other

ST3160811AS 5PT0DR8A
160 GB

Has Issues
SATA 1

Case ID:
76

Settings

BACK

PASSES HASHES MISCELLANEOUS

Pass	Timeout	Jump on error	Max read block size	What to image
1	1s	1 000 000	4 096	All sectors
2	5s	20 000	4 096	
3	5s	4 096	4 096	
4	5s	256	256	
5	60s	1	256	

+ ADD PASS

RESTORE TO DEFAULT


START

Current settings: DEFAULT DAMAGED CUSTOM

Link: PHY Status: ERR INDX CORR DRDQ DRDC PAULT DRDY BUSY Error: AMAR TORP ABST DRDE UNCL ICRC

An imaging pass is highlighted on the Imaging settings page.

3. In the **Edit imaging pass** window, enter new settings or toggle the options you need.
4. Click **Save**.



The Edit imaging pass window shows the current settings for the first pass.

Start and End LBA

For each pass, you can define the starting and ending sectors by entering logical block addresses in the respective fields (**Start LBA**, **End LBA**) or by dragging markers of the slider below.

Alternatively, you can select **All sectors with data** in the **What to image** field. It makes TaskForce search for all known partitions and image only the sectors that contain data. This option is a good use when you are facing a lack of time to take a full image.

Currently supported partitions: NTFS, ext4/3/2, exFAT, XFS, ZFS, Btrfs, APFS (with encrypted volumes), HFS/HFS+, FAT32/16.

Edit imaging pass

Pass
1

Timeout
1

Unit
seconds

Max read block size
4 096 sectors

Jump on error
1 000 000 sectors

What to image
All sectors

Start LBA
0

End LBA
312 581 807

Do not use head map

SELECT

Pattern for unreadable sectors (HEX)
00

LOAD

☐ Reverse direction ?

☐ Disable read look-ahead

CANCEL SAVE

Image in reverse direction

With this function selected, TaskForce approaches skipped areas of the drive from the other side on any selected pass. That means the imaging engine reads a source drive backward and reaches the damaged areas from the opposite direction.

This way, the imaging module can retrieve more data from a drive before entering a damaged zone, which needs to be concentrated on during the following passes. But the speed decreases due to auto disabling of the drive's cache.

Hint: It is one of the best options we recommend you enable to get more data from a severely damaged drive.

Edit imaging pass

Pass
1

Timeout
1

Unit
seconds

Max read block size
4 096

sectors

Jump on error
1 000 000

sectors

What to image
All sectors

Start LBA
0

End LBA
312 581 807

Do not use head map

SELECT

Pattern for unreadable sectors (HEX)
00

LOAD

☐ Reverse direction ?

☐ Disable read look-ahead

CANCEL SAVE

Disable the read look-ahead mode

Most contemporary hard drives have a read look-ahead functionality, which makes the drive read more blocks sequentially than requested by software.

In good drives, this functionality helps the drive operate faster by reading more data and caching it.

But with bad drives, the read look-ahead feature leads to bad areas being addressed more often. This slows down the process and may lead to a complete freeze of the drive. In such cases, we recommend disabling the read look-ahead option.

Edit imaging pass

Pass
1

Timeout
1

Unit
seconds

Max read block size
4 096 sectors

Jump on error
1 000 000 sectors

What to image
All sectors

Start LBA
0

End LBA
312 581 807

Do not use head map SELECT

Pattern for unreadable sectors (HEX)
00 LOAD

☐ Reverse direction ?

☐ Disable read look-ahead


CANCEL SAVE

Add a pass

You can add new passes even when an imaging session is not yet started or when it's paused.

To add a new pass to an imaging session, do the following:

1. On the **Imaging settings** page, click **Change**.
2. Click **Add pass**.




Adding a pass.

3. Adjust the settings for the new pass.
4. Click **Save**.

Delete a pass

You can delete any pass that has not been started yet. To delete an imaging pass, do the following:

1. On the **Imaging settings** page, click **Change**.



The Change button on the Imaging settings page.

2. Click the pass you want to delete.

Edit imaging pass

Pass
3

Timeout

5

Unit

seconds

Max read block size

4 096

sectors

Jump on error

4 096

sectors

What to image

All sectors

Start LBA

0

End LBA

312 581 807

Do not use head map

SELECT

Pattern for unreadable sectors (HEX)

00

LOAD

Reverse direction

?

Disable read look-ahead

CANCEL

SAVE

Imaging freezing drives

The core part of the imaging process is based on reading a source drive by sending multiple **Read sectors** commands and handling the drive's response at the low native I/O level.

When a drive receives and runs a **Read sectors** command but is unable to read any data from that sector. So it goes into **Retry** mode, trying to get data from the damaged area again and again.

After a certain number of tries, it gives up on a particular command and returns an error with timeout.

In TaskForce, you can detect a freezing drive [during diagnostics](#) or [during imaging](#).

Freeze detection and handling during imaging

Imagine that the drive is unable to read data from the damaged sectors and goes into a long-lasting retry mode before it gives up on a particular sector and returns an error.

If TaskForce simply waited for each **Read sectors** command to be completed:

- it would take ages to get an image of a drive with numerous errors;
- it could cause the drive to slip into complete freeze;
- in the worst-case scenario, further damage could be caused to the data on the drive.

The Reset command

To avoid causing further damage to the data on the drive and long waiting periods, TaskForce issues a **Reset** command whenever a drive attempts to read a block of sectors longer than allowed by the pre-configured timeout.

Reset is a device interface operation, using which TaskForce stops the previously sent **Read sectors** (or any other) command and then continues imaging from the next planned block on the drive.

If the device is still running the Read Sectors command, even after the first **Reset** attempt, TaskForce waits 3 seconds and performs the second **Reset** command. At the moment of the second **Reset**, a new entry appears in the imaging **Log** reading

Device hangs while reading block X – Y.

The screenshot shows the Atola TaskForce software interface. The top navigation bar includes 'CASES', 'REPORTS', and 'DEVICES'. The left sidebar contains icons for 'Diagnose', 'Image', 'Logical', 'RAID', 'Wipe', and 'Other'. The main window displays the following information:

- Device:** ST3750630AS 9QK1ZS5T, 750 GB, SATA 1. A yellow 'Has issues' badge is present.
- Case ID:** 96
- Target:** 1 target
- Imaging data...** section shows a progress bar at 7% and a legend.
- Statistics:**
 - Pass: 1 of 5
 - Speed: 77 MB/s
 - Errors: 22
 - Sectors imaged: 103 022 592
 - Sectors scheduled: 1 465 149 168
 - Estimated time: 2 h 30 m
- Log:** A table with columns 'Date' and 'Message'. The entry '2023/03/08 15:41:31 Device hangs while reading block 55,177,024 - 55,181,119' is highlighted with an orange box.
- Buttons:** 'PAUSE' and 'VIEW SETTINGS'.
- Status Bar:** Shows 'Link: PHY', 'Status: CRK', and various error codes like 'ERR', 'INCK', 'CORR', 'DNCK', 'DRSC', 'FALCK', 'DRDY', 'BUSY', 'Error: XAMP', 'TDRP', 'ADRT', 'XONE', 'LUNC', 'ICRC'.

Log entry: Device hangs while reading block X – Y.


Performing a power cycle

If 20 seconds after the second **Reset** command the drive still tries to read the bad block, TaskForce performs the **Power cycle** command by forcibly cutting power to the drive for 5 seconds.

At this point, TaskForce adds two entries to the imaging **Log**:

Performing power cycle... (when the power is cut off) and

Waiting for the device to become ready... (when the power is switched back on).



Log entry: Performing power cycle.

After a successful power cycle

In case the first **Power cycle** command is successful, and the drive become ready to accept the next command, there will be a final log entry for this problematic block of sectors saying:

Cannot read block of data at X - Y (Timeout).

And then TaskForce continues imaging from the next planned block.

After an unsuccessful power cycle

If the first **Power cycle** command is ineffective, and the drive is still in **Busy** state and can't run the next command, TaskForce makes the second **Power cycle**.

If the second **Power cycle** does not help either, imaging is terminated. It can be resumed afterward, and TaskForce will continue to image all remaining sectors.


Freeze detection during diagnostics

TaskForce diagnostics module automatically checks the status of each head (for HDDs) and the condition of the media surface by reading several hundreds of thousands of sectors from the starting/outer, middle, and ending/inner part of the head or drive.

If a drive freezes, unable to read certain sectors, it may or may not be detected during diagnostics because not all sectors of the drive are being read at this stage.

If an HDD has a damaged head causing it to enter a busy state, a diagnostics report notifies you about that with the line

Device freezes when checking head N. Performing power cycle...



Notification “Device freezes when checking head N” in the diagnostics report.

Also, the diagnostics report shows a table with a freeze count and the number of read errors for each HDD head:

Head #	Speed	Sectors processed	Freeze count	Read errors	Head status
0	68 MB/s	390,058	1	1	OK
1	< 1 KB/s	14,336	1	13	Damaged
2	90 MB/s	565,788	0	0	OK
3	79 MB/s	542,205	0	0	OK

According to the media access timing verification test, the head stack is malfunctioning.

Recommendation: disable damaged/degraded heads in imaging settings to image good heads first.

A table with a freeze count for a drive with a damaged head.

Head #	Speed	Sectors processed	Freeze count	Read errors	Head status
0	70 MB/s	445,065	0	0	OK
1	58 MB/s	466,284	0	0	OK
2	72 MB/s	354,934	2	2	OK
3	83 MB/s	500,643	1	0	OK
4	85 MB/s	396,564	0	0	OK
5	83 MB/s	439,825	0	0	OK

According to this test, the device sometimes falls into complete freeze, which will require power cycling.

A table with a freeze count for a drive with good heads.

If the heads status is OK, but some freezes were detected during diagnostics, the result of the diagnostics show the following line:

According to this test, the device sometimes falls into complete freeze, which will require power cycling.

Diagnostics completed. Device has issues

Results

No major hardware or firmware issues found

According to this test, the device sometimes falls into complete freeze, which will require power cycling.

SMART reports that there are defects in the media.

The Diagnostics results notifying that the device sometimes falls into complete freeze.

Imaging a shorted hard drive

Atola TaskForce has built-in short circuit protection and can detect shorted hard drives. In most cases, a drive has become shorted after experiencing overvoltage, either due to a power supply failure or as a result of a user error. Here is what happens to a drive in these scenarios and how to fix this.


How drives become shorted

Most drives are equipped with two TVS diodes to protect the circuit from overvoltage. One diode is located on the 5V rail and another on the 12V rail.

If the drive experiences an overvoltage, the diodes convert the excess electrical power into heat energy and warm up, thus protecting the drive's circuit. Similarly, in the case of reverse polarity, the diode warms up as it conducts the current in the opposite direction.

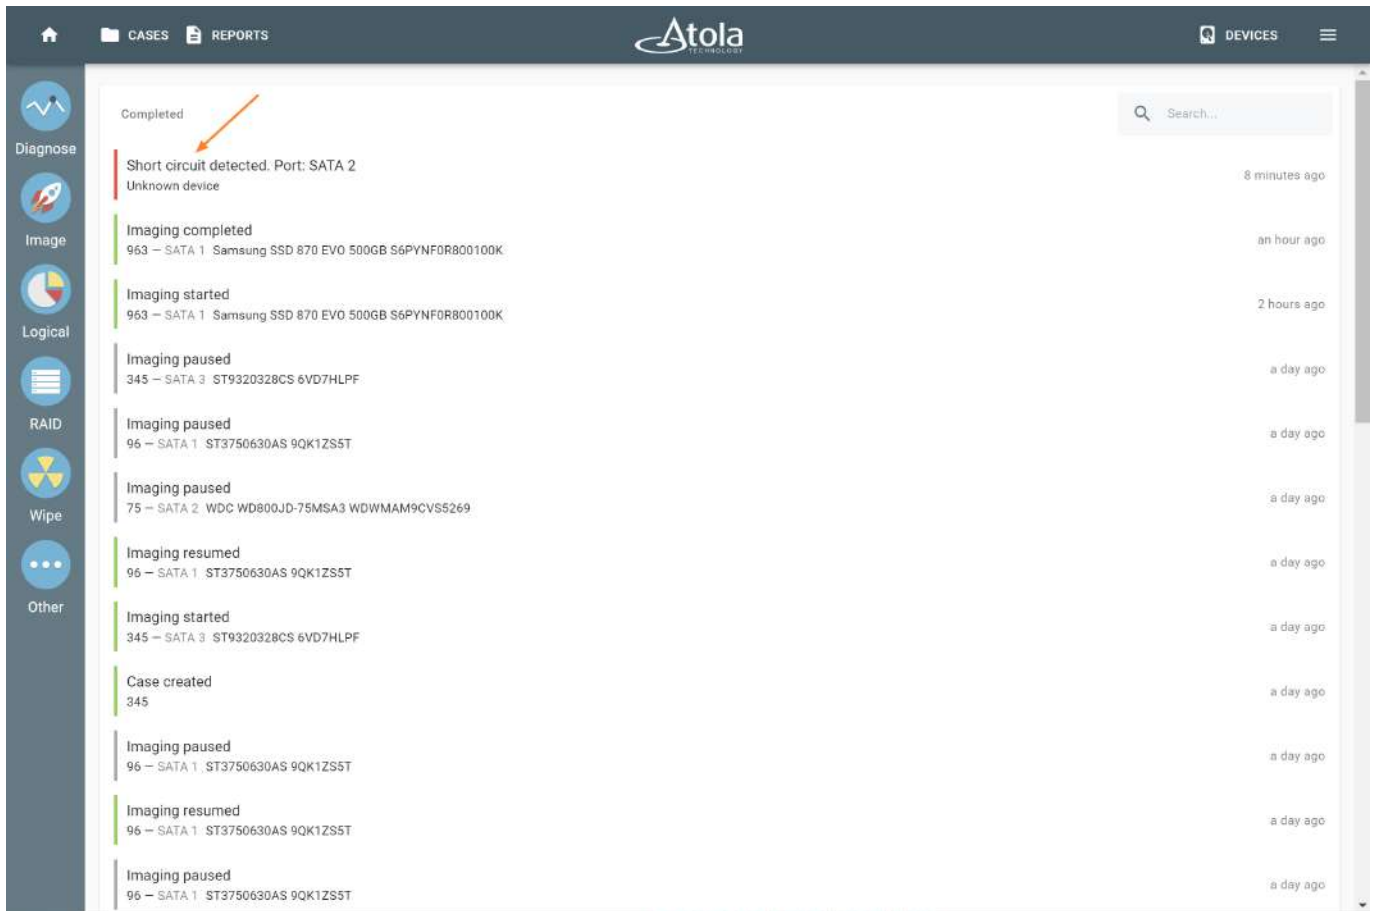
If the overvoltage or reverse polarity event is short and the dissipated energy is not too high, the diodes can recover and continue working. However, if the dissipated energy is too high, the diodes will "sacrifice" themselves and get shorted.

When the drive is subsequently powered, the shorted diodes create a low resistance connection between two nodes, known as a short circuit. This is exactly what happens to a drive when its TVS diodes are shorted.




Detect a shorted drive

When you connect a shorted drive to Atola TaskForce, the **Home** screen shows a short circuit alert to notify the operator about the detected issue.



The Home screen with a short circuit alert.

On the **Select device** panel, TaskForce marks the port to which a shorted drive is connected with the **Short circuit** tag and icon.




The Short circuit tag and icon on the Select device panel.

A drive with a shorted TVS diode cannot be identified, diagnosed, or imaged with TaskForce, until you replace or remove the diode.

Image a shorted drive

If you need to image a shorted drive but do not have new TVS diodes on hand to replace the shorted ones, you can image the drive using Atola TaskForce after removing the diodes. This process is safe because Atola TaskForce has short circuit and overvoltage protection, which guards both the imager and the drives connected to it against circuit failures.

To remove the diodes, heat the area of the drive where they are located with a hot fan (such as in a hot air soldering station) and then gently remove them with tweezers.



Once the diodes have been detached, you can plug the drive into Atola TaskForce and proceed with imaging data from it.

Express mode: self-launching imaging of 17 drive ports

Self-launching imaging mode is a perfect solution when it comes to processing large amounts of data under time pressure, while still allowing gentle treatment of damaged media.

Express mode enables the automatic launch of multiple imaging sessions with predefined settings on all drive ports set to source, except the Extension port (17 in total). Once an examiner plugs drives into TaskForce, imaging sessions start automatically. Drive images are saved as a RAW, E01, or AFF4 file in a specified folder on the local server. Two 10Gb Ethernet ports of TaskForce enable high data throughput.

Additionally, there is the option to group ports in Express mode and assign different imaging settings for each group. TaskForce can be set to automatically diagnose all source drives before imaging and proceed only if diagnostics show no issues.

In this article, you will learn how to:

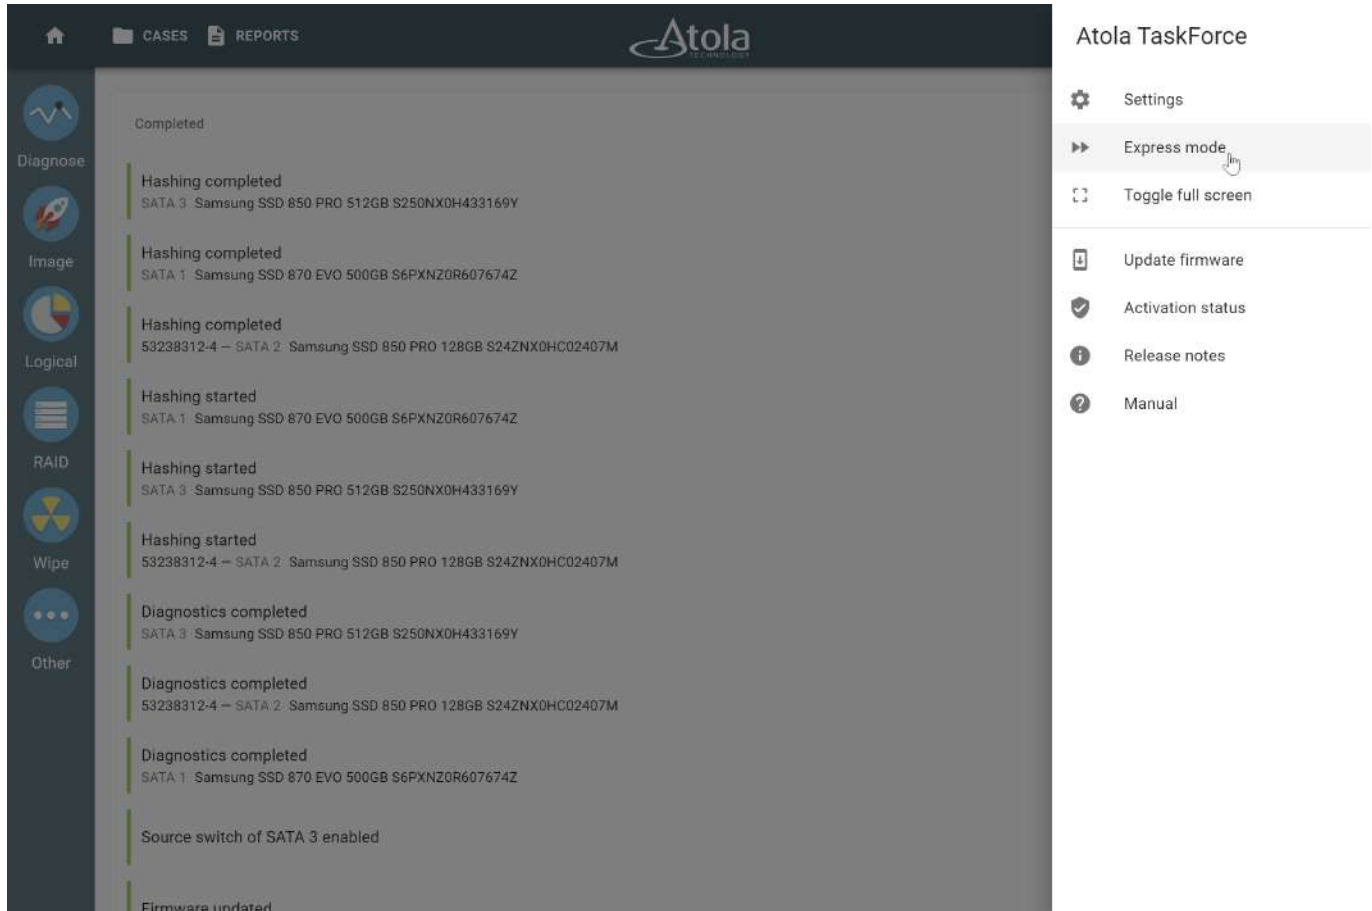
- [Configure Express mode](#)
- [Activate Express mode](#)
- [Image drives in Express mode](#)

- [Deactivate Express mode](#)

Configure Express mode

Before activating Express mode, you need to configure it first: specify the target folder, file image format, and other imaging details.

To open the Express mode settings page, click the **Menu** in the top right and select **Express mode**.




Choosing Express mode from the Menu.

TaskForce can process 17 self-launching imaging sessions in Express mode on almost all of its ports, except for the Extension slot. The ports that can be used for imaging in Express mode are:

- 6 SATA
- 6 SATA/SAS
- 4 USB
- 1 IDE

When activated, Express mode controls all source ports, leaving target ones available for other tasks. If a port is switched from target to source, it also becomes available for imaging in Express mode.

By default, Express mode settings are applied to all drive ports set to source. However, you can combine any number of TaskForce ports into a custom group with specific imaging parameters for Express mode. See [Port groups](#) for details.



Express mode settings.

Port groups

In cases where different types of drives require different approaches to image acquisition, you can assign specific imaging settings for certain groups of drive ports in Express mode.

For example, to balance image acquisition time, you can image newer drives to your faster server and older/smaller SATA drives to slower storage. Or image healthy drives to the "Good" network folder and faulty ones to the "Bad" network folder to treat them differently later.

In Express mode, for each custom group of drive ports, you can specify:


- Network folder where to save an image file
- Which image file format to use
- Which type of hash to calculate
- Whether to diagnose the source device before imaging
- How to fill in case report fields: Case ID, Organization, Investigator

Add a new group of ports

By default, Express mode settings are applied to all drive ports set to source. Thus, the group tab on the Express mode settings page is named **All ports**. All the ports included in the group are listed in the **Ports** section.


To add a new group of ports:

1. Next to the **All ports** group tab, click the **Plus** icon.



The All ports group tab and the Plus icon.


2. In the dialog window, select ports for the new group.



Selecting ports for the new group.

3. Click **Select**.

To avoid confusion and accidental mistakes, TaskForce excludes selected ports from any existing groups and adds them to the new one.



Several port groups in the Express mode settings.

4. The group name changes to clearly show port types and numbers included in it, for example, "SATA 1-8" or "SAS 1-8". Switch between port groups by clicking the tab with the group name.

Now you can proceed with assigning different imaging settings in Express mode to different port groups.

Change a port group

To change the drive ports included in an existing port group:


1. On the Express mode settings page, switch to the port group you want to change by clicking on the tab with its name.
2. On the right, click **Change**.
3. In the dialog window, select the drive ports you want to include in the group or deselect the drive ports you want to exclude.
4. Click **Select**.
5. The group name changes to clearly show port types and numbers included in it.

Target folder

In Express mode, TaskForce stores drive images on a local server.

To activate Express mode, you need to specify a target folder on a local network to store drive images:

1. On the Express mode settings page, find the **Target folder** section and click **Select folder**.
2. In the **Select target folder** window, go to the existing folder and click the **Select folder** button. To create a new folder, click the icon with three dots and select **Create folder**.



The 'Select target folder' window.

What image files to create: type, compression, hash, templates and more

To set up the parameters of a target image file, such as file type, compression method, segment size, and hash type, on the Express mode settings page, find the **What image file to create** section and click **Change**.

Configuration dialog appears.

On the **File** tab of this dialog, configure the target image file.


On the **Template** tab, you can also create the [templates for image file names](#).

Selected parameters apply to an active port group only. Here are the details about each parameter.

Image file type

In Express mode, drive images can be saved as RAW, E01, or AFF4 files.

To change the target image file type for a specific group of ports, select it from the **Type** list on the **File** tab.




Selecting target image file type.

Compression

TaskForce supports the default compression method for the E01 file format and LZ4 or Snappy compression methods for the AFF4 file format.

Choose the compression method from the respective list on the **File** tab. The **Not compressed** option is also available for both E01 and AFF4.

The RAW file format does not support compression.




Selecting the compression method for a target image file.

Segment size

For the E01 file format, TaskForce can split the target image into segments of predefined size. Each segment is stored in a specified folder as a separate file.

The **Segment size** list becomes available on the **File** tab when the E01 file type is selected. You can choose the segment size from the following options:

- Single file
- 4.7 GB (DVD)
- 8.5 GB (DVD DL)
- 25 GB (Blu-ray)
- 50 GB (Ultra HD Blu-ray)
- 100 GB (Blu-ray XL 3L)
- Custom (enter the segment size you want)




Selecting the segment size for a target image file.

Hash type

For AFF4 file format, TaskForce can calculate built-in block hashes during the imaging process in the Express mode.

To select the hash type you want to calculate, use the **Block hash** list on the **File** tab. The available hashing options for the AFF4 file format are:


- MD5
- SHA1
- MD5+SHA1
- SHA256
- SHA512



Selecting the hash type for a target image file.

Create template for image file names

To create a template for naming target subfolders and files according to your organization rules, find the **What image file to create** section, click **Change**, and switch to the **Template** tab. For details, see [Target file templates](#).



Creating a template for a target image files names in Express mode.

Diagnose before imaging and Start imaging only if diagnostics has no issues

As essential as imaging speed is, the proper treatment of evidence drives remains a priority. That's why we recommend that [diagnostics](#) is always performed as soon as a drive is connected to TaskForce for the first time.


To enable the automatic launch of imaging of the healthy devices and avoid potential deterioration of drives in a shaky condition, the Express mode settings have two handy options:

- select **Diagnose source drive before imaging** so that diagnostics start automatically,
- then select **Start imaging only if diagnostics has no issues**.

When enabled in Express mode, Atola's signature automated diagnostics module checks all drive systems:

- hard drive's motor and electronics (PCB),
- head stack,
- media surface,
- all firmware/system areas,
- partitions and file systems.

For more information about TaskForce's diagnostics module, see [Diagnosing a drive with Atola TaskForce](#).



Diagnostics options for Express mode.


Imaging presets

To ensure that all imaging sessions in Express mode fit your organization's demands and procedures, in Express mode you can use one of the preconfigured [imaging presets](#).

TaskForce has two built-in imaging presets, *Default* and *Damaged*, for good and bad drives respectively, but you can also create and configure your own preset. For guidance, see [Imaging presets](#).

To select an imaging preset for Express mode, use the **Type** list in the **Imaging settings preset** section.

Changing the imaging preset is available only when Express mode is deactivated.



The 'Type' list in the 'Imaging settings preset' section.


Set case details

To streamline your acquisition workflow in the Express mode, you can pre-fill case details for each evidence drive using predefined information:

- Case ID
- Investigator's name
- Organization
- Location
- Description


To enter predefined case details for Express mode and allow filling them up:

1. On the Express mode settings page, go to the **Set case details** section and toggle this option on.



2. Click **Change**.

3. In the dialog window, enter case details and click **Save**.




Activate Express mode

Once Express mode settings are specified, click the **Activate** button and connect your evidence drives for an immediate start of imaging.

If any drives are already connected to the source ports, TaskForce starts to image those devices immediately using Express mode settings once you click **Activate**.

To indicate that TaskForce now operates in Express mode:

1. In the user interface, the color of the top panel and the Image icon on the left changes.



2. On the TaskForce's IP screen, the message "Express mode" appears.



Image drives in Express mode

Plug drives

Once all settings are configured and Express mode is activated, simply plug in the drives one by one and watch the imaging sessions start automatically.

Control imaging status in Express mode

- **Home screen:** All diagnostics and imaging processes launched in Express mode appear on the Home screen. Click any process to see its details.
- **IP screen:** TaskForce keeps the user updated by displaying the number of user actions required on the IP screen on the front panel of the unit.

User action required

When TaskForce detects an issue with an imaging session, the **User action required** notification appears on the Home screen for the user to make a decision.

The number of user actions required is also displayed on the IP screen on the front panel of the hardware unit.

TaskForce prompts for user to take action if:


- Diagnostics detect any issues with a device.
- A user pauses the imaging process.
- Imaging finishes.
- A source device has already been imaged, fully or partially.
- A source device has been formatted as storage.

Pause and resume imaging session

As in standard imaging mode, in Express mode you can pause any imaging session and then resume it later.

To **pause** an imaging session in Express mode:


1. On the Home screen, click on the running imaging process you want to pause.
2. On the imaging details page, click the **Pause** button at the bottom.
3. Confirm pausing the imaging process.
4. TaskForce shows the notification that the port has been powered off.
5. **Optional:** Unplug your evidence drive from the port.
6. **Optional:** Plug the next drive or do not plug any.
7. In the dialog, click **Continue**.



Once the imaging process is paused, TaskForce creates an **Imaging paused report** available either from the **Home** screen or from the **Reports** page.

To **resume** the paused imaging session in Express mode:

1. On the Home screen, click the paused session.
2. In the dialog, click **Resume**.



Check imaging reports

Once the imaging process is started, paused, resumed, or completed in Express mode, TaskForce creates the respective report.


Reports contain the various details of the session in its current stage including source and target drive details, imaging settings, hash values, the time when the imaging session started, when it was completed and more.

Reports are available either from the **Home screen** or from the **Reports** page.

Change drives

Once an imaging session in Express mode is completed, TaskForce powers the respective port off and prompts a user to plug in the next evidence device.

To start a new imaging session, plug in the next source device and click **Continue** in the dialog.



Deactivate Express mode

To deactivate Express mode:

1. Click the **Menu** in the top right and select **Express mode**.
2. On the bottom of the Express mode settings page, click **Deactivate**.

To show that Express mode is deactivated, TaskForce changes the color of the top panel and the Image icon on the left back to normal. The message “Express mode” on the IP screen also disappears.

Logical imaging

The **Logical imaging** module saves you time and storage space by copying only selected files and folders from an individual physical drive, image file, or RAID array in a forensically sound manner, without altering the original content of a source device and its metadata.

As a result of logical imaging, you get an image file in L01 or ZIP format, which contains only the data you selected manually or using include/exclude filters. L01 image will also include MD5/SHA1 hashes of the files.

After the logical imaging process is completed, TaskForce generates a comprehensive report with the details about the source, target, number of imaged files and bytes and more.

Key features of logical imaging in TaskForce

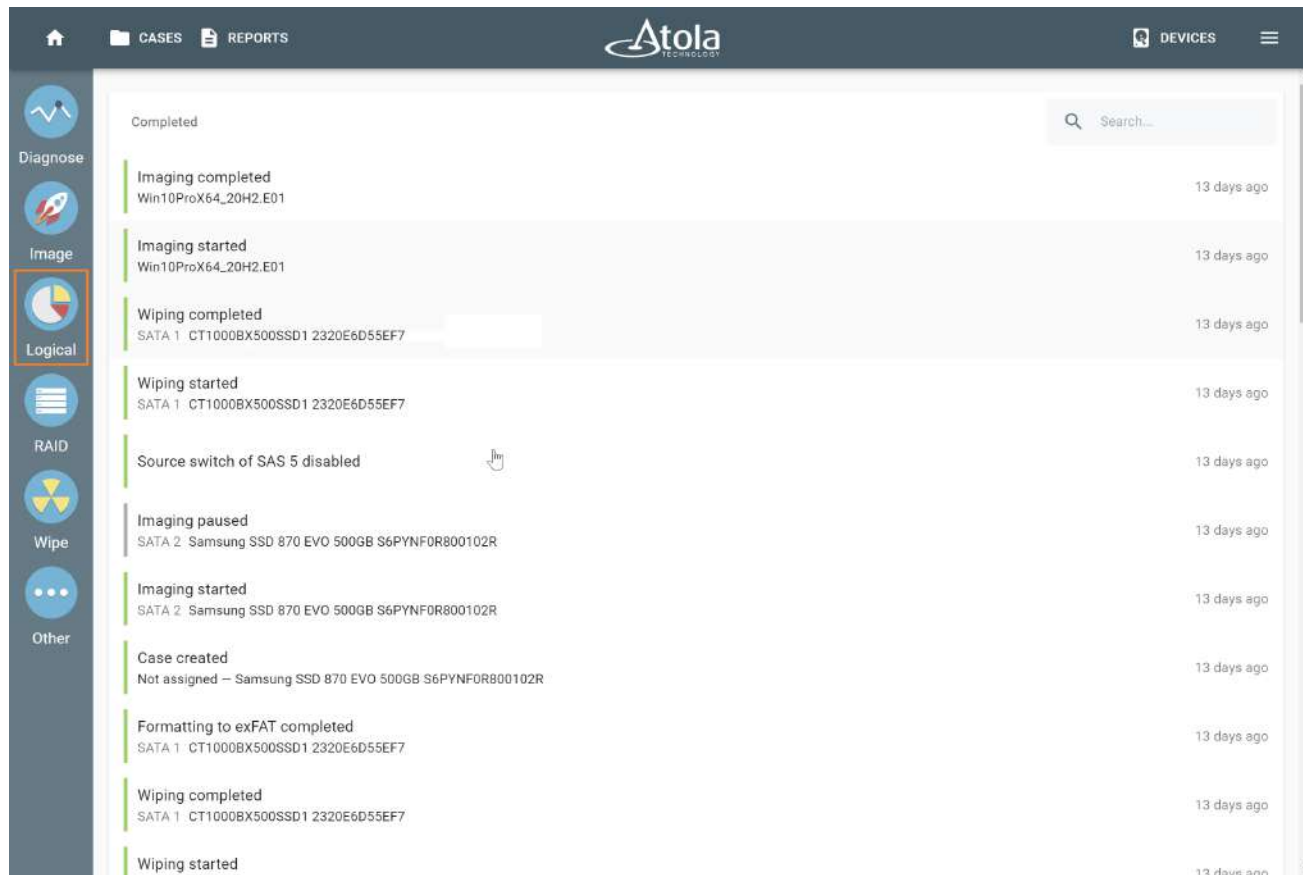
- 25+ parallel logical imaging sessions.
- Pause/Resume for each session.
- L01 or ZIP target images created in a network folder or on a storage drive.

- Compression of a target file is available.
- Supported file systems: NTFS, ext4/3/2, exFAT, XFS, ZFS, Btrfs, APFS (with encrypted volumes), HFS/HFS+, FAT32/16.
- Smart include/exclude filters.

Start logical imaging


There are two ways to start the Logical imaging module in TaskForce.

- **From the Homescreen:** On the Taskbar, click the Logical icon.



The Logical icon in the Taskbar.

- **From the RAID module:** On the RAID Configuration page, click the Go to logical button. For details, see [Logical imaging of a RAID array](#).



The Go to Logical button on the RAID configuration page.

Logical imaging workflow in TaskForce

Once you started the Logical imaging module, follow these steps to create a logical image of your source device:

1. [Select your source device.](#)
2. [Select partitions, folders, or files to image.](#)
3. [Apply smart filters, if needed.](#)
4. [Select target file.](#)
5. Optional: [Pause and resume the imaging process.](#)
6. [View the logical imaging report.](#)

Let's see in detail how to perform each step.

Select source device


After clicking the Logical icon on the Taskbar, select your source device.

TaskForce supports the following source types:

1. Individual physical drive:
 - o SATA,
 - o SAS,
 - o NVMe,
 - o USB,
 - o IDE,
 - o drives with M.2 interface within Apple laptops using the [Apple PCIe SSD extension module](#),
 - o Apple MacBooks with FireWire, Thunderbolt 2, and Thunderbolt 3 interfaces using the [Thunderbolt extension module](#).

2. Image files:
 - IMG, RAW, DD
 - E01
 - AFF4
3. Reassembled RAID arrays.

File systems, supported by the Logical imaging module, are NTFS, ext4/3/2, exFAT, XFS, ZFS, Btrfs, APFS (with encrypted volumes), HFS/HFS+, FAT32/16.




Selecting a source device for logical imaging.

Select partitions, folders, or files to image

After you select your source device for logical imaging, TaskForce takes you to the Logical imaging page. Here you can preview the content of available partitions and folders, apply smart filters or use presets with predefined filtering settings.

Your source device and its case ID are shown at the top of the page.



Selecting partitions, folders, and files for logical imaging.

Preview, select, or deselect partition

The **Partition** section on the left lists all the partitions available on your source device.

- To **preview** the partition contents, click on its name.
- By default, all partitions are selected for imaging.
- To **exclude** a partition, clear the checkbox next to it.
- To **include** a partition again, select the checkbox.

Preview folder, manually include or exclude folders and files

Folders and files of the selected partition are shown in the middle section. By default, TaskForce images all files from all partitions.

- To **preview** a folder's contents, click on the folder name.
- To manually **include** a folder or file for imaging, click on its icon once. The icon turns green. The folder or file appears in the Filters section on the right as manually included.
- To manually **exclude** a folder or file, click on its icon twice. The icon turns orange. The folder or file appears in the Filters section on the right as manually excluded.

Use smart filters and presets for selecting folders or files

Automatically include or exclude folders and files for imaging by using the built-in smart filters or custom filtering presets.


The **Filters** section on the right provides flexible options for fine-tuning your selection of folders and files for imaging.

By default, TaskForce images all files from all partitions.

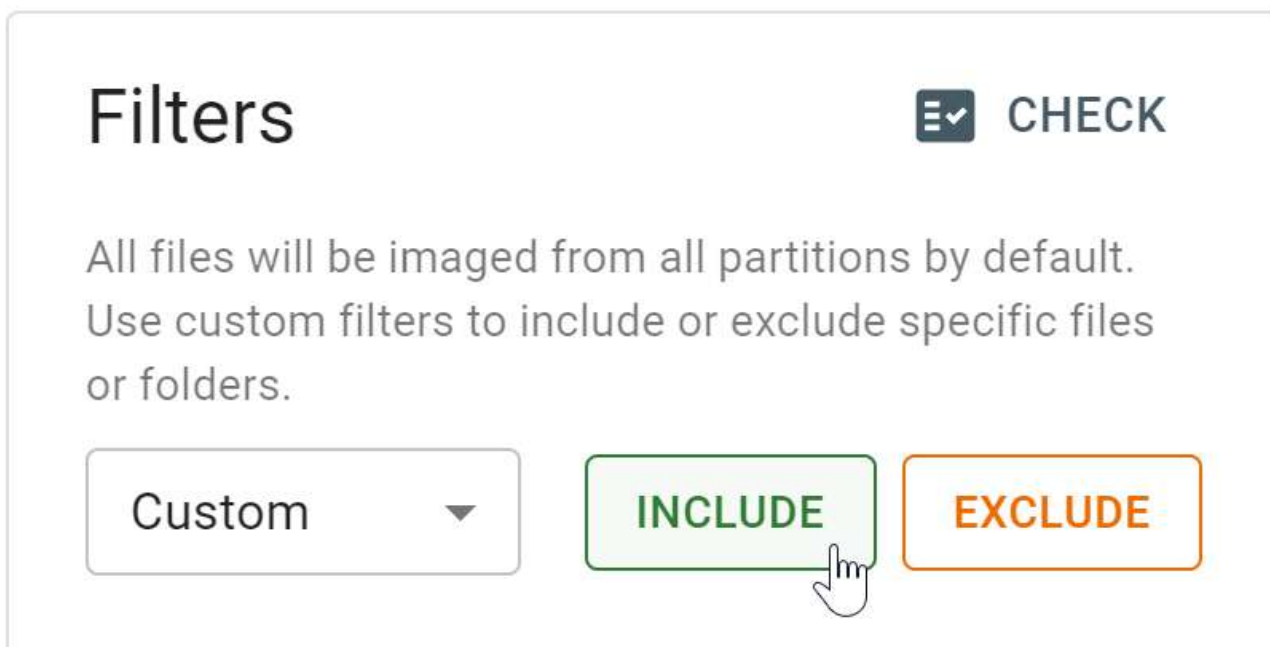
Apply filters

Use smart filters to automatically include or exclude specific files or folders:

1. In the **Filters** section, select the **Custom** option from the list.




2. Click the **Include** or **Exclude** button. The respective dialog appears.



3. Automatically include or exclude folders:

- In the respective dialog, select the **Folders** option from the list at the top.
- Click in the **Folders** field.
- Select one of the predefined options: **User folders** or **OS folders**.
- Or enter the folder paths, separated by commas.
- Click **Apply**.



The screenshot shows a dialog box titled 'Include' with a green header. In the top right corner, there is a dropdown menu labeled 'Folders' with a downward arrow. Below the header, there is a text input field with a placeholder text 'Folders' and a line of examples: 'Examples: /Users/James, /home/zelda, etc'. Below the input field, there are two options: 'User folders' and 'OS folders'. The 'User folders' option is highlighted with a light gray background and includes the text '/Documents And Settings, /home and 8 more' and a hand cursor icon. The 'OS folders' option is below it and includes the text '/Windows, /Program Files, /dev and 13 more'.

4. Automatically include or exclude files:

- In the respective dialog, select the **Files** option from the list at the top.
- Specify **File types** by selecting predefined options from the list: archives, audio, databases, documents, emails, financial, pictures, security keys, video, and virtual machines. You can enter file extensions yourself as well.
- Enter file size range: from 1 byte to infinity.
- Narrow down time spans: when files were modified, accessed, or created.
- Click **Apply**.

Exclude

Files

File types

Audio ×

Pictures ×

Videos ×

Examples: txt...

From

1,000

Bytes

To

Infinity

Modified

Accessed

Created ●

From

01/01/2015



To

12/31/2017



CANCEL

APPLY



All applied Include or Exclude filters appear in the Filters section.

Filters

 CHECK

All files will be imaged from all partitions by default.
Use custom filters to include or exclude specific files or folders.

Custom ▼

INCLUDE

EXCLUDE

Include:

Folders:

User folders



Exclude:

Databases, Documents

Greater or equal to 1,000 bytes

Created between 01/01/2015 and 12/31/2017



Check filters before imaging

To test your filtering settings before imaging and get a list of files that match the selected filters, click **Check** at the top of the Filters section:

Filters



CHECK

All files will be imaged from all partitions by default.
Use custom filters to include or exclude specific files
or folders.

Custom



INCLUDE

EXCLUDE

Include:

Folders:

User folders



Exclude:


Databases, Documents

Greater or equal to 1,000 bytes

Created between 01/01/2015 and 12/31/2017



TaskForce scans the selected partition(s) and lists all files matching the filtering criteria. You can download this file list in CSV format by clicking the **Download** button at the bottom.




Downloading the list of files matching the filtering criteria.

Save filtering settings as a preset

To reuse your filtering settings later or share them with your colleague, save them as a logical imaging preset:


1. In the bottom right corner of the Logical imaging page, click the **More** icon with three dots.
2. Select **Save to**.
3. Enter the name for your preset and click **Save**.



Saving a logical imaging preset.

From the same three-dots menu, you can also:

- Export a preset to use it on another computer.
- Import a preset created on another computer.
- Delete a preset.



To switch between presets, use the **Current settings** list on the bottom right of the logical imaging page.

Atola TECHNOLOGY

Home CASES REPORTS DEVICES

Diagnose Image Logical RAID Wipe Other

CT1000BX500SSD1 2320E6D55EF7
1 TB

Healthy
SATA 1

Case ID:
Logical-imaging-101

Partition	Name	Type	Size
☒ 0	System Reserved	NTFS - 524 MB	
☒ 1	No label	NTFS - 46 GB	
☐ 2	No label	Unknown - 553 MB	

Name	Type	Size
\$Recycle.Bin	Folder	
\$WinREAgent	Folder	
Documents and Settings	Folder	
Intel	Folder	
PerfLogs	Folder	
Program Files	Folder	
Program Files (x86)	Folder	
ProgramData	Folder	
Recovery	Folder	
System Volume Information	Folder	
Users	Folder	
Windows	Folder	
Windows.old	Folder	
DumpStack.log.tmp	tmp	8 KB
\$WINRE_BACKUP_PARTITION.MA...	MARKER	0 B
hiberfil.sys	sys	5 GB

Filters

Custom INCLUDE EXCLUDE

Include:

- Audio, Videos
- Manually selected: Partition 1/Windows, Partition 1/Users, Partition 1/hiberfil.sys

Exclude:

- Manually selected: Partition 1/Program Files, Partition 1/Program Files (x86)

CONTINUE

Current settings: Audio-Video

Selecting a preset from the Current settings list.


Select target file

Once you have adjusted the logical imaging parameters, click **Continue**.

TaskForce takes you to the **Select target folder** window.


The target file in L01 or ZIP format (with or without compression) can be created either on a network drive or on a storage device connected directly to the TaskForce hardware unit:

1. Select the folder for your target file and click **Create file**.



The Select target folder window.

2. Adjust the settings of the target file:
 - a. Name.
 - b. File type: L01 or ZIP.
 - c. Hashing method (Digest type): MD5, SHA1, or MD5+SHA1.
 - d. Compression.
 - e. Case ID and details.
3. Click **Create** to start the imaging process.



Creating logical imaging file.

How to set a drive to Storage mode

To create a target file on a drive connected directly to the TaskForce, you must first add this drive as a storage device:

1. In the **Select image file** window, click **Add storage**.
2. On the **Select device** panel, choose the drive you want to use in Storage mode. TaskForce uses a lighter shade of blue to indicate that a storage drive is being configured.
3. If TaskForce cannot find the appropriate exFAT partition on the selected drive, it offers you to format the device accordingly. In that case, select the **Format device to exFAT** option and click **Next**.
4. To launch target device formatting to exFat with a large cluster size (32 MB), click **Format** and enter **YES** for confirmation. This cluster size enables faster imaging to this drive.
5. Once the target device is formatted, TaskForce perceives it as a Storage target.

Logical imaging process. Pause and resume

The logical imaging process starts immediately when you click Create in the Create logical image file dialog.

Once imaging is launched, TaskForce creates a **Logical imaging started** report. It contains all details about the imaging source and target(s), as well as information about the include or exclude filters. You can find this report on the Home screen, on the Case page, or the Reports page.

Home

CASES

REPORTS

Atola
TECHNOLOGIST

DEVICES

Diagnose

Image

Logical

RAID

Wipe

Other

CT1000BX500SSD1 2320E6D55EF7

1 TB

SATA 1

Case ID:
Logical-imaging-101

Creation date: 05/14/2024 1:50 PM

TaskForce serial: 33688108

Location: SATA 1

Device model: CT1000BX500SSD1

Size: 1 TB (1,000,204,886,016 bytes)

Case investigator: G. I. Jones

Software version: 2024.2

TaskForce IP: 10.0.0.102

Write protection: On

Device serial: 2320E6D55EF7

Device firmware: M6CR056

Case description: Logical imaging 101

Logical imaging started

Selected partitions

All partitions

Include filters

Manually selected folders and files:

Partition 3/\$RECYCLE.BIN

Partition 3/Encase files

Archives, Documents, Pictures, Videos

7z, 7z001, 7zip, apk, ar, arj, bz2, bzip, bzip2, cab, cpio, deb, dmg, gz, gzip, jar, lha, lz, lzh, rar, rpm, tar, tbz2, tgz, z00, z01, zip, chm, csv, doc, docx, gdoc, gdraw, htm, html, kwd, numbers, odoc, odp, ods, odt, osheet, ott, pages, pdf, ppt, pptx, rtf, tex, txt, xhtml, xls, xlsx, xml, xps, 3g2, 3gp, aaf, asf, avi, avif, bik, braw, fcp, flv, h264, hevc, imovieproj, m1v, m2v, m4v, mkv, moov, mov, movie, mp4, mpe, mpeg, mpg, mpv, mswmm, ppj, prproj, rm, swf, veg, veg-bak, vob, vproj, webm, wfp, wimp, wmf, wmv, wve, ai, bmp, cdr, clip, cmx, cpt, cr2, dpx, eps, fpx, gbr, gif, heic, heif, icns, ico, indd, jfif, jng, jp2, jpe, jpeg, jpg, max, nef, odg, ogex, pdd, pdn, pef, png, psd, rw2, svg, tec, tif, tiff, wrl, webp, xcf

Exclude filters

OS folders

/Windows, /WinNT, /Program Files, /Program Files (x86), /ProgramData, /Applications, /System, /bin, /boot, /dev, /sbin, /lib, /usr, /proc, /sys, /unix

Targets

CT1000BX500SSD1 2320E6D55EF7.L01 0 B S:/Images/L01

PRINT

GO TO CASE

GO TO LOGICAL

<

<

>

>

A Logical imaging started report.

During the imaging process, you can track the progress of individual files.

To check the filtering parameters of this running session, click **View filters** at the bottom of the page.

Atola TECHNOLOGY

CT1000BX500SSD1 2320E6D55EF7
1 TB
SATA 1
Healthy

Case ID: Logical-Imaging-101
1 target

Logical imaging

6%

gnsdk_fp.dll 85%

Files scanned: 48,997
Files imaged: 48,967
Speed: 270 MB/s
Bytes imaged: 3.03 GB

Date ↑	Message
2024/05/14 14:56:11	Logical imaging started
2024/05/14 14:56:11	Analyzing: Partition 1 at sectors 1,026,048 - 90,038,951 (type: NTFS)

PAUSE VIEW FILTERS


Link: PHY Status: ERR INDEX CORR DRSC FAULT DRVY BUSY Error: AMNF TONF ABRT - IDNF - UNDF JCRC ?

The logical imaging process.

Pause and Resume

TaskForce lets you pause any logical imaging session and resume it later:

1. While the logical imaging process is running, click the **Pause** button at the bottom of the page.
2. To resume a session, click **Logical** on the Taskbar and then select your source device.
3. On the **Recent logical imaging sessions** page, click **Resume** next to your session.



Resuming the logical imaging session.

Use cases of Pause/Resume

When you have to temporarily power off your TaskForce:

1. Pause logical data acquisition.
2. Turn off the TaskForce.
3. Power it on later.
4. Resume the logical imaging session exactly from the very "pause" moment.

When you want your colleague to continue imaging on another TaskForce:

1. Pause logical data acquisition.
2. Export the case.
3. Transfer the source drive and the exported case to a different location that has another TaskForce.
4. Import the case.
5. Resume the logical imaging session.

Logical imaging completed report

After the logical imaging process is completed, TaskForce generates a comprehensive report with all the details of the imaged data:

- The number of scanned and imaged files
- The volume of imaged data
- Time stamps
- and more

You can find this report on the **Home** screen, on the **Case** page, or on the **Reports** page.

The screenshot displays the Atola software interface. On the left is a vertical sidebar with icons for Diagnose, Image, Logical, RAID, Wipe, and Other. The top navigation bar includes 'CASES', 'REPORTS', and 'DEVICES' tabs. The main content area shows a report for a device with ID CT1000BX500SSD1 2320E6D55EF7, 1 TB, SATA 1. The report title is 'Logical imaging completed'. It includes sections for 'Targets' (listing the device, evidence number 784708-0896, investigator G. I. Jones, and description 'Logical imaging 101'), 'Stats' (showing 353,908 files scanned, 353,896 files imaged, and 44 GB imaged), and a 'Log' section with three entries: 'Logical imaging started', 'Analyzing: Partition 1 at sectors 1,026,048 - 90,038,951 (type: NTFS)', and 'Logical imaging completed'. At the bottom are buttons for 'PRINT', 'GO TO CASE', and 'GO TO LOGICAL', along with navigation arrows.

The 'Logical imaging completed' report.

Imaging only sectors with data

Capacity of an average drive is constantly growing, and selective imaging becomes a way out for many investigators to keep their backlogs smaller.

We at Atola have developed selective imaging functionality to make it possible to image only sectors containing data.

The feature is supported in these file systems: NTFS, APFS, XFS, ext2/3/4, HFS, HFS+, ExFAT, FAT16, FAT32.

To image only sectors with data, do the following:

1. On the **Imaging settings** page, click **Change** to adjust imaging settings.
2. On the **Passes** tab, click on the value in the **What to image** column.

Atola
TECHNOLOGY

CASES

REPORTS

DEVICES

Diagnose

Image

Logical

RAID

Wipe

Other

WDC WD3200JD-22KLB0 WDCAMR1837670

320 GB

Has issues

SATA 4

Case ID: 62

Settings

BACK

PASSES

HASHES

MISCELLANEOUS

Pass	Timeout	Jump on error	Max read block size	What to image
1	1s	1 000 000	4 096	All sectors
2	5s	20 000	4 096	
3	5s	4 096	4 096	
4	5s	256	256	
5	60s	1	256	

+ ADD PASS

RESTORE TO DEFAULT

START

Current settings: DEFAULT DAMAGED DAMAGED...


Link: PHY

Status: ERR INDX CORE DREQ DISC FAULT DIRDY BUSY

Error: AMNR TONR ABRT IDNF UNC ICRC

The value in the What to image column.


3. In the Edit what to image on passes dialog, select Sectors with data.



Selecting the Sectors with data option.

4. To preview the partitions on the source drive, click Show.

Imaging individual partitions is possible: just unselect the partitions you do not want to image. The imaging report will document this setting and indicate, which of the partitions have been imaged and the sector ranges in which the partitions were located.



Partitions preview.

5. In the **Edit what to image on passes** dialog, click **Save**.
6. To proceed with imaging, click **Start**.

Atola
TECHNOLOGY

CASES

REPORTS

DEVICES

Diagnose

Image

Logical

RAID

Wipe

Other

WDC WD3200JD-22KLB0 WDWCAMR1837670

320 GB

Has issues

SATA 4

Case ID: 62

Settings

BACK

PASSES

HASHES

MISCELLANEOUS

Pass	Timeout	Jump on error	Max read block size	What to image
1	1s	1 000 000	4 096	Sectors with data
2	5s	20 000	4 096	
3	5s	4 096	4 096	
4	5s	256	256	
5	60s	1	256	

+ ADD PASS

RESTORE TO DEFAULT

START

Current settings: DEFAULT DAMAGED DAMAGED...


Link: PHY

Status: ERR INDX CORE DREQ DISC FAULT DIRDY BUSY

Error: AMNR TONR ABRT IDNF UNC ICRC

The Start button on the Imaging settings page.

You can see the partitions being imaged in the imaging log. In the imaging bar, the blue areas represent the sectors that are planned to be imaged. These are the sectors that belong to the drive's partitions and contain data.



Information about the partitions imaged in the log.

Imaging cheat sheet

When source drive is damaged

Use these imaging settings and follow the recommendations to cope with severely damaged drives.

Optimal target types for damaged source device

AFF4 image, RAW image file, or target drive plugged into the unit. Best to use segmented hashing with linear hashing disabled.

E01 is a linear format. It limits the use of TaskForce's advanced imaging features, e.g. reverse imaging or manual jumps.

Segmented hashes

Imaging with linear hash: one MD5/SHA1 hash.

Imaging with segmented hashes: many hashes of corresponding LBA ranges of the image.

The sum of these LBA ranges represents the entire image, though not necessarily in sequential order. You can still prove that the entire image has not been modified by verifying all hashes in a set.

Damaged head

If Diagnostics detects a damaged or degraded head, disable the head in the imaging settings for the initial imaging session.

Read more here: [Imaging a drive with a damaged head](#).

Damaged SATA drive on SAS port

TaskForce's SATA module and ports have more in-depth support of damaged SATA drives than the SAS subsystem. For instance, SAS doesn't allow sending so-called vendor ATA commands that TaskForce uses in SATA to detect heads.

To be most effective against damaged SATA drives, it is best to connect them to SATA ports.

USB drive read errors

Use high-quality short USB3 cables. Longer or lower-quality USB3 cables can produce read errors during imaging.

Reverse direction

Imaging pass setting.

When enabled, the imaging engine reads a drive backwards.

Pros:

- disables Read Look-Ahead effect
- reaches damaged areas from the opposite direction

Cons:

- speed decreases due to auto disabling of drive's cache

Manually skip bad sectors

If TaskForce can't read a block from a source device, it will automatically skip sectors and try to image them on later passes. When the imaging process seems to be stuck, you can move the imaging cursor manually by clicking another area on the top progress bar.

Manual sector jumps work only if linear hashing is disabled. You can use segmented hashing instead.

Disable read look-ahead

Imaging pass setting.

When enabled, a source device switches off its read-cache. Disabling read look-ahead decreases speed; but, it can be helpful against damaged drives.

Last imaging pass is too slow

Reduce a max time for a single read block attempt:

1. Pause imaging.

2. Add a new pass.
3. Change its Timeout setting to 500 ms (milliseconds).
4. Use the Go to pass button to switch to the newly added pass.

How imaging works on the last pass

The last pass has a unique feature which does not occur during previous passes: internal auto-reread procedure for error block sector-by-sector. It is defined by an unchangeable Jump size = 1 sector.

How imaging engine works on the last pass:

1. It reads block using Max Block Size pass setting (256 by default)
2. If reading is successful -> proceed to a next non-imaged block
3. If a read error occurs -> re-read the whole error block sector by sector.
4. If a read error occurs and ReadLong setting is enabled -> re-read using ReadLong command.

Faster imaging

If you want to speed up image acquisition, follow these hints.

Diagnose source before imaging

How it is useful:

1. Make sure the drive is in good condition or learn about the type of damage to make an informed decision about your following steps.
2. Prioritize the drive. Diagnostics report tells you if there is any data at all.
3. Use imaging time estimation.

Use faster targets

Good options:

- target SSD
- target NAS
- network server with RAID

When imaging a network, 10Gbit network connection is highly recommended.

When to use Logical imaging

- Time-constrained scenarios
- Large capacity drives (e.g., 10TB and above) or RAID arrays
- When specific files are of interest, not the entire drive

Image output format will be L01 or zip.

Change the timeout for a single read block attempt on-the-fly

You can read the data from your source device with a different timeout by adding another pass to your imaging session when it's paused:

1. Pause the currently running imaging, using the **Pause** button.
2. Click **Image** and select your Source device.
3. On the Recent imaging sessions page, find your paused session and click the **Edit imaging settings** icon next to it.
4. In the Imaging settings dialog, click **Add pass**.
5. Click on the new pass, edit the **Timeout** value and **Save** it.
6. Select the new pass, click **Go to Pass** and **Apply**.
7. Next to your paused session, click the **Resume** button.

Important: The resumed imaging session will complement the data imaged prior to the pause with only the sectors that were not yet copied.

Express mode

Express mode is ideal for processing large amounts of data quickly. It automatically starts an imaging session for a newly inserted drive:

1. In the TaskForce window, go to system **Menu > Express mode**.
2. Specify Express mode **Settings** and click **Activate**.
3. Plug devices into Source ports: SATA, SAS, USB, IDE.
4. Imaging will start automatically.
5. If imaging doesn't start and the corresponding front LED is off, check the Home screen.

Integrate TaskForce into your workflow via Web API

Web API helps external software to communicate with TaskForce. Another software may control TaskForce this way. To put it simply, there are simple operations in TaskForce Web API that any external software can run:

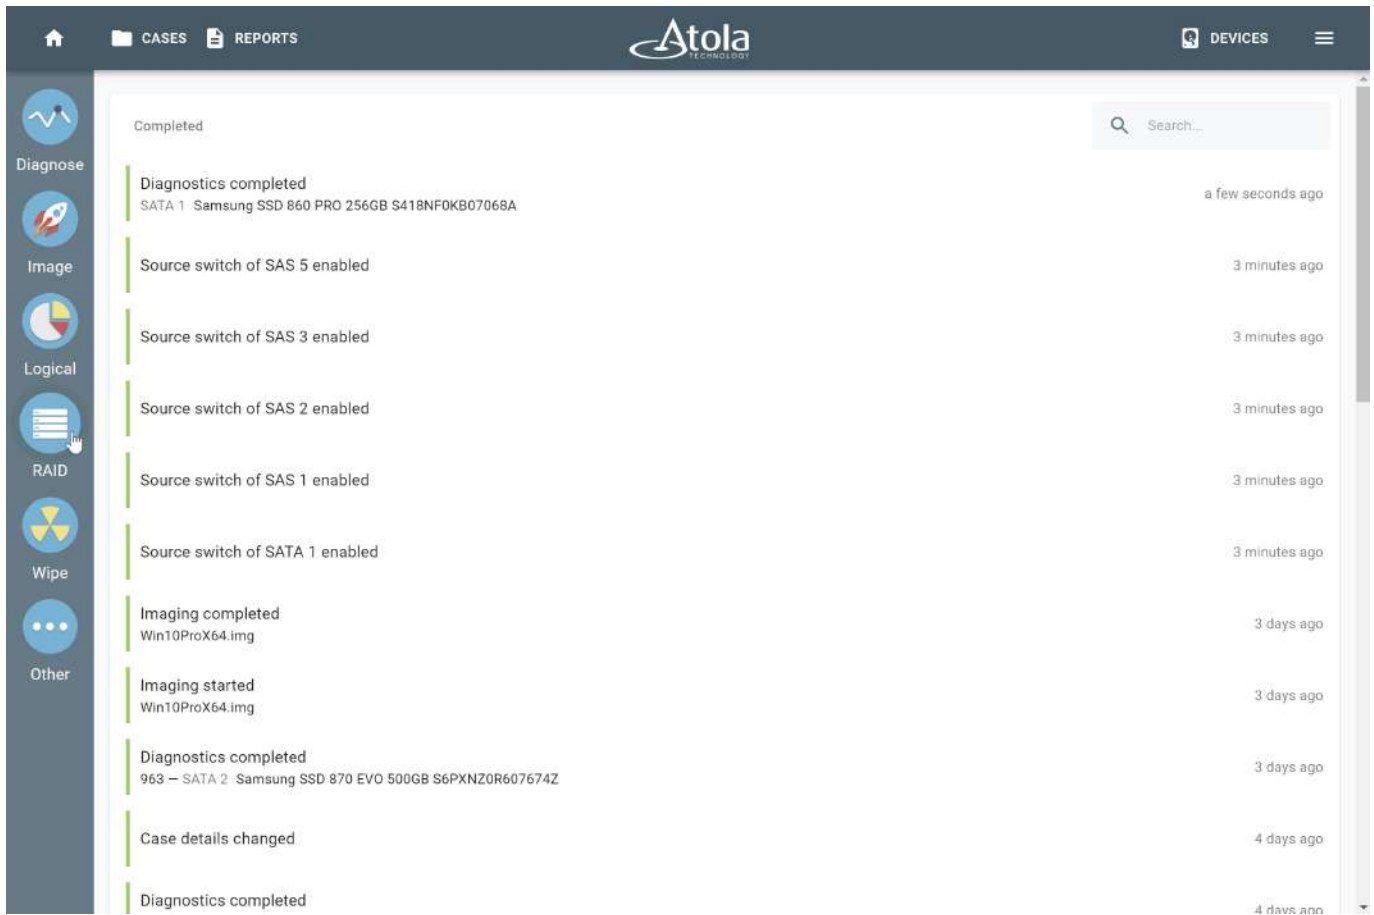
- Get available Source drives
- Optionally: diagnose drives
- Start imaging
- Check task status
- Stop running task (if necessary)
- Download created report

RAID 0 imaging


Atola TaskForce can automatically detect a RAID 0 configuration, assemble and image such RAID array.

Assembling a RAID 0 with unknown configuration

1. Connect the drives to the TaskForce hardware unit's ports that are switched to the Source mode.
2. On the left in the TaskForce main window, click **RAID**.



3. On the **Select source device** panel, select the drives that make up a RAID array and click **Continue**.



Selecting source RAID drives.

You can see key RAID configuration parameters at the top of the page:

- Order of drives/images
- RAID type
- Start LBA
- Block size
- Block order

Before you enter other values either manually or by applying Autodetection module results, the most commonly used values are displayed in the given fields.

The **Autodetection module** starts running immediately.

RAID 0 autodetection

Stage 1: TaskForce is reading data on the drives to identify the RAID type.

To change the order of drives in the array, simply drag a drive to its new position. To remove a device from the current array, grab it and put it into the bin.

The screenshot displays the Atola TaskForce software interface. On the left is a vertical sidebar with icons for Diagnose, Image, Logical, RAID (highlighted), Wipe, and Other. The main window is titled 'RAID configuration' and shows settings for RAID type (RAID 5), Start LBA (0), Block size (64 KB (most frequent)), and Block order (Left asymmetric). Below these settings is a table of six drives:

Order	Device	Capacity	File System
1	GIGABYTE GP-GSTFS31120GNTD SN194908950724	120 GB	ext4
2	WDC WDS120G1G0A-00SS50 171602A01A4D	120 GB	
3	JAJS300M120C 30013348333	120 GB	
4	SATA SSD PH191112003195	120 GB	
5	SPCC Solid State Disk 201813030030070	120 GB	
6	Samsung SSD 860 PRO 256GB S418NF0KB07068A	256 GB	

Below the table are buttons for 'ADD DEVICES', 'ADD MISSING DEVICE', and a 'Drag device here to remove' bin. The bottom section shows 'PARTITIONS' and 'HEX VIEWER' tabs, with a message 'No partitions found' and 'Supported file systems: NTFS, ext4/3/2'. On the right, the 'Autodetection' panel shows 'Stage 1. Analyzing devices' at 14% progress, with 'Processed sectors: 425,984'. A 'PAUSE' button is at the top right of the autodetection panel. At the bottom of the interface are 'GO TO IMAGE' and 'UNMOUNT RAID' buttons.

Analyzing devices.

Stage 2: TaskForce goes through thousands of possible variants of RAID parameters.

Click **Apply** as soon as a **Possible configuration** tile appears.

RAID configuration

RAID type: RAID 5 Start LBA: 0 Block size: 64 KB (most frequent) Block order: Left asymmetric

1	GIGABYTE GP-GSTFS31120GNTD SN194908950724	120 GB	ext4	=
2	WDC WDS120G1G0A-00SS50 171602A01A4D	120 GB		=
3	JAJS300M120C 30013348333	120 GB		=
4	SATA SSD PH191112003195	120 GB		=
5	SPCC Solid State Disk 201813030030070	120 GB		=
6	Samsung SSD 860 PRO 256GB S41BNF0KB07068A	256 GB		=

ADD DEVICES ADD MISSING DEVICE Drag device here to remove

Autodetection PAUSE

Background analysis of RAID devices/images.
It shows hints to help you find RAID configuration faster.

Stage 2. Detecting RAID configuration 26%
Checked variants: 2,413 out of 9,360

Possible configuration APPLY
RAID 0, block size: 256 KB, start LBA: 0

PARTITIONS HEX VIEWER

No partitions found

Supported file systems: NTFS, ext4/3/2

GO TO IMAGE UNMOUNT RAID

Possible RAID configuration.

After you click **Apply**, TaskForce automatically applies the suggested configuration and checks the file system for partitions. At the bottom of the screen, a preview of the partitions is available.

Observing files and folders in the preview confirms the detected RAID 0 configuration is correct.

Imaging RAID 0 array

1. After RAID configuration is successfully applied, click **Go to Image** to proceed with imaging.

RAID configuration

RAID type: RAID 0 (Stripe) Start LBA: 0 Block size: 256 KB (LaCie, ServeRAID-8*)

ID	Device	Capacity	Format
1	GIGABYTE GP-GSTFS31120GNTD SN194908950724	120 GB	ext4
2	SATA SSD PH191112003195	120 GB	
3	SPCC Solid State Disk 201813030030070	120 GB	
4	WDC WDS120G1G0A-00SS50 171602A01A4D	120 GB	
5	Samsung SSD 860 PRO 256GB S418NF0K807068A	256 GB	
6	JAJS300M120C 30013348333	120 GB	

ADD DEVICES Drag device here to remove

Autodetection RESTART

Background analysis of RAID devices/images. It shows hints to help you find RAID configuration faster.

Possible configuration
RAID 0, block size: 256 KB, start LBA: 0 APPLY

PARTITIONS HEX VIEWER


Partition	Start LBA	Name	Type	Size
0 No label ext4 · 719 GB	0	lost+found	Folder	
		Documents	Folder	
		Downloads	Folder	
		Sentinel_LDK_RedHat_and...	Folder	
		Books	Folder	
		Graphics	Folder	
		Manuals	Folder	
		Specifications	Folder	
		mpt3sas	Folder	

GO TO IMAGE UNMOUNT RAID

The Go to image button.


2. Select your target device or network folder and click **Continue**.
3. To launch your imaging session, click **Start**.

RAID imaging may take much longer than imaging a drive. You can optimize the imaging speed by using a fast target or a high-speed server.



How to make raid image.

TaskForce automatically generates reports for every session. The **Imaging completed** report will contain all RAID details as well as timestamps.



RAID 5 forensics: Automated reassembly and imaging

Autodetection of RAID 5 configuration

1. Connect the drives to the TaskForce hardware unit's ports that are switched to the Source mode.
2. On the left in the TaskForce main window, click **RAID**.

Atola TECHNOLOGY

CASES REPORTS DEVICES

Diagnose Image Logical RAID Wipe Other


Completed

Search...

Diagnostics completed	a few seconds ago
SATA 1 Samsung SSD 860 PRO 256GB S418NF0KB07068A	
Source switch of SAS 5 enabled	3 minutes ago
Source switch of SAS 3 enabled	3 minutes ago
Source switch of SAS 2 enabled	3 minutes ago
Source switch of SAS 1 enabled	3 minutes ago
Source switch of SATA 1 enabled	3 minutes ago
Imaging completed	3 days ago
Win10ProX64.img	
Imaging started	3 days ago
Win10ProX64.img	
Diagnostics completed	3 days ago
963 - SATA 2 Samsung SSD 870 EVO 500GB S6PXXZ0R607674Z	
Case details changed	4 days ago
Diagnostics completed	4 days ago

Initiating work with RAID.

3. On the **Select source device** panel, select the drives that make up a RAID array and click **Continue**.



Selecting the drives that make up a RAID.

To assemble a RAID from images instead of drives or to use a combination of drives and images, on the **Select source device** panel, expand the **File** section and click **Select file**. Then browse and select images.

After you selected the drives that make up a RAID array, TaskForce redirects you to the RAID configuration screen. It consists of three parts:

- The selected devices (and/or image files) are shown in the top **RAID configuration** part.
- The **RAID Partitions** viewer below it provides a preview of partitions and files within them, once RAID has been successfully assembled.
- The **Autodetection** module in the right-hand part of the screen immediately starts running and produces an output of RAID configuration suggestions.

Autodetection module reads data from all the selected devices and/or images to detect these RAID parameters:

- Order of drives/images.
- RAID type.
- Start LBA.
- Block size.
- Block order.

Should the configuration be known, these parameters can also be set manually by the operator.

The time required for configuration detection can vary from a few seconds to a few hours depending on the numbers of drives involved, RAID volume and type, and how metadata is distributed on the drives in the RAID. In certain cases, Autodetection may produce several configuration suggestions, which can be applied one by one to find the exact match. TaskForce's Autodetection is based on heuristics algorithms that help speed up the variant check.

4. Click the **Apply** button to apply the configuration suggested by the Autodetection module.

RAID configuration

RAID type: RAID 5 | Start LBA: 0 | Block size: 512 KB (Adaptec, Li...) | Block order: Left asymmetric

Device	Model	Capacity	Format
1	GIGABYTE GP-GSTFS31120GNTD S	120 GB	
2	JAJS300M120C 30013348333	120 GB	
3	SATA SSD PH191112003195	120 GB	ext4
4	SPCC Solid State Disk 2018130300	120 GB	
5	Patriot Burst 5680076A1EEE00402	120 GB	MBR Unknown

ADD DEVICES | ADD MISSING DEVICE | Drag device here to remove

PARTITIONS | HEX VIEWER

Partition	Start LBA	Size
0 1 ext ext4 - 270 GB	2,048	
1 2 ext ext4 - 210 GB	526,725,120	

Partition 0 / home / atola

Name	Type	Size
Downloads	Folder	
Sentinel_LDK_RedHat_and...	Folder	
sg3_utils-1.44	Folder	

Autodetection | RESTART

Background analysis of RAID devices/images. It shows hints to help you find RAID configuration faster.


Possible configuration
 RAID 5, block size: 512 KB, order: Left asymmetric
 APPLY

GO TO IMAGE | UNMOUNT RAID

Detected RAID 5 configuration applied


Imaging selected partitions of RAID 5

1. To adjust the imaging settings and define the target for the image, in the left bottom corner of the screen, click **Go to image**.
2. Select the target for the imaging session. Both a local server and a target device in Storage mode can be used for imaging of a RAID array.
3. Click **Create file**.
4. In the **Create image file** dialog, fill out the image details and click **Create**.




Imaging RAID. E01 details.

5. On the **Settings** page, click **Change** and then click the settings of an imaging pass.



Imaging settings.


6. In **Edit imaging pass** dialog, select the individual partitions to be imaged if selective imaging is required and click **Save**.



Selecting partitions in imaging settings


7. To launch the imaging session, click **Start**.

TaskForce will be imaging RAID 5 array or its partitions as configured in the imaging settings.




Imaging RAID 5 in progress

At the end of the imaging session, TaskForce produces an **Imaging completed** report with all the details of the source drives, the RAID configuration, the target, the partition, the timestamps, etc.



Selecting RAID 5 image files.

5. On the **Select source device** panel, check source image files and then click **Continue**.



Selecting RAID 5 image files.

6. Autodetection module starts running automatically to find a suitable RAID configuration. You know that there is an image file missing from the selection (it may have been lost or damaged), so click **Add missing device** underneath the list of image files.

RAID configuration

RAID type: **RAID 5** Start LBA: **0** Block size: **64 KB (most frequ...** Block order: **Left asymmetric**

1	2.E01	120 GB	MBR	Unknown	≡
2	1.E01	120 GB			≡
3	3.E01	120 GB	ext4		≡
4	4.E01	120 GB			≡

ADD DEVICES **ADD MISSING DEVICE** Drag device here to remove

PARTITIONS **HEX VIEWER**

No partitions found

Supported file systems: NTFS, ext4/3/2

GO TO IMAGE **UNMOUNT RAID**

Autodetection **RESUME**

Background analysis of RAID devices/images. It shows hints to help you find RAID configuration faster. The most probable configurations will appear at the top.

Stage 1. Analyzing devices
Processed sectors: 32,768 1%

Adding missing RAID 5 drive.

- Autodetection will recommence when a new image file or device is added. If the RAID is type 5, TaskForce will be able to identify the right configuration and reassemble the RAID by using the redundancy. When a possible configuration is found, click **Apply**.

RAID configuration

RAID type: RAID 5 Start LBA: 0 Block size: 64 KB (most frequent) Block order: Left asymmetric

1	2.E01	120 GB	MBR	Unknown	=
2	1.E01	120 GB			=
3	3.E01	120 GB	ext4		=
4	4.E01	120 GB			=
5	Missing				=

ADD DEVICES ADD MISSING DEVICE Drag device here to remove

PARTITIONS HEX VIEWER

No partitions found

Supported file systems: NTFS, ext4/3/2

Autodetection PAUSE

Background analysis of RAID devices/images. It shows hints to help you find RAID configuration faster. The most probable configurations will appear at the top.

Stage 2. Detecting RAID configuration 39%
Checked variants: 2,276 out of 5,760

Possible configuration
RAID 5, block size: 512 KB, order: Left asymmetric APPLY

GO TO IMAGE **UNMOUNT RAID**

Applying the found RAID configuration.

The configuration application automatically changes the order of the drives/images, the detected RAID type, block size and order is applied automatically, too. In the bottom part of the screen, the found partitions are available for preview.

- To proceed with imaging, click the **Go to image** button and follow the instruction in subsequent screens to select the target and adjust imaging settings.

Atola
TECHNOLOGY

Diagnose
 Image
 Logical
 RAID
 Wipe
 Other

RAID configuration

RAID type: **RAID 5**
Start LBA: **0**
Block size: **512 KB (Adaptec, L...)**
Block order: **Left asymmetric**

1	1.E01	120 GB	=
2	Missing		=
3	3.E01	120 GB ext4	=
4	4.E01	120 GB	=
5	2.E01	120 GB MBR Unknown	=

ADD DEVICES
 ADD MISSING DEVICE
 Drag device here to remove

PARTITIONS

Partition	Start LBA	Name	Type	Size
0 1 ext ext4 - 270 GB	2,048	lost+found	Folder	
1 2 ext ext4 - 210 GB	526,725,120	Kernel-sources	Folder	
		home	Folder	
		usr	Folder	

GO TO IMAGE

Autodetection

PAUSE

Background analysis of RAID devices/images. It shows hints to help you find RAID configuration faster. The most probable configurations will appear at the top.

Stage 2. Detecting RAID configuration
 51%
 Checked variants: 2,975 out of 5,760


Possible configuration
 RAID 5, block size: 512 KB, order: Left asymmetric

APPLIED

UNMOUNT RAID

Preview the found partitions and proceed to image

During the imaging, TaskForce takes advantage of the redundancy to create an image of the RAID despite one of the devices is missing.



Imaging RAID 5 with a missing drive.

The imaging report clearly indicates which of the devices that make up the RAID array was missing. Other RAID parameters, such as type, block size and order are also clearly reflected in the report.

The screenshot shows the Atola TaskForce software interface. The top navigation bar includes 'CASES', 'REPORTS', and 'DEVICES'. The sidebar on the left contains icons for 'Diagnose', 'Image', 'Logical', 'RAID', 'Wipe', and 'Other'. The main panel displays RAID 5 configuration details:

- Devices:** 5, Block size: 512 KB, Start LBA: 0, Left asymmetric, 480 GB
- Case ID:** Not assigned
- RAID type:** RAID 5
- Creation date:** 08/08/2023 10:58 AM
- TaskForce serial:** 95959595
- Software version:** 2023.7
- TaskForce IP:** 10.0.0.85
- Block size:** 512 KB (524,288 bytes)
- Start LBA:** 0
- Block order:** Left asymmetric
- Case investigator:**
- Case description:**
- Devices in RAID:**
 - 1.E01: 120 GB, S:/images/RAID/Adaptec 52445/750GB RAID-5-ss512 ext4
 - Missing: 120 GB
 - 3.E01: 120 GB, S:/images/RAID/Adaptec 52445/750GB RAID-5-ss512 ext4
 - 4.E01: 120 GB, S:/images/RAID/Adaptec 52445/750GB RAID-5-ss512 ext4
 - 2.E01: 120 GB, S:/images/RAID/Adaptec 52445/750GB RAID-5-ss512 ext4

Below the RAID details, the 'Imaging completed' section shows the following information:

- Targets:** WDC WD100EFAX-68LHPN0 JEK1ATGZ 10 TB SATA 3
- Passes:** 5 passes with user customized settings
- Scheduled sectors:** 937,762,816
- Imaged sectors:** 937,762,816
- Errors:** 0
- Sector size:** 512
- Hash method:** Linear
- MD5:** c862a066e6ec6a2d412d6049dba0275dc

At the bottom, there are buttons for 'PRINT', 'GO TO CASE', and 'GO TO IMAGE', along with navigation arrows.


Imaging report clearly indicates the missing device in RAID 5.

Imaging RAID 5 array with 2 damaged drives

Even if a RAID 5 array contains errors, Atola TaskForce is able to detect its parameters and image such RAID.

Reassemble RAID 5

1. Connect the drives to the ports of the TaskForce hardware unit. Make sure the ports are in *Source mode*.
2. On the left in the TaskForce main window, click **RAID**.
3. Select the drives that make up the RAID array.
4. Click **Continue**.



Selecting RAID members

Detecting errors

The TaskForce autodetection module starts running immediately, after you select the RAID devices.

Stage 1: TaskForce reads data on the drives to identify RAID type. If it runs across an error, TaskForce displays error tags next to the respective RAID member.

To see the number of errors encountered on a RAID member, simply hover the cursor over the error tag.

RAID configuration

RAID type: RAID 5 | Start LBA: 0 | Block size: 64 KB (most frequent) | Block order: Left asymmetric

Device ID	Device Name	Capacity	MBR	Status
1	WDC WD10JPVT-75A1YT0 WXF1C12Y009	1 TB	MBR	Unknown
2	ST500DM002-1BD142 S2A76EWW	500 GB	MBR	Unknown Error
3	ST9500423AS S2V0EP9P	500 GB	MBR	Error
4	HGST HTS541010A9E6...0811JD10424A1	1 TB		
5	TOSHIBA MQ01ABD100 Z5APTWCIT	1 TB		

Read error count: 3

ADD DEVICES | ADD MISSING DEVICE | Drag device here to remove

PARTITIONS

Partition	Start LBA	Name	Type	Size
0		Microsoft reserved partition	Unknown	134 MB

GO TO IMAGE | UNMOUNT RAID

Error tags

Stage 2: TaskForce goes through thousands of possible RAID configurations to identify a suitable one.

Once the configuration is detected, click the **Apply** button.

RAID configuration

RAID type: RAID 5 | Start LBA: 0 | Block size: 64 KB (most frequent) | Block order: Left asymmetric

Device ID	Device Name	Capacity	MBR	Status
1	WDC WD10JPVT-75A1YT0 WXF1C12Y009	1 TB	MBR	Unknown
2	ST500DM002-1BD142 S2A76EWW	500 GB	MBR	Unknown Error
3	ST9500423AS S2V0EP9P	500 GB	MBR	Error
4	HGST HTS541010A9E6...0811JD10424A1	1 TB		
5	TOSHIBA MQ01ABD100 Z5APTWCIT	1 TB		

ADD DEVICES | ADD MISSING DEVICE | Drag device here to remove

PARTITIONS

Partition	Start LBA	Name	Type	Size
0		Microsoft reserved partition	Unknown	134 MB

GO TO IMAGE | UNMOUNT RAID

Autodetection [PAUSE]

Background analysis of RAID devices/images. It shows hints to help you find RAID configuration faster. The most probable configurations will appear at the top.

Stage 1. Analyzing devices
Processed sectors: 192,512 6%

Stage 2. Detecting RAID configuration
Checked variants: 231 out of 5,760 3%

Possible configuration
RAID 5, block size: 64 KB, order: Left symmetric

APPLY

Applying the suggested RAID configuration

After you click **Apply**, TaskForce automatically applies the suggested configuration and checks the file system for partitions.

Despite read errors, TaskForce can mount the partitions for preview by rebuilding the data in the bad sectors using data redundancy inherent to this RAID type.

Imaging RAID 5 array with errors

1. After RAID configuration is successfully applied, click the **Go to image** button to proceed with imaging.

The screenshot displays the Atola TaskForce RAID configuration interface. The RAID configuration section shows a RAID 5 array with 5 drives. Drive 3 (ST9500423AS) and Drive 5 (ST500DM002) show 'Error' status. The Autodetection section shows 'Stage 2. Detecting RAID configuration' at 28% completion, with a 'Possible configuration' of RAID 5, block size: 64 KB, order: Left symmetric, marked as 'APPLIED'. The PARTITIONS section shows a 'New Volume' (NTFS - 429 GB) at LBA 264,192. The 'GO TO IMAGE' button is highlighted with an orange arrow.

RAID type	Start LBA	Block size	Block order
RAID 5	0	64 KB (most frequent)	Left symmetric


Drive	Model	Capacity	Status
1	WDC WD10JPVT-75A1YT0 WXF1C12Y009	1 TB	MBR Unknown
2	TOSHIBA MQ01ABD100 Z5APTWCIT	1 TB	
3	ST9500423AS S2V0EP9P	500 GB	Error
4	HGST HTS541010A9E6...0811JD10424A1	1 TB	
5	ST500DM002-1BD142 S2A76EWW	500 GB	MBR Unknown Error

ADD DEVICES ADD MISSING DEVICE Drag device here to remove

Partition	Start LBA	Name	Type	Size
0	Microsoft reserved partition Unknown - 134 MB	34		
1	New Volume NTFS - 429 GB	264,192		


GO TO IMAGE UNMOUNT RAID

2. Select your target device or folder and click **Continue**.
3. To launch your imaging session, click **Start**.
4. When the TaskForce system encounters an error, it automatically reconstructs the missing data, using the data in the parity blocks on the remaining RAID members.



Coping with errors.

Thus TaskForce can recover the full image. No operator involvement is necessary.



Imaging RAID 5 array.

In case 2 drives of the RAID 5 array contain errors in the same sector, TaskForce can still recover the data from those, using the [multipass imaging system](#).

Imaging report

TaskForce automatically generates the **Imaging completed** report with all RAID details and timestamps.

TaskForce managed to successfully reconstruct and image all the data. That is why the number of errors in the **Imaging completed** report is zero.

The screenshot displays the Atola TaskForce software interface. On the left is a vertical sidebar with icons for Diagnose, Image, Logical, RAID, Wipe, and Other. The main window is titled 'REPORTS' and shows a detailed report for a RAID 5 array. The report includes the following information:

- Devices:** 5, Block size: 64 KB, Start LBA: 0, Left symmetric, 2 TB, RAID 5
- Case ID:** Not assigned
- Creation date:** 08/12/2023 3:42 PM
- TaskForce serial:** 95959595
- RAID type:** RAID 5
- Block size:** 64 KB (65,536 bytes)
- Case investigator:**
- Devices in RAID:**
 - WDC WD10JPVT-75A1YT0 WXFIC12V0097 1 TB SATA 3
 - TOSHIBA M001ABD100 Z5APTWCIT 1 TB SATA 5
 - ST9500423AS S2VGE3P 500 GB SATA 1
 - HGST HTS541010A9E680 160811JD10424A14BLPS 1 TB SATA 4
 - ST500DM002-1BD142 S2A78EWW 500 GB SATA 6

Imaging completed

Targets

HGST HDN726060ALE610 NAGE1P7X 6.0 TB SATA 2

Passes: 5 passes with user customized settings

Scheduled sectors: 3,907,092,480

Imaged sectors: 3,907,092,480

Errors: 0

Sector size: 512

Hash method: Linear

MD5: eebf0e81ea8a7af367e877a91404938f

Calculated range: 0 - 3,907,092,479

At the bottom of the report, there are buttons for 'PRINT', 'GO TO CASE', and 'GO TO IMAGE', along with navigation arrows.

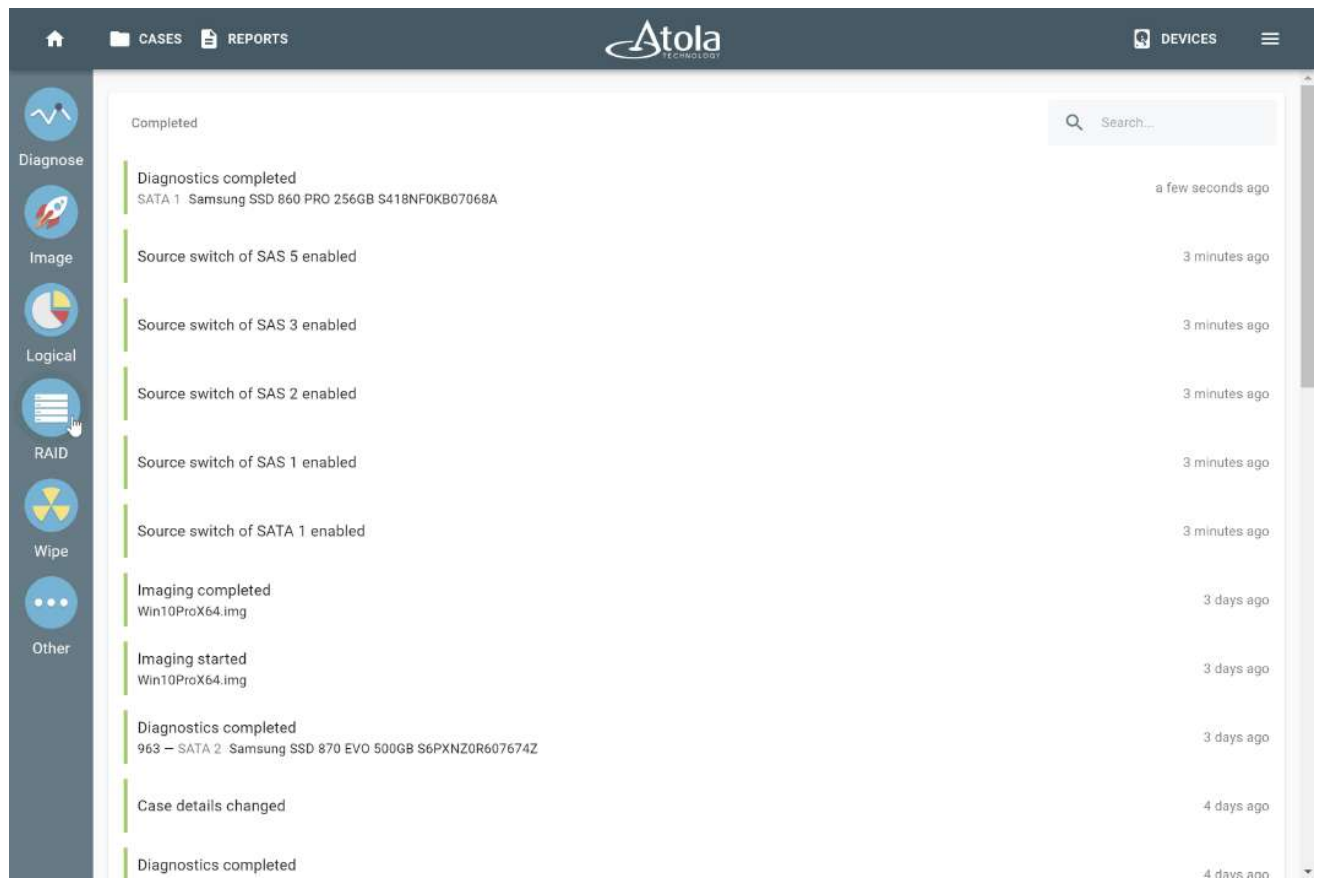
Imaging report.

RAID 6 with unknown configuration

Atola TaskForce lets you automatically detect all parameters of a RAID 6 array, preview its contents, and then create its full physical image or perform logical imaging of only selected partitions, folders, and files.


To reassemble and image a RAID 6 array with an unknown configuration, do the following:

1. On the left in the TaskForce main window, click **RAID**.



Initiating work with RAID.

2. On the **Select source device** panel, select the devices that make up a RAID array. They can be physical drives of all [supported types](#), connected to the TaskForce hardware unit, image files (raw, E01 or AFF4 image files), or a combination of both.



Selecting RAID members.

3. Once all the RAID members are selected, click **Continue**.
4. TaskForce redirects you to the **RAID configuration** screen and immediately starts the **Autodetection** module.

The screenshot displays the RAID configuration interface in Atola TaskForce. The top navigation bar includes 'CASES', 'REPORTS', the 'Atola' logo, and 'DEVICES'. A left sidebar contains icons for 'Diagnose', 'Image', 'Logical', 'RAID', 'Wipe', and 'Other'. The main content area is titled 'RAID configuration' and includes settings for RAID type (RAID 5), Start LBA (0), Block size (64 KB), and Block order (Left asymmetric). Below these settings is a table of detected devices:

Order	Device	Capacity	MBR	Unknown	Order
1	TOSHIBA HDWD110 709NZGYMS	1 TB	MBR	Unknown	=
2	WDC WD10EZRX-22HTKB0 WDWCC	1 TB	MBR	Unknown	=
3	TOSHIBA MQ01ABD100 Z5APTWC1	1 TB	MBR	Unknown	=
4	TOSHIBA HDWD110 89EJY0PFS	1 TB			=
5	HGST HTS541010A9E6...0811JD10-	1 TB			=
6	TOSHIBA MQ01ABD100 Y3Q7P74Y	1 TB			=

Below the device list are buttons for 'ADD DEVICES', 'ADD MISSING DEVICE', and a 'Drag device here to remove' instruction. The 'PARTITIONS' tab is active, showing a table with one partition:

Partition	Start LBA	Name	Type	Size
0	34	Microsoft reserved par...		Unknown - 17 MB

The 'HEX VIEWER' tab is also visible. At the bottom, there are buttons for 'GO TO IMAGE', 'GO TO LOGICAL', and 'UNMOUNT RAID'. On the right, the 'Autodetection' section shows 'Stage 1. Analyzing devices' at 23% progress, with 'Processed sectors: 700,416'.

The RAID configuration screen in Atola TaskForce.

5. **Optional:** If you know some of the RAID parameters, such as *RAID type*, *Start LBA*, *Block size*, and *Block order*, you can select them manually from the lists at the top of the **RAID configuration** section. To change the device order, drag the devices up or down, as needed.

The screenshot shows the Atola software interface. On the left is a vertical sidebar with icons for Diagnose, Image, Logical, RAID, Wipe, and Other. The main window is titled 'RAID configuration' and features several dropdown menus: RAID type (RAID 6), Start LBA (0), Block size (64 KB (most fr...)), Block order (Left asymmetr...), and Parity block order (PQ). Below these is a table of detected RAID members:

Order	Device	Size	MBR	Unknown	Parity
1	TOSHIBA HDWD110 709NZGYMS	1 TB	MBR	Unknown	=
2	WDC WD10EZRZ-22HTXB0 WDWCC	1 TB	MBR	Unknown	=
3	TOSHIBA MQ01ABD100 Z5APTWC	1 TB	MBR	Unknown	=
4	HGST HTS541010A9E6...0811JD10	1 TB			
5	TOSHIBA HDWD110 89EJY0PFS	1 TB			=
6	TOSHIBA MQ01ABD100 Y3Q7P74Y	1 TB			=

Below the table are buttons for 'ADD DEVICES', 'ADD MISSING DEVICE', and a 'Drag device here to remove' area. At the bottom of the RAID configuration section are 'GO TO IMAGE' and 'GO TO LOGICAL' buttons. To the right, the 'Autodetection' section shows 'Stage 1. Analyzing devices' at 75% completion, with 'Processed sectors: 2,297,856'. Below this is an 'UNMOUNT RAID' button. The bottom of the interface shows a 'PARTITIONS' section with a table:

Partition	Start LBA	Name	Type	Size
0	34	Microsoft reserved par...		Unknown - 17 MB

Below the partitions table is a message: 'Unknown file system. It cannot be parsed.'

Lists with RAID parameters for manual selection.

- Optional: If you know that one or two RAID members are missing, choose *RAID 6* from the *RAID type* list and then click the *Add missing device* button once or twice respectively. For details, see [RAID 6 with two missing devices](#).

RAID configuration

RAID type: RAID 6 | Start LBA: 0 | Block size: 64 KB (most fr...) | Block order: Left asymmetric | Parity block order: PQ

Device ID	Device Name	Capacity	MBR	Unknown	Parity
1	TOSHIBA HDWD110 89EJY0PFS	1 TB	MBR	Unknown	=
2	WDC WD10EZRZ-22HTKB0 WDWC	1 TB	MBR	Unknown	=
3	TOSHIBA HDWD110 709NZGYMS	1 TB			=
4	HGST HTS541010A9E6...0811JD104	1 TB			=
5	TOSHIBA MQ01ABD100 Z5APTWCIT	1 TB			=

ADD DEVICES | **ADD MISSING DEVICE** | Drag device here to remove

PARTITIONS | HEX VIEWER

Partition	Start LBA	Name	Type	Size
0		Microsoft reserved par...		
		Unknown - 134 MB		34

Unknown file system. It cannot be parsed.

GO TO IMAGE | GO TO LOGICAL | UNMOUNT RAID

Adding a missing device to a RAID.

- Wait for a couple of minutes while the module is checking thousands of variants to find a possible RAID configuration. The Autodetection progress is displayed on the right side of the screen. To learn more about our heuristic RAID autodetection algorithm, see [How autodetection of an unknown RAID works](#).
- Once TaskForce has detected a possible configuration, click **Apply**. TaskForce automatically changes the order of the devices, applies the detected RAID type and other parameters, and then parses and verifies the file system of an array.

RAID configuration

RAID type: RAID 6 | Start LBA: 0 | Block size: 64 KB (most fr...) | Block order: Left asymmetr... | Parity block order: PQ

Device ID	Device Name	Capacity	MBR	Unknown	Order
1	TOSHIBA HDWD110 709NZGYMS	1 TB	MBR	Unknown	=
2	WDC WD10EZRX-22HTKB0 WDWCC	1 TB	MBR	Unknown	=
3	TOSHIBA MQ01ABD100 Z5APTWC1	1 TB	MBR	Unknown	=
4	TOSHIBA HDWD110 89EJY0PFS	1 TB			=
5	HGST HTS541010A9E6...0811JD10	1 TB			=
6	TOSHIBA MQ01ABD100 Y3Q7P74Y	1 TB			=

ADD DEVICES | ADD MISSING DEVICE | Drag device here to remove

Autodetection [PAUSE]

Background analysis of RAID devices/images. It shows hints to help you find RAID configuration faster. The most probable configurations will appear at the top.

Stage 2. Detecting RAID configuration 0%
Checked variants: 17 out of 51,840

Possible configuration
RAID 6, block size: 64 KB, order: Left symmetric [APPLY]

PARTITIONS | HEX VIEWER

Partition	Start LBA	Name	Type	Size
0	34	Microsoft reserved par...		Unknown - 17 MB

Unknown file system. It cannot be parsed.

GO TO IMAGE | GO TO LOGICAL | UNMOUNT RAID

Applying a possible RAID configuration.

- To make sure that the suggested RAID configuration is resulting in tangible data, check the **Partitions** section at the bottom left and preview the contents of a reassembled RAID, including partitions, folders, and files. To help you decide whether the applied configuration is valid or not, TaskForce displays a small tag with the accuracy percentage next to the partition. In some cases, TaskForce will suggest alternative configurations, and you can try to improve the accuracy percentage by applying these alternative configurations.

The screenshot shows the Atola RAID configuration software interface. The top navigation bar includes 'CASES', 'REPORTS', and 'DEVICES'. The left sidebar contains icons for 'Diagnose', 'Image', 'Logical', 'RAID', 'Wipe', and 'Other'. The main area is titled 'RAID configuration' and shows a RAID 6 setup with 6 drives. The autodetection section on the right shows 'Stage 2. Detecting RAID configuration' at 0% and a 'Possible configuration' of RAID 6, block size: 64 KB, order: Left asymmetric, which has been 'APPLIED'. Below the RAID configuration, there are tabs for 'PARTITIONS' and 'HEX VIEWER'. The 'PARTITIONS' tab shows a table of partitions, including a 'backups' partition (ExFAT - 4.0 TB, 100% full) and a 'System Volume Information' folder. The 'HEX VIEWER' tab shows a table of files, including 'System Volume Information', '\$WinREAgent', and '1'. At the bottom, there are buttons for 'GO TO IMAGE', 'GO TO LOGICAL', and 'UNMOUNT RAID'.

RAID configuration

RAID type: RAID 6 | Start LBA: 0 | Block size: 64 KB (most fr...) | Block order: Left asymmetr... | Parity block order: PQ

Drive	Model	Capacity	File System	Parity
1	WDC WD10EZRX-22HTKB0 WDWCC	1 TB	MBR	Unknown
2	HGST HTS541010A9E6...0811JD10	1 TB		
3	TOSHIBA MQ01ABD100 Y3Q7P74Y	1 TB	exFAT	
4	TOSHIBA HDWD110 89EJY0PFS	1 TB		
5	TOSHIBA MQ01ABD100 Z5APTWC1	1 TB	MBR	Unknown
6	TOSHIBA HDWD110 709NZGYMS	1 TB	MBR	Unknown

ADD DEVICES | ADD MISSING DEVICE | Drag device here to remove

Autodetection [PAUSE]

Background analysis of RAID devices/images. It shows hints to help you find RAID configuration faster. The most probable configurations will appear at the top.

Stage 2. Detecting RAID configuration 0%
Checked variants: 37 out of 51,840

Possible configuration
RAID 6, block size: 64 KB, order: Left asymmetric APPLIED

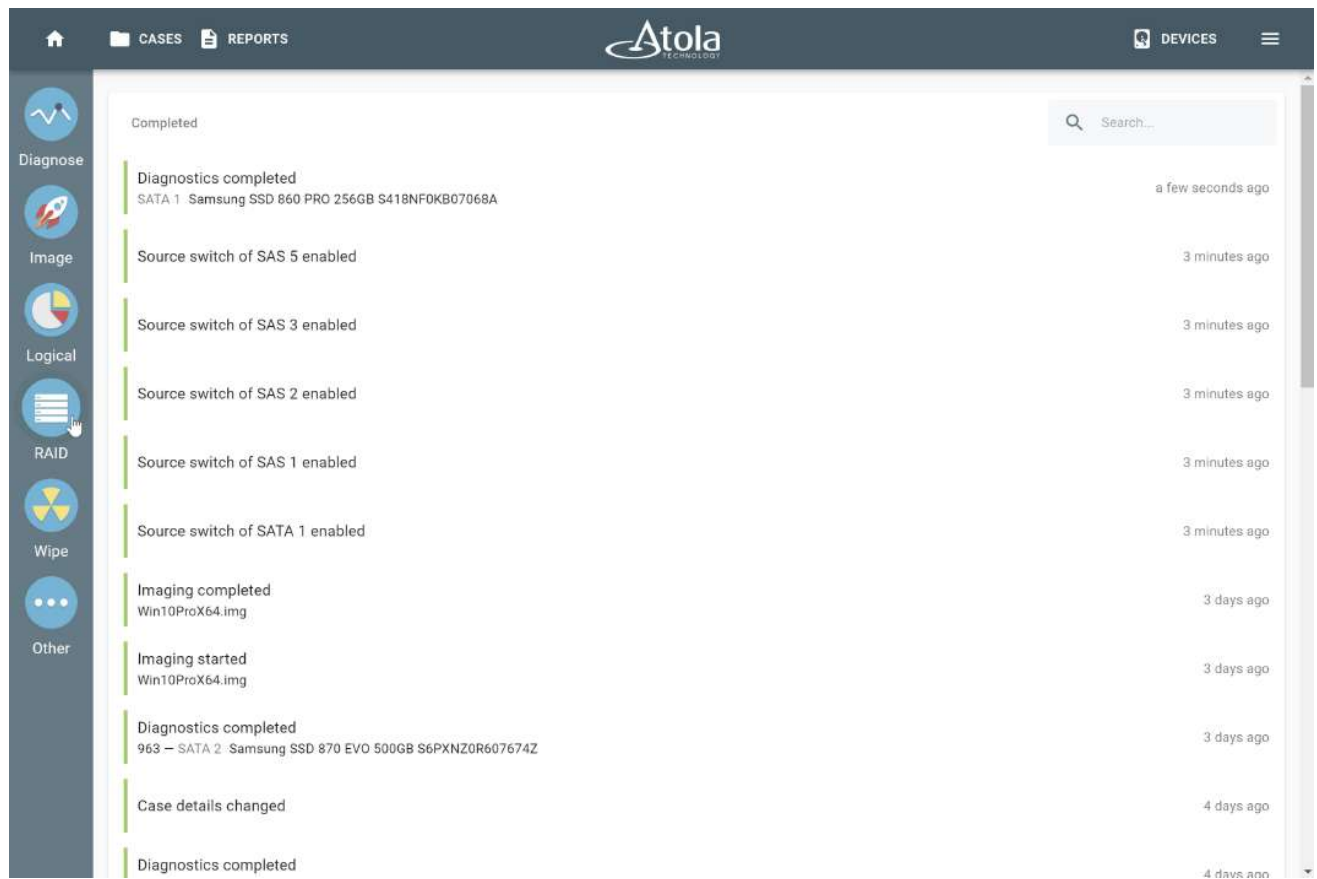
PARTITIONS | HEX VIEWER

Partition	Checked files	Valid files	Invalid files	Start LBA	Name	Type	Size
0	100	100	0	34	System Volume Informa...	Folder	
1	backups			32,768	\$WinREAgent	Folder	
	ExFAT - 4.0 TB	100%			1	Folder	

GO TO IMAGE | GO TO LOGICAL | UNMOUNT RAID


The preview of the RAID contents.

- Click **Go to image** to acquire a full bit-by-bit copy of the array to E01, AFF4, or raw image. Or click **Go to logical** to [copy only selected partitions, folders, or files](#) to L01 or ZIP.
- After imaging is finished, TaskForce automatically generates an **Imaging completed** report. It includes all the details of the source drives, the RAID configuration, the target, the partition, the timestamps and more.




Initiating work with RAID.

2. On the **Select source device** panel, select the devices that make up a RAID array. They can be physical drives of all [supported types](#), connected to the TaskForce hardware unit, image files (raw, E01 or AFF4 image files), or a combination of both.



Selecting RAID members.

- Once the available RAID members are selected, click **Continue**.
- TaskForce redirects you to the RAID reassembly screen and immediately starts the **Autodetection** process to find the suitable configuration.



The RAID configuration screen in Atola TaskForce.

- Since you know that one or two RAID members are missing, select *RAID 6* from the **RAID type** list and click the **Add missing device** button once or twice respectively. When checking variants of possible RAID configuration, TaskForce will take into account that there are missing RAID members.

RAID configuration

RAID type	Start LBA	Block size	Block order	Parity block order
RAID 6	0	64 KB (most fr...	Left asymmetr...	PQ

Device ID	Device Name	Capacity	MBR	Unknown	Order
1	TOSHIBA HDWD110 89EJY0PFS	1 TB	MBR	Unknown	=
2	WDC WD10EZRX-22HTKB0 WDWC	1 TB	MBR	Unknown	=
3	TOSHIBA HDWD110 709NZGYMS	1 TB			=
4	HGST HTS541010A9E6...0811JD10	1 TB			=
5	TOSHIBA MQ01ABD100 Z5APTWC	1 TB			=
6	Missing				=
7	Missing				=

Autodetection

Background analysis of RAID devices/images. It shows hints to help you find RAID configuration faster. The most probable configurations will appear at the top.

Stage 1. Analyzing devices 18%

Processed sectors: 565,248

PARTITIONS

Partition	Start LBA	Name	Type	Size
0	Unknown - 134 MB	Microsoft reserved par...		34

Buttons: GO TO IMAGE, GO TO LOGICAL, UNMOUNT RAID

Lists with RAID parameters for manual selection.

- Wait while the module is checking thousands of variants to find a possible RAID configuration. Autodetection progress is displayed on the right side of the screen. To learn more about our heuristic RAID autodetection algorithm, see [How autodetection of an unknown RAID works](#).
- Once TaskForce has detected a possible configuration with two missing devices, click **Apply**. TaskForce automatically changes the order of the devices, applies the detected RAID type and other parameters, and then parses and verifies the file systems of an array.

1

CASES

REPORTS

Atola

DEVICES

Diagnose

Image

Logical

RAID

Wipe

Other

RAID configuration

RAID type

Start LBA

Block size

Block order

Parity block order

RAID 6

0

64 KB (most fr...

Left asymmetr...

PQ

1	TOSHIBA HDWD110 89EJY0PFS	1 TB	MBR	Unknown	=
2	WDC WD10EZRX-22HTKB0 WDWC	1 TB	MBR	Unknown	=
3	TOSHIBA HDWD110 709NZGYMS	1 TB			=
4	HGST HTS541010A9E6...0811JD10	1 TB			=
5	TOSHIBA MQ01ABD100 Z5APTWC	1 TB			=
6	Missing				=
7	Missing				=

ADD DEVICES

ADD MISSING DEVICE

Drag device here to remove

PARTITIONS

HEX VIEWER

Partition	Start LBA	Name	Type	Size
0	Microsoft reserved par... Unknown - 134 MB	34		

GO TO IMAGE

GO TO LOGICAL

Autodetection

II

PAUSE

Background analysis of RAID devices/images. It shows hints to help you find RAID configuration faster. The most probable configurations will appear at the top.

Stage 2. Detecting RAID configuration

52%

Checked variants: 95,530 out of 181,440

Possible configuration

RAID 6, block size: 256 KB, order: Left symmetric, parity order: PQ

APPLY

UNMOUNT RAID

Applying a possible RAID configuration.

- To make sure that the suggested RAID configuration is correct, check the **Partitions** section at the bottom left and preview the contents of the reassembled RAID, including partitions, folders, and files. To help you decide whether the applied configuration is valid or not, TaskForce displays a small tag with the accuracy percentage near the partition.

RAID configuration

RAID type: RAID 6 | Start LBA: 0 | Block size: 256 KB (LaCie...) | Block order: Left symmetric | Parity block order: PQ

Drive	Model	Capacity	File System	Status
1	Missing			Missing
2	Missing			Missing
3	TOSHIBA MQ01ABD100 Z5APTWC1	1 TB		Present
4	TOSHIBA HDWD110 709NZGYMS	1 TB	NTFS	Present
5	HGST HTS541010A9E6...0811JD10	1 TB		Present
6	TOSHIBA HDWD110 89EJY0PFS	1 TB	MBR	Unknown
7	WDC WD10EZRZ-22HTKB0 WDWCC	1 TB	MBR	Unknown

Autodetection | PAUSE

Background analysis of RAID devices/images. It shows hints to help you find RAID configuration faster. The most probable configurations will appear at the top.

Stage 2. Detecting RAID configuration | 61%

Checked variants: 111,074 out of 181,440

Possible configuration

RAID 6, block size: 256 KB, order: Left symmetric, parity order: PQ | APPLIED

PARTITIONS | HEX VIEWER

Partition	Type	Size	Start LBA
0	Micro Unknown	34	
1	NTFS - 2.4 TB	264,192	

File Explorer:

Name	Type	Size
\$RECYCLE.BIN	Folder	
RAID_Manage_Win_v2.5...	Folder	
RR272x-win-v1.0.9.1229	Folder	


Buttons: GO TO IMAGE, GO TO LOGICAL, UNMOUNT RAID

The preview of the RAID contents.

- Click **Go to image** to acquire a full bit-by-bit copy of the array or **Go to logical** to copy [only selected partitions, folders, or files](#).
- After imaging is finished, TaskForce automatically generates an **Imaging completed** report. It includes all the details of the source drives, the RAID configuration, the target image, the timestamps and more.

To mount an unknown RAID 10 and image it:

1. On the left in the TaskForce main window, click **RAID**.
2. Select the drives that make up the RAID array and click **Continue**.



RAID 10. Selecting RAID members.

3. When the autodetection module (in the right part of the screen) comes up with a suggestion of RAID configuration, click **Apply**.

In rare cases, there can be more than one suggestion. Try them out one by one by clicking **Apply** next to each suggestion. If you know the type and parameters of the RAID, you can enter them manually in the drop-downs.

RAID configuration

RAID type: RAID 5 | Start LBA: 0 | Block size: 64 KB (most frequent) | Block order: Left asymmetric

Device	Model	Capacity	File System	Role
1	TOSHIBA HDWD110 709NZQYMS	1 TB	NTFS	Mirror A
2	ST2000DM008-2FR102 WFL4KFG0	2 TB	NTFS	Mirror A
3	Samsung SSD 850 PRO 1TB S2BBNWAG200266R	1 TB		Mirror B
4	TOSHIBA HDWD110 89EJY0PFS	1 TB		Mirror B

ADD DEVICES | ADD MISSING DEVICE | Drag device here to remove

PARTITIONS | HEX VIEWER

No partitions found

Supported file systems: NTFS, ext4/3/2, exFAT, XFS, HFS, HFS+, HFSX

Autodetection | RESTART

Background analysis of RAID devices/images. It shows hints to help you find RAID configuration faster. The most probable configurations will appear at the top.

Possible configuration
RAID 10, block size: 64 KB, RAID 1 groups: [2, 2]

APPLY

GO TO IMAGE | UNMOUNT RAID

RAID 10. Applying the detected configuration.

- TaskForce arranges the drives into respective groups and applies other RAID settings to mount the partitions, which you can browse through in the Partition preview below.

RAID configuration

RAID type: RAID 10 (Stripe of mirrors) | Start LBA: 0 | Block size: 64 KB (most frequent) | RAID 1 groups: 2

Device	Model	Capacity	File System	Role
1	TOSHIBA HDWD110 709NZQYMS	1 TB	NTFS	Mirror A
2	ST2000DM008-2FR102 WFL4KFG0	2 TB	NTFS	Mirror A
3	Samsung SSD 850 PRO 1TB S2BBNWAG200266R	1 TB		Mirror B
4	TOSHIBA HDWD110 89EJY0PFS	1 TB		Mirror B

ADD DEVICES | Drag device here to remove

PARTITIONS | HEX VIEWER

Partition	Start LBA	Partition 0 / photo																		
0	2,048	<table border="1"> <thead> <tr> <th>Name</th> <th>Type</th> <th>Size</th> </tr> </thead> <tbody> <tr> <td>D:\JL0189.DNG</td> <td>DNG</td> <td>19 MB</td> </tr> <tr> <td>D:\JL0190.DNG</td> <td>DNG</td> <td>19 MB</td> </tr> <tr> <td>D:\JL0206.DNG</td> <td>DNG</td> <td>19 MB</td> </tr> <tr> <td>D:\JL0209.DNG</td> <td>DNG</td> <td>19 MB</td> </tr> <tr> <td>D:\JL0210.DNG</td> <td>DNG</td> <td>19 MB</td> </tr> </tbody> </table>	Name	Type	Size	D:\JL0189.DNG	DNG	19 MB	D:\JL0190.DNG	DNG	19 MB	D:\JL0206.DNG	DNG	19 MB	D:\JL0209.DNG	DNG	19 MB	D:\JL0210.DNG	DNG	19 MB
Name	Type	Size																		
D:\JL0189.DNG	DNG	19 MB																		
D:\JL0190.DNG	DNG	19 MB																		
D:\JL0206.DNG	DNG	19 MB																		
D:\JL0209.DNG	DNG	19 MB																		
D:\JL0210.DNG	DNG	19 MB																		

GO TO IMAGE | UNMOUNT RAID

Autodetection | RESTART


Background analysis of RAID devices/images. It shows hints to help you find RAID configuration faster. The most probable configurations will appear at the top.

Possible configuration
RAID 10, block size: 64 KB, RAID 1 groups: [2, 2]

APPLIED


RAID 10. Previewing the partitions.

- Click the Go to image button and select the target.




RAID 10. Selecting target.

6. Check the settings. Please note that you can choose to image only part of the data (only one of the partitions).



RAID 10. Imaging.

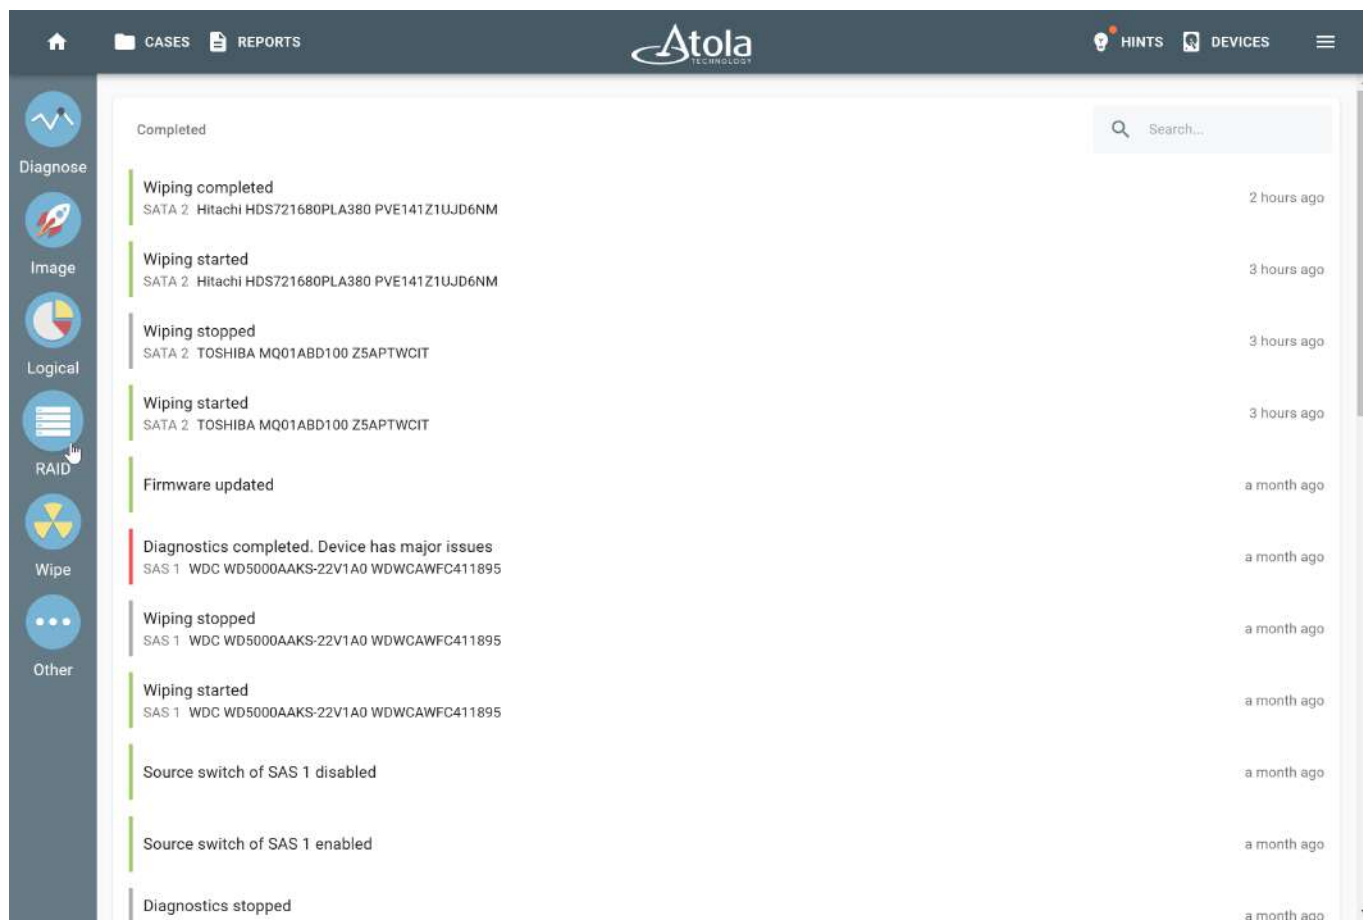
7. The automatically generated **Imaging report** contains the details of the RAID and its members, the target, hashes, signatures found as well as the timestamps.



Synology NAS RAID auto reassembly

Atola TaskForce can automatically detect the configuration of a Synology NAS RAID array and perform immediate forensic RAID reconstruction.


1. On the left in the TaskForce main window, click **RAID**.



Initiating work with RAID.

2. On the **Select source device** panel, select the devices that make up a RAID array.


They can be physical drives of all [supported types](#), connected to the TaskForce hardware unit, image files (raw, E01 or AFF4 image files), or a combination of both.



Selecting RAID members.

3. Once all the RAID members are selected, click **Continue**.
4. TaskForce redirects you to the **RAID configuration** screen and immediately starts the **Autodetection** module.

Despite the complexity of the Synology NAS RAID structure, the autodetection module instantly identifies the array by detecting controller metadata and applies the RAID configuration.



The 'Imaging started' report.

- After the imaging is finished, TaskForce generates an **Imaging completed** report. It includes all the details of the source devices in a RAID, the target(s), the number of imaged sectors, the timestamps and more.

Instant forensic RAID reconstruction: Linux mdadm

Atola TaskForce can automatically detect configuration of software RAID created with mdadm in Linux and perform immediate forensic RAID reconstruction.

Autodetection of the mdadm-created software RAID

- Connect the drives to ports in Source mode.
- On the left in the TaskForce main window, click **RAID**.
- Select the drives that make up a RAID array and click **Continue**.

Autodetection module has instantly recognized and applied the RAID configuration. TaskForce automatically identifies mdadm-created RAID arrays with great precision by detecting controller metadata.

This RAID's Start LBA is different from 0. TaskForce's autodetection module is trained to detect this parameter for different types of RAID arrays and mdadm versions.

A partition is displayed in the bottom part of the screen, confirming that the applied configuration is correct.

1

ADD CASE

CASES

REPORTS

Atola
TECHNOLOGY

DEVICES

Diagnose

Image

Logical

RAID

Wipe

Other

RAID configuration

RAID type

RAID 5

Start LBA

0

Block size

64 KB (most freque...

Block order

Left symmetric

1	WDC WDS240G2G0A-00JH30 20470	240 GB	MBR	Unknown	=
2	WDC WDS240G2G0A-00JH30 20485	240 GB	NTFS		=
3	WDC WDS240G2G0A-00JH30 20485	240 GB			=
4	WDC WDS240G2G0A-00JH30 20485	240 GB			=
5	WDC WDS240G2G0A-00JH30 20485	240 GB	MBR	Unknown	=

ADD DEVICES

ADD MISSING DEVICE

Drag device here to remove

PARTITIONS

HEX VIEWER

Partition	Start LBA
1	No label NTFS - 57 GB 100% 104,448
2	No label NTFS - 534 MB 111,595,520
3	Storage NTFS - 900 GB 100% 112,640,000

Partition 3

Name	Type	Size
PDFs	Folder	
Public forensic images	Folder	
RAID_data	Folder	

GO TO IMAGE

GO TO LOGICAL

UNMOUNT RAID

Autodetection

PAUSE

Background analysis of RAID devices/images. It shows hints to help you find RAID configuration faster. The most probable configurations will appear at the top.

Stage 2. Detecting RAID configuration

Checked variants: 5,209 out of 5,760

90%

Possible configuration


RAID 5, block size: 64 KB, order: Left

APPLIED

symmetric

The suggested RAID configuration applied, partitions preview shows partitions.


2. In the Logical imaging module, adjust your selection if needed.



Specify the partitions, files and folders to included or excluded in the imaging.


By default, logical imaging is set to image all files from a drive or a RAID. To fine-tune your selection, include or exclude what you need:

- All or selected partitions.
- Manually selected files or folders.
- Specified file types: archives, emails, documents, databases, financial, virtual machine, audio, video, pictures, security keys.
- Folder types: only user or only OS folders.
- Time spans: when files were accessed, created, modified.
- File size: from 1 byte to infinity.



Multiple search parameters allow you to fine-tune your selection.

3. Once you have adjusted the logical imaging parameters, click **Continue**.
4. Select the Target for your L01 file and click **Create file**.
5. Adjust the settings of the L01 file and click **Create** to start the imaging session.



Settings of an L01 file.

When the imaging is running, you can track the progress of imaging of the individual files. As for the overall imaging progress, the upper graph indicates the amount of imaged data as related to the whole space of the RAID volume, not to the selected volume.

1

ADD CASE

CASES

REPORTS

Atola
TECHNOLOGY

DEVICES

Diagnose

Image

Logical

RAID

Wipe

Other

Devices 5, Block size: 64 KB, Left symmetric

960 GB

RAID 5

Case ID:

54963-5755

1 target

Logical imaging

stereook.mp4

59%

Files scanned: 4248

Files imaged: 468

Speed: 57 MB/s

Bytes imaged: 11.8 GB

1%

Date ↑	Message
2023/07/20 05:34:28	Logical imaging started
2023/07/20 05:34:28	Analyzing: partition 1 at sectors 104,448 - 111,595,495 (type: NTFS)
2023/07/20 05:34:29	Analyzing: partition 3 at sectors 112,640,000 - 1,871,181,823 (type: NTFS, name: Storage)

PAUSE

VIEW FILTERS

Logical imaging of a RAID array.

Imaging report

TaskForce generates the **Imaging completed** report with all the details of the imaged selection from this RAID:

- The number of scanned and imaged files
- The volume of imaged data
- Time stamps

ADD CASE

CASES

REPORTS

Atola

DEVICES

Diagnose

Image

Logical

RAID

Wipe

Other

Devices 5, Block size: 64 KB , Start LBA: 0, Left symmetric

960 GB

Case ID:

54963-5755

Creation date:

07/20/2023 5:38 AM

Software version:

2023.7

TaskForce serial:

83336646

TaskForce IP:

192.168.1.65

RAID type:

RAID 5

Start LBA:

0

Block size:

64 KB (65,536 bytes)

Block order:

Left symmetric

Case Investigator:

Jason Beck

Case description:

Devices in RAID:

WDC WDS240G2G0A-00JH30 204705802235

240 GB

SAS 3

WDC WDS240G2G0A-00JH30 204859801245

240 GB

SAS 5

WDC WDS240G2G0A-00JH30 204859801231

240 GB

SAS 1

WDC WDS240G2G0A-00JH30 204859801251

240 GB

SAS 2

WDC WDS240G2G0A-00JH30 204859804508

240 GB

SAS 4

Logical imaging completed

Targets

Devices 5, Block size 64 KB, Left symmetric_2.L01

14 GB (Samsung SSD 850 PRO 512GB S250NX0H433169Y)/Partition 1

Evidence number: (empty)

Investigator: Jason Beck

Description: (empty)

Stats

Files scanned

4,614

Files imaged

679

Bytes imaged

14 GB (13,860,000,160 bytes)

Log

Date ↑

Message

2023/07/20 05:34:28

Logical imaging started

2023/07/20 05:34:28

Analyzing: partition 1 at sectors 104,448 - 111,595,495 (type: NTFS)

2023/07/20 05:34:29

Analyzing: partition 3 at sectors 112,640,000 - 1,871,181,823 (type: NTFS, name: Storage)

2023/07/20 05:38:30

Logical imaging completed

PRINT

GO TO CASE

|<

<

>

>|

Imaging completed report of a logical imaging session

Imaging started report is created the moment the imaging launched; it is a detailed report of all the settings including the elements selected for this acquisition (partitions, files and folders).

ADD CASE

CASES

REPORTS

Atola

DEVICES

Diagnose

Image

Logical

RAID

Wipe

Other

Devices 5, Block size: 64 KB, Start LBA: 0, Left symmetric

960 GB

RAID 5

Case ID: 54963-5755

Creation date: 07/20/2023 5:38 AM

TaskForce serial: 83336646

RAID type: RAID 5

Block size: 64 KB (65,536 bytes)

Case investigator: Jason Beck

Devices in RAID:

WDC WDS240G2G0A-00JH30 204705802235

WDC WDS240G2G0A-00JH30 204859801245

WDC WDS240G2G0A-00JH30 204859801231

WDC WDS240G2G0A-00JH30 204859801251

WDC WDS240G2G0A-00JH30 204859804508

Software version: 2023.7

TaskForce IP: 192.168.1.65

Start LBA: 0

Block order: Left symmetric

Case description:

240 GB SAS 3

240 GB SAS 5

240 GB SAS 1

240 GB SAS 2

240 GB SAS 4

Logical imaging started

Selected partitions

Partition 1 at sectors 104,448 - 111,595,495 (type: NTFS)

Partition 3 at sectors 112,640,000 - 1,871,181,823 (type: NTFS, name: Storage)

Include filters

Audio, Pictures, Videos

ai, bmp, cdr, clip, cmx, cpt, cr2, dpx, eps, fpx, gbr, gif, heic, heif, icns, ico, indd, jfif, jng, jp2, jpe, jpeg, jpg, max, nef, odg, ogex, pdd, pdn, pef, png, psd, rw2, svg, tec, tif, tiff, wrl, webp, xcf, 16svx, 8svx, aac, ac3, aif, aiff, ape, ast, bwf, caff, cdda, flac, m1a, m2a, m3u, m4a, m4p, m4r, mid, midi, mp1, mp2, mp3, mpa, mpe, ogg, voc, wav, wma, 3g2, 3gp, aaf, asf, avi, avif, bik, brow, fcp, flv, h264, hevc, imovieproj, m1v, m2v, m4v, mkv, moov, mov, movie, mp4, mpe, mpeg, mpg, mpv, mswmm, ppj, prproj, rm, swf, veg, veg-bak, vob, vproj, webm, wfp, wimp, wmf, wmv, wve

Archives, Databases, Documents, Emails, Financial

7z, 7z001, 7zip, apk, ar, arj, bz2, bzip, bzip2, cab, cpio, deb, dmg, gz, gzip, jar, lha, lz, lzh, rar, rpm, tar, tbz2, tgz, z00, z01, zip, zdb, accde, accdt, db, dbf, dbf-wal, db3, dbc, dbf, fdb, fp, fp3, fp5, fp7, frm, gdb, idx, itdb, kexl, mda, mdb, mdf, nsf, ntf, odb, sdf, sql, sqlite, sqlite-wal, sqlite3, v12, wmdb, chm, csv, doc, docx, gdoc, gdraw, htm, html, kwd, numbers, odoc, odp, ods, odt, osheet, ott, pages, pdf, ppt, pptx, rtf, tex, txt, xhtml, xls, xlsx, xml, xps, email, eml, eml, mbox, msg, oft, ost, pst, sc2, vcf, wab, wab, iif, myo, myob, ofx, qbj, qbo, qfx, tax, ytab

Between 1 KB and 50 MB

Modified between 02/01/2002 and 02/01/2020

Exclude filters

OS folders

/Windows, /WinNT, /Program Files, /Program Files (x86), /ProgramData, /Applications, /System, /bin, /boot, /dev, /sbin, /lib, /usr, /proc, /sys, /unix

Documents, Emails

chm, csv, doc, docx, gdoc, gdraw, htm, html, kwd, numbers, odoc, odp, ods, odt, osheet, ott, pages, pdf, ppt, pptx, rtf, tex, txt, xhtml, xls, xlsx, xml, xps, email, eml, eml, mbox, msg, oft, ost, pst, sc2, vcf, wab

Modified between 05/01/2005 and 01/01/2008

Targets

Devices 5, Block size 64 KB, Left symmetric_2.L01

0 B (Samsung SSD 850 PRO 512GB S250NX0H433169Y)/Partition 1

PRINT

GO TO CASE

<

<

>

>

Imaging started report of a logical imaging session.

Imaging selected partitions of a RAID array

With Atola TaskForce, you can image individual partitions of an assembled RAID array to avoid imaging excessive amounts of data.

Once the RAID array has been reassembled, you can preview the contents of its partitions. It helps you conclude which of the partitions may contain the critical evidence and which are irrelevant.

Image selected partitions

To image only selected partitions of an assembled RAID array, do the following:

1. On the **RAID configuration** page, click **Go to image**.

The screenshot shows the Atola Technology RAID configuration interface. The top navigation bar includes 'ADD CASE', 'CASES', 'REPORTS', and 'DEVICES'. The left sidebar contains icons for 'Diagnose', 'Image', 'Logical', 'RAID', 'Wipe', and 'Other'. The main content area is titled 'RAID configuration' and displays the following settings:

- RAID type: RAID 10 (Stripe of ...)
- Start LBA: 0
- Block size: 64 KB (most freque...)
- RAID 1 groups: 2

Below the settings, there are two RAID arrays:

- RAID 1 array #1:
 - 1: ST1000LM048-2E7172 ZDEN417Z, 1 TB, MBR, Unknown, Mirror A
 - 2: ST1000LM048-2E7172 ZDEN40S5, 1 TB, MBR, Unknown, Mirror A
- RAID 1 array #2:
 - 3: ST1000LM048-2E7172 ZKP1BSW9, 1 TB, Mirror B
 - 4: ST1000LM048-2E7172 ZKP1BSYA, 1 TB, Mirror B

At the bottom of the RAID configuration section, there are buttons for 'ADD DEVICES' and 'Drag device here to remove'.

The 'PARTITIONS' section is active, showing a table of partitions:

Partition	Start LBA	Name	Type	Size
1 NTFS2048 NTFS - 4 GB 100%	264,192	BOOKS	Folder	
2 exFAT_32KB exFAT - 5 GB 100%	8,456,192	ENGLISH	Folder	
3 FAT32_8K FAT32 - 10 GB 100%	19,046,048	guitar lessons	Folder	
		MOVIES	Folder	


Below the partitions table, there are buttons for 'GO TO IMAGE' and 'GO TO LOGICAL'. An orange arrow points to the 'GO TO IMAGE' button.

On the right side, the 'Autodetection' section shows a 'RESTART' button and a 'Possible configuration' box with a green checkmark, indicating that the RAID 10 configuration is applied.

At the bottom right, there is a button labeled 'UNMOUNT RAID'.


The suggested RAID configuration applied, partitions preview shows partitions.

2. Select your target device and click **Continue**.
3. On the imaging **Settings** page, click the imaging pass to adjust the imaging range.




The suggested RAID configuration applied, partitions preview shows partitions.

4. In the Edit imaging pass dialog, click on the What to image and select Sectors with data.



Selecting only sectors with data from the What to image list.


5. Unselect the partitions you do not need and click **Apply**. Then click **Save** to keep this change.



Select partitions for imaging

The selective imaging does not linear hashing of the imaged range, [segmented hashing](#) is suggested for such imaging.

6. To launch your imaging session, click **Start**.



The suggested RAID configuration applied, partitions preview shows partitions

Imaging report

TaskForce generates the Imaging completed report with all the details of the RAID:

- The imaged range.
- Time stamps.
- A link to the file with the segmented hash, calculated for the imaged range.

ADD CASE

CASES

REPORTS

Atola

DEVICES

Diagnose

Image

Logical

RAID

Wipe

Other

Devices 4, Block size: 64 KB , Start LBA: 0, RAID 1 groups: [2, 2]

Case ID: 5892-225

2 TB

RAID 10

Creation date: 07/18/2023 11:25 AM

TaskForce serial: 83336646

RAID type: RAID 10

Block size: 64 KB (65,536 bytes)

Case investigator: George Jacobs

Devices in RAID:

ST1000LM048-2E7172 ZDEN417Z

ST1000LM048-2E7172 ZDEN40S5

ST1000LM048-2E7172 ZKP1BSW9

ST1000LM048-2E7172 ZKP1BSYA

Software version: 2023.7

TaskForce IP: 192.168.1.65

Start LBA: 0

RAID 1 groups: [2, 2]

Case description:

1 TB SATA 3

1 TB SATA 6

1 TB SATA 2

1 TB SATA 4

Imaging completed

Targets

Passes: 1 passes with user customized settings

Scheduled sectors: 376,224,906

Imaged sectors: 376,224,906

Errors: 0

Source sector size: 512

Hash method: Segmented

Segment size: 4 GB

Hash type: MD5

Segmented hashes: MD5_HashSegments-raid-5295f31d65f1069e36e310ab493ef08e76f451e46c58cbbede3ca2235f5ffc74.csv (Internal storage)

Signatures: 2,570

Top 10 found signatures

File descriptor	Bytes	Count
JPEG Image File	FF D8 FF	2486
Portable Document Format File	25 50 44 46 2D	32
Microsoft Office File	D0 CF 11 E0 A1 B1 1A E1	8
Compressed Archive File	50 4B 03 04	6
Multimedia File	00 00 00 20 66 74 79 70	4
DVD Movie Information File	44 56 44	4
Qcow Disk Image	51 46 49	4
Archive	78 9C ED	4
Adobe Photoshop Actions Settings File	00 00 00 10 00 00 00	2
Bzip2 Compressed File Archive	42 5A 68	2

Found signatures.csv (Internal storage)

Log

Date ↑	Message
2023/07/18 11:24:42	Imaging started
2023/07/18 11:24:43	Imaging partition 2 at sectors 8,456,192 - 180,946,047 (type: exFAT, name: exFAT_32KB)
2023/07/18 11:25:14	Imaging partition 3 at sectors 180,946,048 - 384,681,098 (type: FAT32, name: FAT32_8K)
2023/07/18 11:25:47	Pass #1 completed
2023/07/18 11:25:47	Imaging completed

PRINT

GO TO CASE

GO TO IMAGE

⏪

⏴

⏵

⏩

The Imaging completed report


How to unmount an assembled RAID

After you selected specific drives or images as RAID participants, they constitute virtual RAID array and become unavailable for individual tasks. For instance, you cannot recalculate hash of one of such drives.

To use the individual drives that are a part of a mounted RAID array, you need to unmount the RAID first.

Follow these steps:


1. At the top right, click **Devices**.



The screenshot shows the Atola Technology web interface. The top navigation bar includes a home icon, 'CASES', 'REPORTS', the Atola Technology logo, and a 'DEVICES' button with a magnifying glass icon. A left sidebar contains icons for 'Diagnose', 'Image', 'Logical', 'RAID', 'Wipe', and 'Other'. The main content area is titled 'Completed' and features a search bar. A list of tasks is displayed, each with a green vertical bar on the left and a timestamp on the right.

Task	Time
Case created	2 hours ago
Source switch of SATA 2 disabled	6 hours ago
Source switch of SATA 2 enabled	a day ago
Diagnostics completed SATA 1 Samsung SSD 860 PRO 256GB S418NF0KB07068A	a day ago
Source switch of SAS 5 enabled	a day ago
Source switch of SAS 3 enabled	a day ago
Source switch of SAS 2 enabled	a day ago
Source switch of SAS 1 enabled	a day ago
Source switch of SATA 1 enabled	a day ago
Imaging completed Win10ProX64.img	4 days ago
Imaging started	4 days ago

2. Scroll down to the bottom of the **Devices** panel and in the **RAID** section select the currently connected and mounted RAID.



Assembled RAID array.

3. At the bottom of the RAID configuration page, click the Unmount RAID button.

The screenshot displays the Atola TaskForce software interface for RAID configuration. The top navigation bar includes 'CASES', 'REPORTS', and 'DEVICES'. The left sidebar contains icons for 'Diagnose', 'Image', 'Logical', 'RAID', 'Wipe', and 'Other'. The main area is titled 'RAID configuration' and shows a RAID 0 (Stripe) configuration with 6 disks. The 'Autodetection' section on the right shows a 'Possible configuration' that has been applied. An orange arrow points to the 'UNMOUNT RAID' button at the bottom right.

Partition	Start LBA	Name	Type	Size
0	No label EXT4 - 719 GB 100%	Books	Folder	
		Graphics	Folder	
		Manuals	Folder	

Unmount RAID with a single click.

Now you can proceed with other sessions, using any of the drives connected to TaskForce.

Supported RAID types and parameters

RAID stands for "*redundant array of independent disks*" or "*redundant array of inexpensive disks*". It is a technology of data storage that combines two or more physical disk drives into one or more logical units (**arrays**). This approach improves data storage reliability, performance, or both.





The data can be distributed across physical drives of a RAID array using different standardized methods, which are called **RAID types** or **RAID levels**.



Atola TaskForce can [reassemble](#) and image the following RAID types:

- [JBOD](#)
- [RAID 0 \(Stripe\)](#)
- [RAID 1 \(Mirror\)](#)
- [RAID 10 \(Stripe of mirrors\)](#)
- [RAID 5 \(Distributed parity\)](#)
- [RAID 6 \(Dual parity\)](#)

This article covers currently supported RAID types and their parameters in more detail.

RAID types comparison

RAID type and method	Scheme of data distribution	Description	Common usage	Performance	Fault tolerance	Capacity	Minimum number of disks
JBOD: Spanning		Multiple physical drives are spanned, or concatenated, together to form a single logical disk with no redundancy.	Increasing capacity. Cost-effective storage with no performance or security benefits.	No benefits. Array performance equals the performance of each drive depending on data location.	Very low. Redundancy is not provided. If one drive fails, a failure is typically isolated to this drive.	Array capacity is the sum of the capacities of its members.	2
RAID 0: Striping		Data is split ('striped') evenly across the drives forming a logical array, without parity or mirroring.	Improving reading and writing speed.	The highest read and write speed because data is striped across all disks, and the controller can read and write several drives in parallel with no delay due to parity calculation.	Very low. If one drive fails, data is lost.	Volume is the sum of the capacities of the drives in the set. No overhead is provided by parity blocks.	2
RAID 1: Mirroring		Mirroring without parity or striping. Data is written identically to two or more drives, producing a "mirrored set" of drives.	Standart app servers; applications that require higher fault tolerance, than storage capacity or performance.	High read speed: the controller can read several drives in parallel. Lower compared to RAID 0 write speed: all data must be written to two disks.	High. If one drive fails, data is restored from the "mirror" drive(s).	Storage capacity is reduced by 1/2 as all data is written twice.	2
RAID 10 (1+0): Striping and mirroring		Striped set from a series of mirrored drives. Data is shared between disks and duplicated.	Highly utilized database, email, web servers.	High read speed similar to RAID 1 due to improved read rates through simultaneous disk reads. Moderate to high write speed because data is mirrored across pairs, but the striping allows for higher write speeds compared to just RAID 1.	Very high fault tolerance. The array can sustain multiple drive losses so long as no mirror loses all its drives.	Storage capacity is reduced by 1/2 as all data is written twice.	4

RAID type and method	Scheme of data distribution	Description	Common usage	Performance	Fault tolerance	Capacity	Minimum number of disks
RAID 5: Striping and parity		Data is split ('striped') evenly across the drives with added distributed parity.	Normal file storage and app servers.	High read speed but slightly lower than RAID 0 because of the overhead from parity calculations and distribution across all disks. Moderate write speed due to the need for parity calculations and writing both data and parity information.	High. If one drive fails, data can still be calculated from the distributed parity.	Due to parity blocks, storage capacity equals the smallest drive capacity multiplied by (N-1) drives.	3
RAID 6: Striping and double parity		Data is split ('striped') evenly across the drives with two distributed parity blocks, instead of one.	Large file storage and app servers.	High read speed similar to RAID 5, but can be slightly slower due to additional parity calculations. Write speed is lower than RAID 5 because it requires two parity blocks to be written.	Very high fault tolerance due to double parity. Data can still be calculated even if two drives fail.	Due to dual parity blocks, storage capacity equals the smallest drive capacity multiplied by (N-2) drives.	4

RAID parameters

Each individual RAID is defined by a combination of several main parameters:

- [RAID type \(or RAID level\)](#)
- [Number of devices](#)
- [Device order](#)
- [Start LBA \(or Disk offset\)](#)
- [Block size](#)
- [Block order](#)
- [Parity block order](#)
- [File system](#)

To identify RAID type and exact parameters, TaskForce [automatically checks](#) up to 200,000,000 RAID parameter combinations. However, if you know the exact RAID configuration, you can enter it manually.

RAID configuration

RAID type	Start LBA	Block size	Block order
RAID 5	0	512 KB (Adaptec, Li...	Left asymmetric

1	GIGABYTE GP-GSTFS31120GNTD S	120 GB	=
2	JAJS300M120C 30013348333	120 GB	=
3	SATA SSD PH191112003195	120 GB ext4	=
4	SPCC Solid State Disk 2018130300	120 GB	=
5	Patriot Burst 5680076A1EEE00402	120 GB MBR Unknown	=

ADD DEVICES ADD MISSING DEVICE

Drag device here to remove

Depending on the RAID type, different sets of parameters can be applied:

Number of devices	✓	✓	✓	✓	+ missing	+ missing
Device order	✓	✓	✓	✓	✓	✓
Start LBA	✓	✓	✓	✓	✓	✓
Block size	-	✓	-	✓	✓	✓
RAID 1 groups	-	-	-	✓	-	-
Block order	-	-	-	-	✓	✓
Parity block order/type	-	-	-	-	-	✓

Here's what each of the RAID parameters means.

RAID type (or RAID level)

RAID type is defined by the method in which data is distributed across physical drives combining the array. TaskForce currently supports the following RAID types:

- [JBOD](#)
- [RAID 0 \(Stripe\)](#)
- [RAID 1 \(Mirror\)](#)
- [RAID 10 \(Stripe of mirrors\)](#)
- [RAID 5 \(Distributed parity\)](#)

- [RAID 6 \(Dual parity\)](#)

In TaskForce, the RAID type can be automatically identified by the **Autodetection** module along with other parameters.

However, if you know the RAID type, you can select it manually from the **RAID type** list in the **RAID configuration** section.

RAID configuration

RAID type: **RAID 5** Start LBA: **0** Block size: **64 KB (most frequ...** Block order: **Left asymmetric**

1	Patriot Burst 5680076A1EEE00402C	120 GB	MBR	Unknown	=
2	JAJS300M120C 30013348333	120 GB			=
3	GIGABYTE GP-GSTFS31120GNTD S	120 GB			=
4	SATA SSD PH191112003195	120 GB	ext4		=
5	SPCC Solid State Disk 2018130300C	120 GB			=

ADD DEVICES ADD MISSING DEVICE Drag device here to remove

PARTITIONS HEX VIEWER

No partitions found

Supported file systems: NTFS, ext 4/3/2

Autodetection PAUSE

Background analysis of RAID devices/images. It shows hints to help you find RAID configuration faster.

Stage 2. Detecting RAID configuration 79%
Checked variants: 4,537 out of 5,760

Possible configuration
RAID 5, block size: 512 KB, order: Left asymmetric **APPLY**

GO TO IMAGE UNMOUNT RAID

The RAID type list in the RAID configuration section.


Number of devices

Applies to: all supported RAID types.

The minimum number of disks needed to build a RAID array is:

- two for JBOD, RAID 0, and RAID 1,
- three for RAID 5,
- four for RAID 10 and RAID 6.

The actual number of RAID members is determined by you when you select in the TaskForce RAID module physical disks (SATA/SAS/USB drives) or their images (raw, E01 or AFF4 image files) as Sources that make up a RAID array.




Selecting source RAID drives.

Missing RAID members. For the RAID types that use redundancy (parity blocks), TaskForce can reassemble an array and create its image [even with a missing device](#):

- with **one missing device** for RAID 5 that uses distributed parity,
- with **two missing devices** for RAID 6 that uses double distributed parity.

See [RAID 5 with a missing drive](#) for more details.



Adding missing RAID 5 drive.

Device order

Applies to: all supported RAID types.

The drive order is the order in which physical devices (or their images) follow each other within a logical array (think pages within a book). That order is defined by a RAID controller during the setup of an array. To reassemble and image a RAID array, you need to determine which of its members go first, second, third and so on.

In TaskForce, the drive order can be automatically detected by the Autodetection module along with other parameters.

However, if you know the drive order, arrange the RAID members manually by dragging them up or down in the **RAID configuration** section of the TaskForce RAID module.

RAID configuration

RAID type

RAID 5

Start LBA

0

Block size

512 KB (Adaptec, Li...

Block order

Left asymmetric

1	GIGABYTE GP-GSTFS31120GNTD S	120 GB	=
2	JAJS300M120C 30013348333	120 GB	=
3	SATA SSD PH191112003195	120 GB ext4	=
4	SPCC Solid State Disk 2018130300	120 GB	=
5	Patriot Burst 5680076A1EEE00402	120 GB MBR Unknown	=

ADD DEVICES ADD MISSING DEVICE

Drag device here to remove

Start LBA (or Disk offset)

Applies to: all supported RAID types.

LBA stands for Logical Block Address. It is a scheme used for addressing a particular area on a storage device in a linear way using sequential integer numbers, starting from zero: 0, 1, 2, 3... Thanks to LBA a single drive or a RAID controller identifies the exact location of bytes it needs to read from or write to.

When we talk about RAID's, **Start LBA**, also known as **Device offset**, refers to an address of a logical block on each physical RAID member where user data begins. **Start LBA** of each member of the same RAID is identical.

If RAID's Start LBA is 0, RAID blocks (chunks) start from the first logical block on each drive in an array, without any offset (it is the most common case).

If RAID's Start LBA is greater than 0, a space created by this offset is skipped. RAID blocks start from the specified Start LBA on each drive in an array. Typically, the skipped space can be used to store RAID metadata (data about data).

In some cases, TaskForce can identify RAID's Start LBA during autodetection using a heuristic algorithm.

If you know Start LBA, you can enter it in the **Start LBA** field of the **RAID configuration** section.

RAID configuration

RAID type

RAID 5

Start LBA

0

Block size

512 KB (Adaptec, Li...

Block order

Left asymmetric

1	GIGABYTE GP-GSTFS31120GNTD S	120 GB	=
2	JAJS300M120C 30013348333	120 GB	=
3	SATA SSD PH191112003195	120 GB ext4	=
4	SPCC Solid State Disk 2018130300	120 GB	=
5	Patriot Burst 5680076A1EEE004023	120 GB MBR Unknown	=

ADD DEVICES ADD MISSING DEVICE

Drag device here to remove

When we talk about logical partitions, Start LBA refers to the logical block address where a particular partition starts in an already assembled RAID array. In this case, Start LBA of each partition is different.

PARTITIONS

HEX VIEWER

Partition	Start LBA
0 1 ext ext4 - 270 GB	2,048
1 2 ext ext4 - 210 GB	526,725,120

Partition 0 / home / atola

Name	Type	Size
Downloads	Folder	
Sentinel_LDK_RedHat_and...	Folder	
sg3_utils-1.44	Folder	

Block size

Applies to: RAID 0, RAID 10, RAID 5, RAID 6 (types that use striping method).

Some RAID types use striping method for distributing data across physical drives: data is split into consecutive logical blocks (also called "stripes" or "chunks"), which are stored on different physical storage devices.

The **Block size** parameter in TaskForce RAID module refers to the size of these logical blocks (or "stripes") measured in bytes. TaskForce supports RAID block sizes from 512 bytes to 1 megabyte and can detect that parameter automatically using a heuristic algorithm. The most frequently used block size is 64 kilobytes.

If you know the block size of a RAID, you can select it manually from the **Block size** list in the **RAID configuration** section.

RAID configuration

RAID type

RAID 5

Start LBA

0

Block size

512 KB (Adaptec, Li...

Block order

Left asymmetric

1	GIGABYTE GP-GSTFS31120GNTD S	120 GB	=
2	JAJS300M120C 30013348333	120 GB	=
3	SATA SSD PH191112003195	120 GB ext4	=
4	SPCC Solid State Disk 20181303001	120 GB	=
5	Patriot Burst 5680076A1EEE004023	120 GB MBR Unknown	=

ADD DEVICES ADD MISSING DEVICE

Drag device here to remove 

Block order

Applies to: RAID 5, RAID 6 (types that use block-level striping with distributed parity).

The **Block order** parameter defines the layout (or pattern), in which RAID logical blocks ("stripes") are distributed among the devices in the array. It is the sequence of writing data and [parity blocks](#) across the RAID members and depends on:

1. The direction of data blocks writing: left to right or right to left on the disk array.
2. The placement of the parity blocks: at the beginning or the end of the stripe.
3. The location of the first block of a stripe relative to the parity of the previous stripe.

For RAID 5, the Block order parameter can be:

- **Left symmetric.** This is the default RAID 5 layout under Linux.
- **Left asymmetric.** This is the 'standard' RAID 5 layout.
- **Right symmetric.**
- **Right asymmetric.**

Left symmetric

1	2	3	P
5	6	P	4
9	P	7	8
P	10	11	12
Device 1	Device 2	Device 3	Device 4

Right symmetric

P	1	2	3
6	P	4	5
8	9	P	7
10	11	12	P
Device 1	Device 2	Device 3	Device 4

Left asymmetric

1	2	3	P
4	5	P	6
7	P	8	9
P	10	11	12
Device 1	Device 2	Device 3	Device 4

Right asymmetric

P	1	2	3
4	P	5	6
7	8	P	9
10	11	12	P
Device 1	Device 2	Device 3	Device 4

For RAID 6, the Block order parameter can be following:


Left symmetric



Left symmetric (shift 1)



Right symmetric



Left asymmetric



Left asymmetric (shift 1)



Right asymmetric



Parity block order

Applies to: RAID 6.


Unlike RAID 5, RAID 6 uses not one, but two types of parity blocks. The **Parity block order** parameter defines, which parity block type comes first and which follows:

- **PQ:** XOR parity block comes first, Reed-Solomon parity block comes second.
- **QP:** Reed-Solomon parity block comes first, XOR parity block comes second.

File system

After reassembling RAID using other parameters, TaskForce automatically identifies its file system and shows a preview of its partitions, folders, and files.

Currently supported file systems are: .



The preview of the RAID file system with partitions, folders, and files.

JBOD

JBOD is an abbreviation of "*Just a Bunch Of Drives*" and refers to a method for **concatenation** (or **spanning**) of multiple physical drives into a single logical disk with no redundancy. Data is written consecutively, from the beginning to the end of the first drive, and then spans to the second, the third and so on.

JBOD provides cost-effective storage with no performance or security benefits. It needs at least 2 drives to function. Array capacity is the sum of the capacities of its members.

If one drive fails, all data stored across the entire array may be compromised or lost, as JBOD does not implement any form of redundancy or fault tolerance.

JBOD advantages:

- Increasing capacity.
- Cost-effective storage.

JBOD disadvantages:

- No performance benefits.
- Redundancy is not provided.

JBOD parameters

When reassembling JBOD in Atola TaskForce, the following parameters are available for manual selection:

- [Device order](#)
- [Start LBA](#)

JBOD



RAID 0 (Stripe)

In RAID 0, data is evenly distributed into stripes across two or more drives, forming a logical array, without parity or mirroring. RAID 0 provides higher reading and writing speed, as it can read and write multiple drives in parallel. As RAID 0 does not use mirroring or parity blocks, there's no redundancy and fault tolerance: if one drive fails, data is lost.

RAID 0 needs at least 2 drives to function. Array volume is the sum of the capacities of the drives in the set, there's no overhead provided by parity blocks.

RAID 0 advantages:

- Highest reading and writing speed.
- No overhead provided by parity blocks.
- Maximum capacity: sum of all RAID members.

RAID 0 disadvantages:

- No fault tolerance: if one drive fails, data is lost.

RAID 0 parameters

When reassembling RAID 0 in Atola TaskForce, the following parameters are available for manual selection:

RAID 0 (Stripe)



- [Device order](#)
- [Start LBA](#)
- [Block size](#)

See also

- [Imaging RAID 0 with unknown configuration](#)

RAID 1 (Mirror)


RAID 1 creates an exact copy of the drive's data ("mirror") on another drive or drives. Thus RAID 1 usually consists of mirrored drives which are exact copies of each other. No spanning, striping, or parity is being used. The capacity of the array is limited to the size of the smallest member drive.

If one member drive fails, data is restored by copying from the mirror drive. Because of that RAID 1 can still work even if only one member drive is in good condition.

RAID 1 reads multiple drives in parallel, increasing overall reading performance. But writing speed is lower, as it takes additional time to duplicate data on the mirror drive.

When TaskForce detects a mirrored pair of devices (drives or image files), it labels both of them with the "Mirror" tag.

RAID 1 (Mirror)



RAID configuration

RAID type

RAID 1 (Mirror)

Start LBA

0

1	Samsung SSD 860 EV...3YKNS0N625541K.E0	51 GB	Mirror A	exFAT	=
2	Samsung SSD 860 PR...42VNGAK204945N.E0	51 GB	Mirror A	exFAT	=

RAID 1 advantages:

- Improved reading speed.
- High fault tolerance. Data can be restored from the “mirror” drive.
- Simplest RAID storage system.

RAID 1 disadvantages:


- Storage capacity is limited to the size of the smallest member drive.
- Writing speed is lower due to data duplication.

RAID 1 parameters

When reassembling RAID 1 in Atola TaskForce, the following parameters are available for manual selection:

- [Device order](#)
- [Start LBA](#)

RAID 10 (Stripe of mirrors)



RAID 10 (Stripe of mirrors)

RAID 10 is called “stripe of mirrors” as it combines [RAID 1](#) (“mirroring”) and [RAID 0](#) (“striping”) methods within a single logical array. Data is shared between drives and duplicated. It is a RAID 0 array consisting of mirrors and thus requires a minimum of four drives: two for striping and two for storing mirrored data of each drive.

High read and write performance is achieved by striping RAID 1 mirrored segments.

If one member drive fails, data is restored by copying from the mirror drive. The overall storage capacity is reduced by 1/2 as all data is written twice.

RAID 10 typical use is highly loaded databases, email, or web servers.

When TaskForce detects a mirrored pair of devices (drives or image files), it labels both of them with the “Mirror” tag.

RAID configuration

RAID type	Start LBA	Block size	RAID 1 groups
RAID 10 (Stripe of mirrors) ▼	0	64 KB (most frequent) ▼	2 ▼

RAID 1 array #1

1	1.img	240 GB	Mirror A	NTFS	=
2	2.img	240 GB	Mirror A	NTFS	=

RAID 1 array #2

3	3.img	240 GB	Mirror B		=
4	4.img	240 GB	Mirror B		=

RAID 10 advantages:

- High reading and writing speed.
- High fault tolerance. Data can be restored from the “mirror” drive.

RAID 10 disadvantages:

- Storage capacity is reduced by 1/2 as all data is written twice.
- Limited scalability.

RAID 10 parameters

When reassembling RAID 10 in Atola TaskForce, the following parameters are available for manual selection:

- [Device order](#)
- [Start LBA](#)
- [Block size](#)
- RAID 1 groups

See also

- [Reassembly and imaging of RAID 10](#)

RAID 5 (Distributed parity)

In RAID 5, data in the form of block-level stripes is evenly distributed across at least 3 drives, along with parity information which is used to restore data.

RAID 5 can still operate without one of its members. If one drive fails, its data is calculated from the parity blocks distributed across other members. Atola TaskForce can detect the parameters of the [RAID 5 array with two damaged drives](#) and successfully image such RAID5 regardless of errors.


The storage capacity of RAID 5 is reduced due to parity blocks.

Reading and writing speed is high since RAID 5 can read from and write to all array members in parallel.

RAID 5 is typically used for file servers, database servers, and application servers.

RAID 5 can have different layouts, or block orders, depending on the pattern in which RAID 5 data blocks (“stripes”) and parity blocks are distributed among the devices in the array. The layout is defined by:

RAID 5




1. The direction of data blocks writing: left to right or right to left on the disk array.
2. The placement of the parity blocks: at the beginning or end of a stripe.
3. The location of the first block of a stripe relative to the parity of the previous stripe.


There are four RAID 5 layouts (or block orders):

1. Left symmetric.
2. Right symmetric.
3. Left asymmetric.
4. Right asymmetric.

Left symmetric



Right symmetric



Left asymmetric



Right asymmetric



RAID 5 advantages:

- High read speed but slightly lower than RAID 0 because of the overhead from parity calculations and distribution across all disks.
- High fault tolerance. Data from one failed drive can be restored using parity information, stored on other array members.

RAID 5 disadvantages:

- Storage capacity is reduced due to parity blocks. Space efficiency is described by formula: $1 - 1/n$, where n is the number of devices in an array. For instance, for RAID 5 consisting of 5 drives with total capacity of 1 TB, available space will be reduced by 1/5 (or 200 MB), to 800 MB.
- Moderate write speed due to the need for parity calculations and writing both data and parity information.

RAID 5 parameters

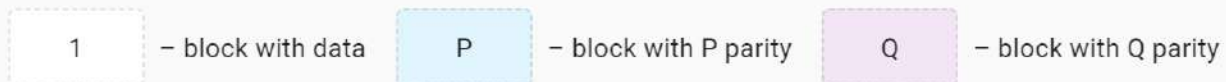
When reassembling RAID 5 in Atola TaskForce, the following parameters are available for manual selection:

- [Device order](#)
- [Start LBA](#)
- [Block size](#)
- [Block order](#)

See also

- [RAID 5 with unknown configuration](#)
- [RAID 5 with a missing drive](#)
- [Imaging RAID 5 array with 2 damaged drives](#)

RAID 6



RAID 6 (Dual parity)

RAID 6 uses block-level striping (data is shared between drives) with two parity blocks, instead of one, distributed across all member disks. This gives extra redundancy to an array: RAID 6 can read and write data even if two drives fail at the same time.

This RAID type needs at least 4 drives to function. Storage capacity is reduced because of the dual parity scheme.

RAID 6 is typically used for large file storage, file servers, database servers, app servers.

Thanks to RAID 6 extra redundancy coming from two parity block types, TaskForce can rebuild a RAID 6 array even if two of its members are missing or damaged.

RAID 6 can have different layouts, or block orders, depending on the pattern in which RAID 6 data blocks ("stripes") and parity blocks are distributed among the devices in the array. The layout is defined by:

1. The direction of data blocks writing: left to right or right to left on the disk array.
2. The placement of the parity blocks: at the beginning or end of the stripe.
3. Parity block order.
4. The location of the first block of a stripe relative to the parity of the previous stripe.

RAID 6 advantages:

- Very high data fault tolerance.
- Data can be restored even if two drives fail.

- High reading speed similar to RAID 5, but can be slightly slower due to additional parity calculations.

RAID 6 disadvantages:

- Storage capacity is reduced because of the dual parity scheme.
- Reduced writing speed due to the usage of two types of parity blocks.
- Minimum 4 drives needed.

RAID 6 parameters

When reassembling RAID 6 in Atola TaskForce, the following parameters are available for manual selection:

- [Device order](#)
- [Start LBA](#)
- [Block size](#)
- [Block order](#)
- [Parity block order](#)

RAID read and write speeds

The different RAID levels offer various trade-offs in terms of read and write speed, redundancy, and storage efficiency. Below is a comparison of RAID 0, 1, 10, 5, and 6 regarding their read and write speeds.

RAID type	Read Speed	Write Speed
RAID 0	Highest because data is striped across all disks, allowing simultaneous reads.	Highest because data is striped, allowing simultaneous writes. No parity calculation delay.
RAID 1	High because data can be read from either of the mirrored disks, potentially doubling the read rate compared to a single disk.	Lower than RAID 0 because all data must be written to two disks, causing a slight overhead.
RAID 10	High , similar to RAID 1 because it combines striping and mirroring, offering improved read rates through simultaneous disk reads.	Moderate to High because data is mirrored across pairs, but the striping allows for higher write speeds compared to just RAID 1.
RAID 5	High but slightly lower than RAID 0 because of the overhead from parity calculations and distribution across all disks.	Moderate due to the need for parity calculations and writing both data and parity information, which introduces some overhead.
RAID 6	High , similar to RAID 5, but can be slightly slower due to additional parity calculations.	Lower than RAID 5 because it requires two parity blocks to be written, further increasing the overhead compared to RAID 5.

Key Points:

- **RAID 0** offers the best performance both in reads and writes but no redundancy.
- **RAID 1** provides good read performance and redundancy at the cost of available capacity.
- **RAID 10** combines the benefits of RAID 0 and RAID 1, offering a good balance of speed and redundancy.
- **RAID 5 and RAID 6** offer a good balance between storage efficiency, read speed, and redundancy, with RAID 6 providing higher fault tolerance at the cost of write speed due to additional parity calculations.

RAID tags and what they mean

When reassembling a RAID in TaskForce, start by selecting the drives and/or images it consists of. The system reads the first 3 million sectors of each RAID member, analyzes the data and compares the members against each other. TaskForce then labels the RAID members with appropriate tags to help you identify the condition, relation and possible position.

1	disk-1.vhd.raw	512 GB	MBR	Mirror A	=
2	disk-3.vhd.raw	512 GB			=
3	disk-4.vhd.raw	512 GB	MBR	Mirror A	=
4	disk-2.vhd.raw	512 GB	Spare		=

Spare

The **Spare** tag informs you that 99.95% of the drive's initial 3M sectors are filled with zeros.

A spare (aka hot spare) drive in RAID arrays is used as a standby drive reserved to replace a RAID member in case it fails: in this situation the RAID controller uses redundancy data to reconstruct the data from the failed disk to the spare one. Spare drives are used in RAID 1, RAID 5 and RAID 6.

Mirror

With the **Mirror** tag, TaskForce informs you that 99.5% out of the analyzed 3M sectors are identical to the same sectors on a different RAID member. It works as a hint that you are dealing with a RAID 1, RAID 10 or RAID 50.

MBR

MBR (Master Boot Record) points to the drive(s) that contain an MBR. Therefore, it is likely that the drive with detected MBR can be placed first in the given array.

File system tags

The initial 3M sectors of each drive/image are analyzed to identify a boot sector of any known file system. When a partition boot sector is detected, the corresponding file system tag is added. To see LBA offset and sector count of the partition, hover mouse cursor over the tag.

Supported file system tags: NTFS, ext4/3/2, APFS, HFS, HFS+, exFAT, FAT32, FAT16, XFS.

Additionally, **Unknown** tag may appear if the file system identification is not supported yet.

1	2.E01	120 GB	MBR		=
2	1.E01	120 GB			=
3	3.E01	120 GB	ext4		=
4	4.E01				=
5	5.E01	120 GB			=

LBA offset: 0
Sector count: 526,723,072 (270 GB)

Error

The **Error** tag informs you whether there are read errors encountered in the process of RAID autodetection. To see the exact number of encountered errors, hover mouse over the **Error** tag to get the tooltip.

1	HGST HTS545050A7E680 RBE50AM52D2U0P	500 GB	
2	Seagate-E-errors 4124512312	250 GB	Error

ADD DEVICES

Read error count: 4

RAID cheat sheet

When RAID Autodetection finds nothing

Follow these tips if the automatic RAID reassembly module doesn't detect a RAID configuration.

If you know the RAID configuration

If you know the exact RAID configuration, skip waiting for Autodetection to complete and manually select the RAID parameters instead.

If you only know the device order

If you know the exact order of the RAID members, select them in the correct sequence: the RAID Autodetection module prioritizes the device order specified by the user.

If you know the RAID offset (Start LBA)

If you're aware of a specific disk offset for the RAID members, specify the LBA in the Autodetection module's settings (click the gear icon). This will enable the Autodetection to identify a suitable RAID configuration in cases where it previously failed.

Synology NAS RAID detection

If you are certain the devices belong to a Synology NAS RAID but the RAID configuration isn't detected right away and Autodetection is still running, consider the following possibilities:

- One or more selected devices may not be part of this RAID array.
- One or more devices could be missing from the array.

To resolve this, stop the Autodetection process and try selecting a different combination of drives that belong to the RAID array.

Multiple configurations found

If the Autodetection module suggests multiple possible RAID configurations, the best match will always appear first.

Additionally, pay attention to the small tags displaying the accuracy percentage next to the RAID partitions — the higher the percentage, the better the match.

Parity order in RAID 6

For RAID 6, the 'Parity Order' parameter is only relevant when a device is missing.

Handling RAID 5 / RAID 6 with missing devices

If you know that a device is missing from your RAID 5 or RAID 6, click 'Add missing device' to speed up the Autodetection process.

For RAID 6, you can specify up to two missing devices.

Exclude 'Spare' devices

RAID members marked with the 'Spare' tag most likely contain only zeroes. It's recommended to remove these devices from the list by dragging them to the bin icon.

RAID configuration not detected

If the Autodetection module is unable to find a valid RAID configuration, try adjusting the parameters manually.

To help you evaluate the validity of the applied configuration, TaskForce displays a small evaluate the validity of next to each partition.

If a file system is not currently supported, an Unknown tag may appear next to the corresponding RAID member.

Tricky cases

RAID member s tagged with 'Error'

If a RAID member is marked with an **Error** tag:

- Image the drive first, then attempt to reassemble the RAID using the acquired image.
- For **RAID 5** or **RAID 6**, remove the device with 'Error' from the list, then click 'Add missing device'. TaskForce will use redundancy to rebuild the RAID.
- For **RAID 10**, remove the faulty from the list. TaskForce will automatically use the mirrored copy to rebuild the array.

Imaging RAID takes too long

RAID capacity can exceed dozens of terabytes, making full physical imaging a time-consuming process.

To speed up access to critical data, start with logical imaging using smart filters to extract only the files of interest. You can perform a full physical imaging later by clicking the 'Go to image' button.

Not enough drive ports for all RAID members

If you're dealing with a large array and don't have enough drive ports for all RAID members:

1. image some of the drives to image files.
2. Go to RAID > Select source device > Files > Select files and use these image files instead of physical drives.

Using Web API in a browser

Web API is built into TaskForce, and it helps optimize your workflow in many ways.

Web API is extremely handy as it allows you to use it in scripts, via CLI tools like curl, and simply by typing commands in the browser address bar.

Here's how to use Web API in a browser:

1. Scan devices plugged to all source ports. The command powers up all ports and returns the list of drive on each port in Source mode as well as the model and the serial number of the drive on each port.



A screenshot of a web browser window. The address bar shows the URL `10.0.0.93/api/scan-devices`. The page content displays a JSON response from the API. The JSON object contains three fields: `"scannedSourcePorts": 5`, `"totalSourcePorts": 9`, and `"foundSourceDevices": [ ... ]`. The array contains five strings, each representing a drive with its interface, brand, model, capacity, and serial number.

```
{
  "scannedSourcePorts": 5,
  "totalSourcePorts": 9,
  "foundSourceDevices": [
    "SAS3 WDC WDS120G2G0A-00JH30 180287804729",
    "SAS1 Samsung SSD 850 PRO 128GB S24ZNX0HC02407M",
    "SATA3 WDC WDS120G1G0A-00SS50 171710A0035D",
    "SATA5 Samsung SSD 860 PRO 256GB S418NF0KB07067E",
    "SATA1 Samsung SSD 860 PRO 256GB S418NF0KB07041E"
  ]
}
```

Devices identification.

2. Start imaging a source drive plugged into TaskForce SATA 4 port.




A screenshot of a web browser window. The address bar shows the URL `10.0.0.65/api/start-image?source=SATA4&targetFolder=//Vitaliy/Share`. The page content displays a single string response from the API, which is an image key.

```
image_9_SanDisk_SDSSDA120G_171108456213_Z32080RL
```

Start imaging.

3. Track imaging session status using task key received in response to the command above.



Check task status.

For more information about these and other commands, see [API specification](#) that we made available to public.

Instantly starting 12 imaging sessions using Web API

Imagine you have 12 TaskForce ports switched to Source mode and source drives plugged into them. Now you can instantly launch 12 imaging sessions simply starting the script.

Python script utilizes `/start-image` API request and prints task keys of all launched imaging sessions.

```
import sys

if sys.version_info[0] < 3:
    raise Exception("Please use Python 3 to run this script")

import urllib.request

ports = ["SATA1", "SATA2", "SATA3", "SATA4", "SATA5", "SATA6", "SAS1", "SAS2", "SAS3", "SAS4", "SAS5", "SAS6"]
tasks = []
errors = {}

for port in ports:
    try:
        res = urllib.request.urlopen("http://10.0.0.4/api/start-image?source=%s&targetFolder=//Vitaliy/Share" % (port))
        tasks.append(res.read().decode('utf-8'))
    except urllib.error.HTTPError as e:
        errors[port] = e.read()

print("IDs of started imaging tasks:")
print('\n'.join(tasks))
```

The script works in any operating system. To run, perform the following actions:

1. Save the script into `image12.py` file.
2. Replace `10.0.0.4` with IP address of your TaskForce.
3. Replace `//Vitaliy/Share` with your shared network folder path.
4. Execute the script in the console: `python image12.py`.

For more information about these and other commands, please look up the [API documentation](#) that we made available to public.

Autostart image analysis when imaging is completed

With TaskForce, you can track the status of the started imaging sessions using `/check-task` API request. It reports the imaging progress enabling you (or your code) to notice when the task gets completed. Once this notification is received, it makes perfect sense to automatically start the forensic analysis of the target image.

Powershell script below shows how one can create this kind of automation flow:

1. Start imaging a source drive on TaskForce SATA port 4 to the target folder `\\Vitaliy\Share`.
2. Wait for imaging completion using `/check-task`.
3. Launch Autopsy Ingest via command-line when the target image is ready.

Important: Instead of Autopsy, you are free to use any Magnet Forensics products, X-Ways Forensics, or any other forensic analysis toolkit that supports console launch with arguments.

```
try {
    $r = Invoke-WebRequest "http://10.0.0.65/api/start-image?source=SATA4&targetFolder=\\Vitaliy\Share"
}
catch {
    Write-Output "$($_.Exception.Message)"
    exit $_.Exception.Response.StatusCode
}

$taskKey = $r.Content
do {
    $check = (Invoke-WebRequest "http://10.0.0.65/api/check-task?taskKey=$taskKey").Content | ConvertFrom-Json
    Start-Sleep -s 1
} while ($check.state -eq "progress")

$windowsPath = "C:\Share\" + ($check.target -replace '[\|]', '\|' | Split-Path -leaf)
$caseName = "Case123"
$autopsyArguments = '" --createCase --caseName="' + $caseName + ' --caseBaseDir="C:\Work\Cases"'
                  + ' --addDataSource --dataSourcePath="' + $windowsPath + '" --runIngest --generateReports'

Start-Process -FilePath "C:\Program Files\Autopsy\bin\autopsy64.exe" -ArgumentList $autopsyArguments
```

The script works in Windows with Powershell. To run it, please perform the following actions:

1. Install [Autopsy](#).
2. Create C:\Share folder.
3. Save the script into image.ps1 file.
4. Replace 10.0.0.65 with IP address of your TaskForce.
5. Replace `\\Vitaliy\Share` with your shared network folder path.
6. Execute the script in the console: `powershell -ExecutionPolicy Bypass -File image.ps1 .`

Autopsy Ingest v4.11 does not work with network file paths from the command line. That's why this example shows a shared folder located at PC where PowerShell script is executed. Therefore `\\Vitaliy\Share` points to C:\Share folder.

For more information about these and other commands, see [API specification](#) that we made available to public.


Multi-launch of single-device operations

To wipe a bunch of target drives for subsequent imaging sessions or verify hash values on multiple drives in your archive, use multi-launch functionality in TaskForce.

The function is currently supported for single-drive tasks: Wiping, Diagnostics, Hashing.


To wipe multiple drives:

1. On the left, click **Wipe**.
2. On the **Select target devices** panel, enable **Multi-launch** and select the drives.




The Multi-launch option.

3. On the wiping settings page, adjust the sessions parameters. They will be applied to all the currently selected drives. If you wish to double-check the list of selected devices, click the top panel with the number of selected drives. The drop-down provides info about the drives, their health status and case ID.



The Settings page.

4. Click **Start**.



Multiple sessions initiated via multi-launch.


When the **Check if device contains data** option is enabled, TaskForce scans all selected devices. If any data may be overwritten, the imager will warn you.

The procedure is the same for other single-drive processes. When hashing, multi-launch can be applied to both the devices plugged into the system and the locally stored image files.

The [case management system](#) automatically saves separate reports into the individual cases.

Wiping 18 drives simultaneously

While cumulative imaging speed in TaskForce constitutes 15 TB/h, the engine can wipe up to 18 drives connected to it and achieve a cumulative speed of 15TB/h, 20TB/h or even more.



TaskForce's task-oriented and efficient user interface is designed to enable the launch of every operation in just a couple of clicks to expedite work with multiple evidence drives.

TaskForce has 18 ports:

- 6 SATA,
- 6 SATA/SAS,
- 4 USB,
- 1 IDE,
- 1 Extension slot for Atola Thunderbolt, Apple PCIe SSD and M.2 NVMe/PCIe/SATA SSD extension modules,

all of which can be used for simultaneous wiping sessions.

TaskForce can wipe 18 devices simultaneously at their top native speeds when using the standard wiping method.

Launch multiple wiping sessions

To perform multiple wiping sessions:

1. Connect the drives to TaskForce.
2. Switch the ports, to which the drives are connected, to Target mode by using the individual Source switches on each port.



3. In the TaskForce user interface, click the **Wipe** icon on the left.
4. On the **Select target device** panel, select a drive.
To launch the wiping process for multiple drives simultaneously, select **Multi-launch** at the top right, and select the drives you want to wipe.

Wiping

Start LBA: 0 End LBA: 468 862 127

Wiping method:

- ☒ Fill sectors with the following pattern
- ☐ Put LBA number at the beginning of every sector
- ☐ NIST 800-88
- ☐ DoD 5220.22-M
- ☐ Random
- ☐ Secure Erase

Pattern: 00 ☒ HEX

Method description

Linear overwrite of the whole disk space from LBA 0 to the LBA reported by IDENTIFY DEVICE command. Data is overwritten with the specified pattern. Maximum pattern length is 512 bytes. If fewer than 512 bytes are provided, the pattern will be expanded to 512 bytes by repeating the entered pattern multiple times.

Additional actions

- ☒ Check if device contains data
- ☐ Automatically unclip HPA/DCO/AMA

START


Link: **PLAY** Status: ERR INDX CORR DREQ **DISC** FAULT **DADY** BUSY Error: AMRE TONE ABRT -- ICNF -- LANC ICRC ?

The wiping process consecutively starts for each device. Repeat the same with all the drives you want to wipe.

Track the wiping progress


Once the operations have started, track the progress of all tasks on the **Home** screen. It displays the percentage of wiped drive area and the time left until the end of the planned session. To see more details on the progress of an individual wiping session, click on that session.

To reveal the current overall speed of wiping, click the Atola logo in the center of the top bar. In this case, we were able to achieve 18 TB/h. This high-speed wiping capability allows a forensic expert to complete the process of preparing drives for wiping in minimal time.



A wiping session can take longer if a different wiping method is selected. E.g. the NIST 800-88 method implies not only wiping but also rereading of the wiped range. The DoD 5220.22-M method wipes the same range three times.

To ensure maximum transparency and effectiveness, TaskForce documents every operation by creating detailed reports and logs. Click **Reports** at the top and find the report in the list or use the **Search bar** at the top of the page.



How TaskForce verifies that wiping is successful

To verify that wiping is successful, TaskForce applies the following algorithm:

1. One hundred sectors, evenly spaced across the wiping range, are filled with the verification ASCII pattern '*AtolaTechnology*'.
2. The whole drive or a selected range of sectors on it is wiped by applying the method selected by the user ('*Fill sectors with the following pattern*' by default).
3. Finally, TaskForce reads the hundred sectors that were filled during the first step and makes sure that none of them contains the pattern '*AtolaTechnology*'.

Unclip or change HPA, DCO, AMA limitations

DCO (device configuration overlay), HPA (host protected area), or AMA (accessible max address) features were created by hard drive manufacturers as hidden areas reserved for storing vendor utilities or simply to make a drive appear to have a certain number of sectors (smaller than the actual drive capacity).

But it is many years ago that end users learned to modify and write to these areas of hard drives with the help of open source and freely available tools. For digital forensics specialists, it means that without the ability to identify such hidden areas of a drive and image the full physical image including data in these areas, the evidence they get may be incomplete and lead to inaccurate investigative conclusions.

Atola TaskForce helps you detect, unclip, or change HPA, DCO, AMA limitations.

Detect DCO, HPA, or AMA limitations

When you connect a hard drive to the TaskForce unit, in addition to the standard *Identify device* command, TaskForce software automatically sends two commands to look up the drive size as set in drive's firmware: *Read native max address* and *Device configuration identify*. If drive size has been limited by DCO, HPA, or AMA, TaskForce will draw your attention to these changes by adding the note in red color in the device menu.

Home

CASES

REPORTS

Diagnose

Image

Logical

RAID

Wipe

Other

Select device

Power on all devices

SATA

1

WDC WD10JPVT-75A1YT0
W XF1C12Y0097
973 GB - limited via HPA

Case: Not assigned

source

2

Device not detected

target

3

WD2500AAKS-00F0A0
W CAT1F625038
250 GB

Case: Not assigned

source

4

ST320LT014-9YK142
W0Q326G4
307 GB - limited via DCO

Case: Not assigned

target

5

ST16000NM001G-2KK103
W L2013X5
15.9 TB - limited via AMA

Case: Not assigned

target

6

HGST HTS545050A7E680
R B697AMB G0ULNN
500 GB

Case: Not assigned

target

SAS

1

Device not detected

target

2

Device not detected

target

3

Device not detected

target

4

Device not detected

target

5

Device not detected

target

6

Device not detected

target

USB

IDE

EXT

FILE

RAID

Notification about HPA, DCO, and AMA in device menu.

To get more details about the modifications that have been made to the drive's firmware, run **Diagnose** and see the **Firmware** section of the **Diagnostics report**.

Atola TECHNOLOGY

Home | CASES | REPORTS | DEVICES

Diagnose
Image
Logical
RAID
Wipe
Other

Device Information:
 ST16000NM001G-2KK103 WL2013X5
 16 TB
 SATA 3
 Case ID: 23

Metadata:
 Creation date: 09/02/2023 5:29 PM
 TaskForce serial: 28882888
 Location: SATA 3
 Device model: ST16000NM001G-2KK103
 Size: 16 TB (16,000,900,608,512 bytes)
 Case investigator:
 Software version: 2023.7
 TaskForce IP: 10.0.0.246
 Write protection: Off
 Device serial: WL2013X5
 Device firmware: SN03
 Case description:

Diagnostics completed. Device has issues

Results
 No major hardware or firmware issues found
 AMA is active. To be able to access the entire disk space you need to disable AMA. Go to Other => Hidden drive areas.
 Found partition of an unknown type at sector 34, Label: Microsoft reserved partition, Partition size: 134 MB
 Found partition of an unknown type at sector 7,079,936, Label: ext4_2048, Partition size: 537 MB
 Found partition of an unknown type at sector 139,352,064, Label: xfs_v4_64k, Partition size: 458 MB
 99% of the disk is not associated with any partition.
 Estimated imaging time: 20 hours 52 minutes

1. Circuit board
 Device is powered on. A power cycle is needed...
 Powering down the device...
 Applying power and watching currents...
 Current oscillogram (5V):

PRINT **GO TO CASE** |< < > >|

AMA limitation is indicated in the Diagnostics summary.

There you will see three lines indicating the drive's Max Address according to different records in the drive's firmware:

1. The *Max Address according to device ID* line shows the max address from the ID sector, affected by DCO and HPA/AMA restrictions if those are applied.
2. *Native Max Address* indicates max address ignoring HPA/AMA limitation that may have been enabled, yet affected by DCO restriction.
3. *Max Address from DCO* is the line that gives you the actual drive size.

A **Diagnostics report** of a drive that does not have HPA/AMA or DCO activated will have the same value in all three lines.

Atola

Diagnose Image Logical RAID Wipe Other

CASES REPORTS

DEVICES

ST320LT014-9YK142 WDQ326G4
282 GB - limited via HPA, DCO

Has issues
SATA 4

Case ID:
Not assigned

Hidden drive areas

DCO HPA AMA

DCO max address: 625 142 447

Native max address: 600 000 000
Can be limited via DCO

ID sector max address: 550 000 000
Can be limited via HPA

☐ Change ID sector max address temporarily (until power cycle)

It will unclip or change hidden areas of the ATA drive.

DCO (device configuration overlay), HPA (host protected area) or AMA (accessible max address) features were created by hard drive manufacturers as hidden areas reserved for storing vendor utilities or simply to make a drive appear to have a certain number of sectors (smaller than the actual drive capacity). However, these areas may be used by skilled users to hide files they don't want to be easily accessible.

Not all drives support hidden areas.

The DCO and HPA can co-exist on the same drive; max address limited via HPA should be less than DCO. Expectedly, AMA is supported by new drives and can't exist if DCO or HPA is supported, and vice versa.

UNCLIP

Link: PHY Status: ERR INDEX CORN DRC2 DRDC FAULT DREV BLBY Error: AMR0 T0M0 ABRT SCNF SINC SCRC ?

Remove HPA and DCO by clicking the Unclip button.

TaskForce lifts HPA/AMA and DCO restrictions in a matter of seconds and enables access to all data on the drive.

The screenshot shows the Atola TaskForce software interface. On the left is a vertical sidebar with icons for Diagnose, Image, Logical, RAID, Wipe, and Other. The main window displays a report for a 3 TB SATA 4 drive (ST3000DM001-9YN166 Z1F168R5). The report includes the following details:

Creation date:	09/02/2023 5:50 PM	Software version:	2023.7
TaskForce serial:	28882888	TaskForce IP:	10.0.0.246
Location:	SATA 4	Write protection:	Off
Device model:	ST3000DM001-9YN166	Device serial:	Z1F168R5
Size:	3 TB (3,000,592,982,016 bytes)	Device firmware:	CC4B
Case investigator:		Case description:	

The report also shows the results of the unclip operation:

- Unclip hidden area completed**
- HPA restriction removed**
 - Old max address: 1,115,973,254
 - New max address: 1,200,457,120
- DCO restriction removed**
 - Old max address: 1,200,457,120
 - New max address: 5,860,533,167

At the bottom of the report, there are buttons for PRINT and GO TO CASE, and a set of navigation arrows.

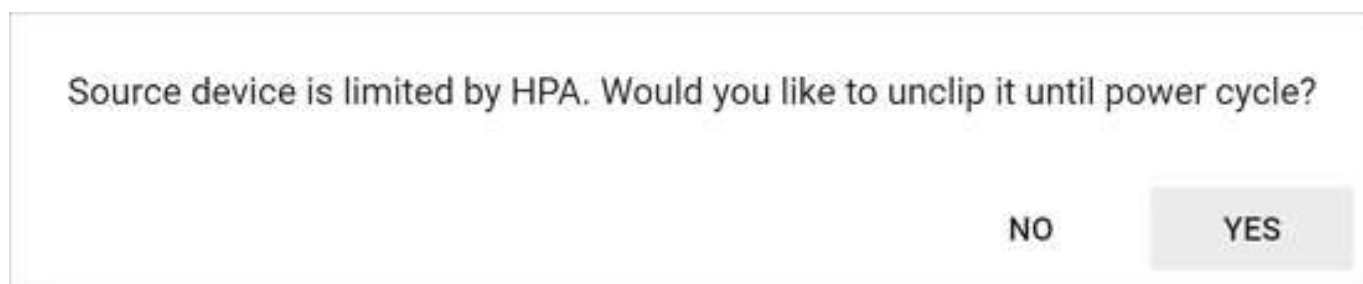
HPA and DCO unclip report.

Unclip HPA temporarily (until power cycle)

To ensure the forensically sound process, it can be necessary to avoid making any changes to the drive. Therefore it is prohibited to disable HPA and DCO restrictions and access data in the hidden areas. With TaskForce it is possible to lift HPA restriction until the next power cycle. This helps avoid permanent changes to the drive.


To unclip HPA on the source drive until power cycle before imaging:

1. Go to **Image**.
2. Select source device.
3. Select target devices and click **Continue**.
4. In **Confirmation** dialog, suggesting you unclip the drive until power cycle, click **Yes**.



This will allow temporary access to the data in HPA-protected area, but as soon as you power off or unplug the drive, the HPA will be back again.

After you confirm unclipping HPA until power cycle, the imaging process starts and the following message appears in the imaging log: *Source device HPA was set to native max address until power cycle.*




The message in the imaging log informs that source device HPA was set to native max address until power cycle.

If a drive freezes in the course of imaging TaskForce forcibly performs power cycles to continue imaging the drive. However, such power cycles will not affect the temporarily disabled HPA. TaskForce will be temporarily removing HPA max address restriction after each power cycle performed during imaging. The HPA-protected area will remain accessible throughout the imaging process.

Set or change HPA, DCO, AMA limitations

Not all drives support hidden areas. Limitation type supported by the particular drive will be shown in green on the **Hidden drive areas** page.



The drive supports DCO and HPA limitations, but does not support AMA.

The DCO and HPA can co-exist on the same drive: max address limited via HPA should be less than DCO.

Expectedly, AMA is supported by new drives and can't exist if DCO or HPA is supported, and vice versa.

If your target device is larger than your source device, but you need hash values for the source and for the target devices to be identical, see [Clip target drive to source evidence size](#).

To set or change DCO limitation:

1. Make sure that the drive is in the Target mode.
2. On the left, click **Other** and then **Hidden drive areas**.
3. Select device.
4. Enter new **Native max address**. Notice that **ID sector max address** changes accordingly.
5. Click **Change**.

Atola

CASES REPORTS

Diagnose Image Logical RAID Wipe Other

ST320LT014-9YK142 WDQ326G4 320 GB

Has issues SATA 4

Case ID: Not assigned

Hidden drive areas

DCO HPA AMA

DCO max address: 625 142 447

Native max address: 600 000 000
Can be limited via DCO

ID sector max address: 600 000 000
Can be limited via HPA

☐ Change ID sector max address temporarily (until power cycle)

CHANGE

It will unclip or change hidden areas of the ATA drive.

DCO (device configuration overlay), HPA (host protected area) or AMA (accessible max address) features were created by hard drive manufacturers as hidden areas reserved for storing vendor utilities or simply to make a drive appear to have a certain number of sectors (smaller than the actual drive capacity). However, these areas may be used by skilled users to hide files they don't want to be easily accessible.

Not all drives support hidden areas.

The DCO and HPA can co-exist on the same drive: max address limited via HPA should be less than DCO. Expectedly, AMA is supported by new drives and can't exist if DCO or HPA is supported, and vice versa.

UNCLIP

Link: PHY Status: ERR INDEX CORN DRC DRC FAULT DREV BUSY Error: AMR TORP ABRT SCNF SINC SCRC ?

Setting new DCO limitation.

To set or change HPA limitation:

1. Make sure that the drive is in the Target mode.
2. On the left, click **Other** and then **Hidden drive areas**.
3. Select device.
4. Enter new **ID sector max address**.
5. Optional: Check **Change ID sector max address temporarily (until power cycle)** if needed.
6. Click **Change**.

Atola

CASES REPORTS

Diagnose Image Logical RAID Wipe Other

ST320LT014-9YK142 WDQ326G4
320 GB

Has issues
SATA 4

Case ID:
Not assigned

Hidden drive areas

DCO HPA AMA

DCO max address: 625 142 447

Native max address: 625 142 447
Can be limited via DCO

ID sector max address: 500 000 000
Can be limited via HPA

☐ Change ID sector max address temporarily (until power cycle)

CHANGE

It will unclip or change hidden areas of the ATA drive.

DCO (device configuration overlay), HPA (host protected area) or AMA (accessible max address) features were created by hard drive manufacturers as hidden areas reserved for storing vendor utilities or simply to make a drive appear to have a certain number of sectors (smaller than the actual drive capacity). However, these areas may be used by skilled users to hide files they don't want to be easily accessible.

Not all drives support hidden areas.

The DCO and HPA can co-exist on the same drive: max address limited via HPA should be less than DCO. Expectedly, AMA is supported by new drives and can't exist if DCO or HPA is supported, and vice versa.


UNCLIP

Link: PHY Status: ERR INDEX CORR DRC DRDC FAULT DREV BLIST Error: AMR TORP ABRT SCNF SINC SCRC ?

Setting new HPA limitation.

To set or change AMA limitation:

1. Make sure that the drive is in the Target mode.
2. On the left, click **Other** and then **Hidden drive areas**.
3. Select device.
4. Enter new **ID sector max address**.
5. Click **Change**.




Setting new AMA limitation.

Calculating hash during imaging

Atola TaskForce supports hash calculation of both the evidence drive and the image in conjunction with imaging. We have developed highly flexible functionality to help optimize evidence acquisition process to fit one's internal procedures, while avoiding further damage to fragile media.

To calculate hash of both the evidence and the image:

1. On the left, click **Image**.
2. Select the source and target device.
3. On the imaging **Settings** page, click **Change**.
4. On the **Hashes** tab, toggle one or several options how you want to calculate hash:
 - Pre-hash source device
 - Hash source during imaging
 - Post-hash target device



Selecting hash methods.

Pre-hash source drive option must be used with caution: although pre-hashing can be required by an investigator's internal procedures, when dealing with drives that have been diagnosed with hardware failure, this operation may cause further damage to the drive before essential data is imaged.

Hash source during imaging is the most appropriate way to calculate the hash of a fragile source evidence drive. In this case, TaskForce only needs to read the data on the drive once to both image and calculate the hash, thus minimally using the drive's hardware.

Linear hash can only be calculated by reading data in sectors consecutively in one pass. When it encounters a bad sector, linear hash calculation is discontinued. For bad drives we recommend using Segmented hashing.

Post-hash target device option lets you to properly record the calculated hash in the case.

Alternatively, you can look up the metadata stored in the E01 file itself:

1. At the top, click **Devices**.
2. Expand the **File** category and click **Select file**.
3. Select the E01 file in the file browser.

The screenshot displays the Atola Technology software interface. The top navigation bar includes a home icon, 'CASES', 'REPORTS', the Atola Technology logo, 'DEVICES', and a menu icon. A left sidebar contains icons for 'Diagnose', 'Image', 'Logical', 'RAID', 'Wipe', and 'Other'. The main content area shows details for a file named 'TEAM T253240GB.E01' (240 GB) located at '//S/shared/Home'. It includes a dropdown for 'Case ID' (Not assigned) and fields for 'LBA sectors: 468 862 128' and 'Path: //S/shared/Home/TEAM T253240GB.E01'. Below this is a section titled 'E01 metadata' containing a table of file information. Two orange arrows point to the 'Acquisition MD5' and 'Acquisition SHA1' rows. The 'Reports' section at the bottom shows a search bar and a table with columns 'Date' and 'Title', currently displaying 'No reports yet'.

E01 metadata	
Model	TEAM T253240GB
Serial number	TPBF2208080020101272
Acquisition MD5	c4ceaa819d7a0bc3c1a0078b2cce1d09
Acquisition SHA1	bad86efb1aa2a5ba77e7d4f78e8ea271ff97edce
Evidence number	266-23
Examiner name	F. M. MacNelly
Drive type	Fixed
Acquisition date	Sep 4, 2023 9:37:35 AM
System date	Sep 4, 2023 9:37:35 AM
Case number	266
Description	Drive subpoenaed in accordance with summon order #3141592 from witness Istar S. White
Notes	
Software version	Atola230700
Compression	Best speed

Reports	
Date ↓	Title
No reports yet	

Hash calculated during imaging stored in E01 file's metadata

Recalculate the hash for an E01 file

To ensure the integrity of the data in the file, you can recalculate its hash.

1. On the left taskbar, click **Other**, and then **Hash**.
2. To choose the file for which you want to calculate hash, expand the **File** category and click **Select file**.
3. Select the E01 file in the file browser.
4. Adjust hashing settings. Make sure to select the same hashing types (MD5, SHA1, etc.)
5. Click **Start**.

Atola TECHNOLOGY

CASES REPORTS

Diagnose Image Logical RAID Wipe Other

TEAM T253240GB.E01
240 GB

//S:/shared/Home

Case ID:
Not assigned

Hashing

Start LBA: 0 End LBA: 468 862 127

Hash unreadable sectors using pattern: 00

LOAD

☒ MD5
☒ SHA1
☐ SHA256
☐ SHA512

START

Start hash calculation

When the hash calculation is completed, you can make sure that the two sets of hashes are identical.

	A	B	C	D	E	F
1	07516243b1448188c21588058171fc68	0	8388607			
2	a60a5a885d8105de4b4fe2d709bc39d9	8388608	16777215			
3	ee049cc72bfe704b39bc3bc470b0ac84	16777216	25165823			
4	60012c7e0a346574107e79d81055442e	25165824	33554431			
5	90f52962bbc20e0134b4acd06b89099c	33554432	41943039			
6	e5314db38ca4ebc5076488566f703b49	41943040	50331647			
7	c1a07f9c9d19a2051d3becdb962fedfb	50331648	58720255			
8	6b3b1f2586b7b60c064ed77d7bcb2497	58720256	67108863			
9	c9a04a7d84c68231635676b1c7b1057f	67108864	75497471			
10	37a7c995b411f9aab6701311e984ca3e	75497472	83886079			
11	acba3e4cfff1f2949aed9751506aba26	83886080	92274687			
12	b809ed82b43501c44b64cb32d5de6032	92274688	100663295			
13	3310c61aace8edbe15f1f1928e0bf92e	100663296	109051903			
14	5ddaeac511afbfe716b9cefdc2a704a0	109051904	117440511			
15	3e0ebd9f01eb5b02b730308e9bfa99b3	117440512	125829119			
16	86e8555e17b524a9f9bf3531e546790b	125829120	134217727			
17	d471779fa3d1706a5f38e2d4d4a4736c	134217728	142606335			
18	118881def5f6b51350765df175ffbf2a	142606336	150994943			
19	bd920def4b75e5a858db4619beac86af	150994944	159383551			
20	ee316fa1b478f5fcc6a038c126e3e15c	159383552	167772159			
21	b6a1ece24f831ed411e6ae4fa769a27f	167772160	176160767			
22	a60a5a885d8105de4b4fe2d709bc39d9	176160768	184549375			
23	ee049cc72bfe704b39bc3bc470b0ac84	184549376	192937983			
24	60012c7e0a346574107e79d81055442e	192937984	201326591			
25	90f52962bbc20e0134b4acd06b89099c	201326592	209715199			
26	e5314db38ca4ebc5076488566f703b49	209715200	218103807			
27	c1a07f9c9d19a2051d3becdb962fedfb	218103808	226492415			
28	6b3b1f2586b7b60c064ed77d7bcb2497	226492416	234881023			
29	c9a04a7d84c68231635676b1c7b1057f	234881024	243269631			
30	37a7c995b411f9aab6701311e984ca3e	243269632	251658239			
31	acba3e4cfff1f2949aed9751506aba26	251658240	260046847			
32	b809ed82b43501c44b64cb32d5de6037	260046848	268435455			
33	3310c61aace8edbe15f1f1928e0bf92e	268435456	276824063			
34	5ddaeac511afbfe716b9cefdc2a704a0	276824064	285212671			
35	3e0ebd9f01eb5b02b730308e9bfa99b3	285212672	293601279			
36	86e8555e17b524a9f9bf3531e546790b	293601280	301989887			

<
+
MD5_HashSegments-Samsung SSD 860 PRO 256GB_S418NF0KB07041E_RVM01B6Q


Segmented hashing vs regular hashing

With conventional hashing method you get a single hash for the entire image, while segmented hashing allows getting many hashes of corresponding LBA ranges of the image. The sum LBA ranges represents the entire image.

Verifying all hashes in a set allows you to prove that the entire image has not been modified.


Segmented hashing and post-hashing of the target for immediate image verification

1. Go to the imaging **Settings**.
2. On the **Hashes** tab, select **Segmented** hashing method.
3. To obtain both sets of hashes for the evidence drive and the image, toggle **Post-hash target devices**.




Selecting hashing method.

Hashing while imaging does not slow down the imaging session:



TaskForce imaging session.

Post-hashing will commence as soon as the imaging session is completed:



Atola TaskForce: Post-hashing.

In the **Imaging completed** report you can see imaging results with the link to the file with segmented hashes.

In case you select the post-hashing of the target, you also get the results of cross-checking between the hash sets of the evidence drive and the image.

	A	B	C	D	E	F
1	07516243b1448188c21588058171fc68	0	8388607			
2	a60a5a885d8105de4b4fe2d709bc39d9	8388608	16777215			
3	ee049cc72bfe704b39bc3bc470b0ac84	16777216	25165823			
4	60012c7e0a346574107e79d81055442e	25165824	33554431			
5	90f52962bbc20e0134b4acd06b89099c	33554432	41943039			
6	e5314db38ca4ebc5076488566f703b49	41943040	50331647			
7	c1a07f9c9d19a2051d3becdb962fedfb	50331648	58720255			
8	6b3b1f2586b7b60c064ed77d7bcb2497	58720256	67108863			
9	c9a04a7d84c68231635676b1c7b1057f	67108864	75497471			
10	37a7c995b411f9aab6701311e984ca3e	75497472	83886079			
11	acba3e4cfff1f2949aed9751506aba26	83886080	92274687			
12	b809ed82b43501c44b64cb32d5de6032	92274688	100663295			
13	3310c61aace8edbe15f1f1928e0bf92e	100663296	109051903			
14	5ddaeac511afbfe716b9cefdc2a704a0	109051904	117440511			
15	3e0ebd9f01eb5b02b730308e9bfa99b3	117440512	125829119			
16	86e8555e17b524a9f9bf3531e546790b	125829120	134217727			
17	d471779fa3d1706a5f38e2d4d4a4736c	134217728	142606335			
18	118881def5f6b51350765df175ffbf2a	142606336	150994943			
19	bd920def4b75e5a858db4619beac86af	150994944	159383551			
20	ee316fa1b478f5fcc6a038c126e3e15c	159383552	167772159			
21	b6a1ece24f831ed411e6ae4fa769a27f	167772160	176160767			
22	a60a5a885d8105de4b4fe2d709bc39d9	176160768	184549375			
23	ee049cc72bfe704b39bc3bc470b0ac84	184549376	192937983			
24	60012c7e0a346574107e79d81055442e	192937984	201326591			
25	90f52962bbc20e0134b4acd06b89099c	201326592	209715199			
26	e5314db38ca4ebc5076488566f703b49	209715200	218103807			
27	c1a07f9c9d19a2051d3becdb962fedfb	218103808	226492415			
28	6b3b1f2586b7b60c064ed77d7bcb2497	226492416	234881023			
29	c9a04a7d84c68231635676b1c7b1057f	234881024	243269631			
30	37a7c995b411f9aab6701311e984ca3e	243269632	251658239			
31	acba3e4cfff1f2949aed9751506aba26	251658240	260046847			
32	b809ed82b43501c44b64cb32d5de6037	260046848	268435455			
33	3310c61aace8edbe15f1f1928e0bf92e	268435456	276824063			
34	5ddaeac511afbfe716b9cefdc2a704a0	276824064	285212671			
35	3e0ebd9f01eb5b02b730308e9bfa99b3	285212672	293601279			
36	86e8555e17b524a9f9bf3531e546790b	293601280	301989887			

<
+
MD5_HashSegments-Samsung SSD 860 PRO 256GB_S418NF0KB07041E_RVM01B6Q

Benefits of segmented hashing

With the conventional **linear hashing** method you get a single hash for the entire image. As the linear hashing stops upon encountering the first bad sector, it is impossible to calculate hash for the entire space of the source evidence drive.


With **segmented hashing**, hashing can be performed during the [multipass imaging](#) of a damaged drive. Hashes are calculated only for the successfully imaged areas, while all bad sectors are excluded from the calculation.

If an acquired evidence image is damaged at some point in the future, with the regular linear hashes you will get a hash mismatch upon [verification](#), and the entire image becomes useless. With segmented hashes only the hash of the damaged segment will become invalid.

Calculate segmented hashes of a source and target devices


To calculate segmented hashes of a damaged drive during imaging and post-hash the target for immediate image verification, do the following:

1. Go to the imaging **Settings**.
2. On the **Hashes** tab, select **Segmented** hashing method and specify **Segment size**.
3. To obtain both sets of hashes for the evidence drive and the image, toggle **Post-hash target devices**.




Selecting hashing method.

Hashing while imaging does not slow down the imaging session:



TaskForce imaging session.

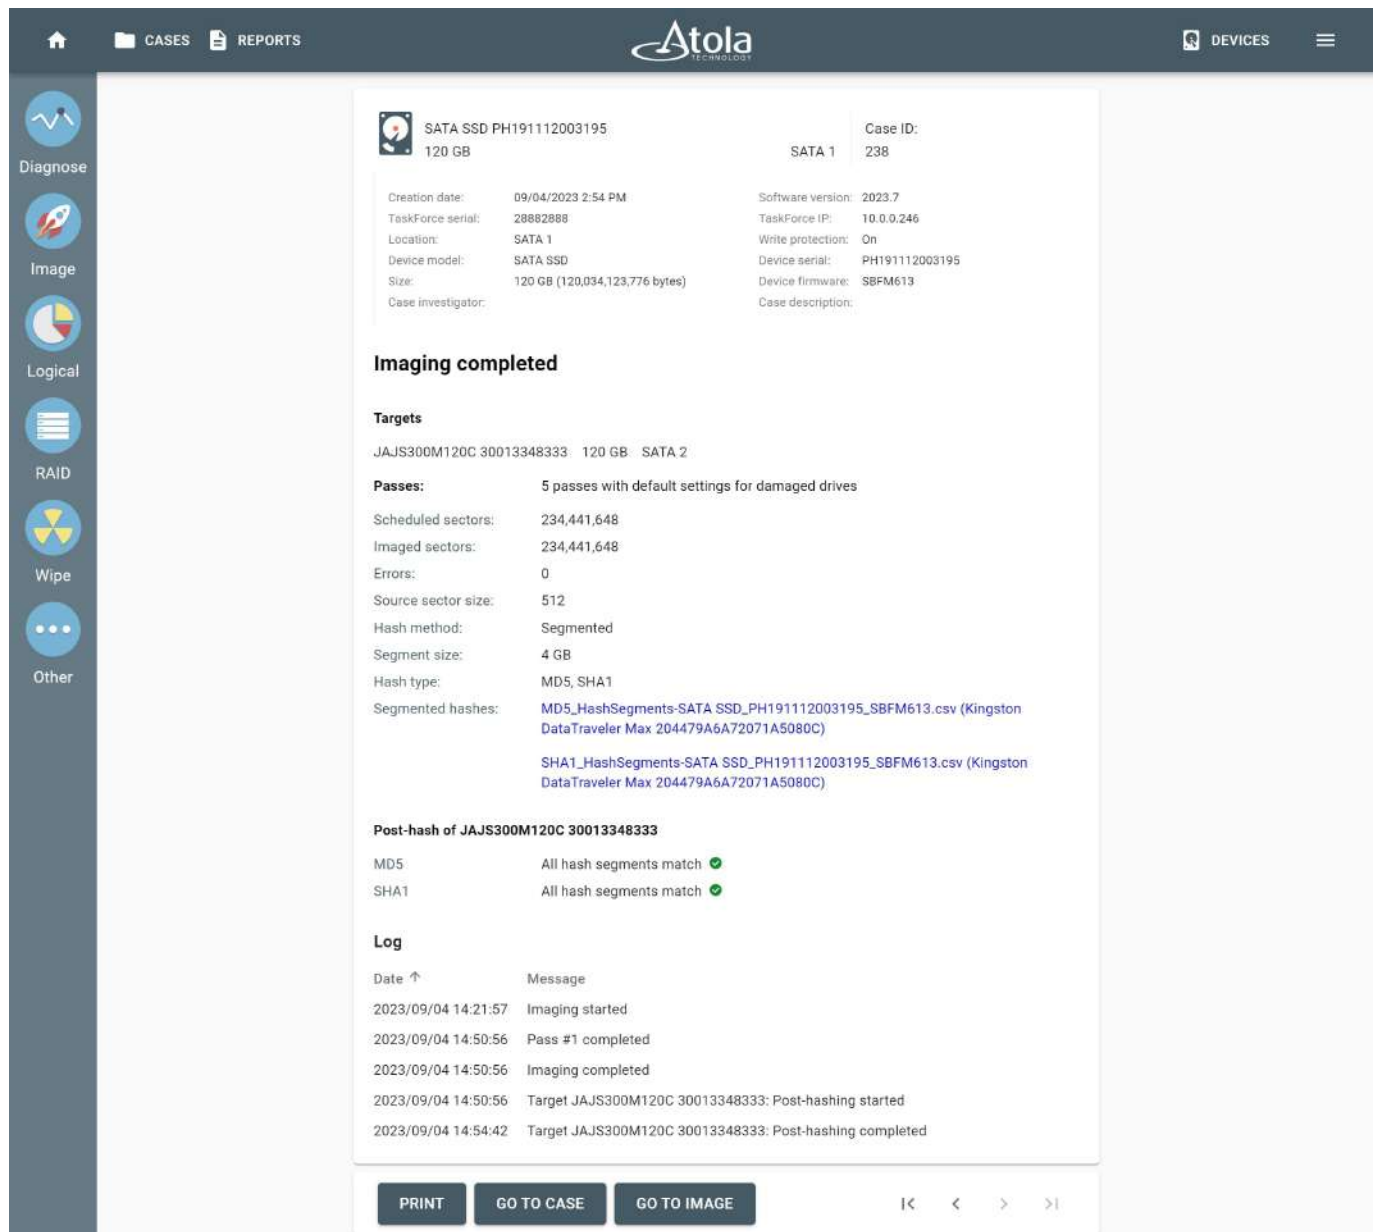
Post-hashing commences as soon as the imaging session is completed:



Atola TaskForce: Post-hashing.


In the **Imaging completed** report, you can see imaging results with the link to the file with segmented hashes.

In case you select the post-hashing of the target, you also get the results of cross-checking between the hash sets of the evidence drive and the image.




Verify an image of a drive with segmented hashing

1. In the TaskForce window, go to **Other > Verify segmented hashes**.



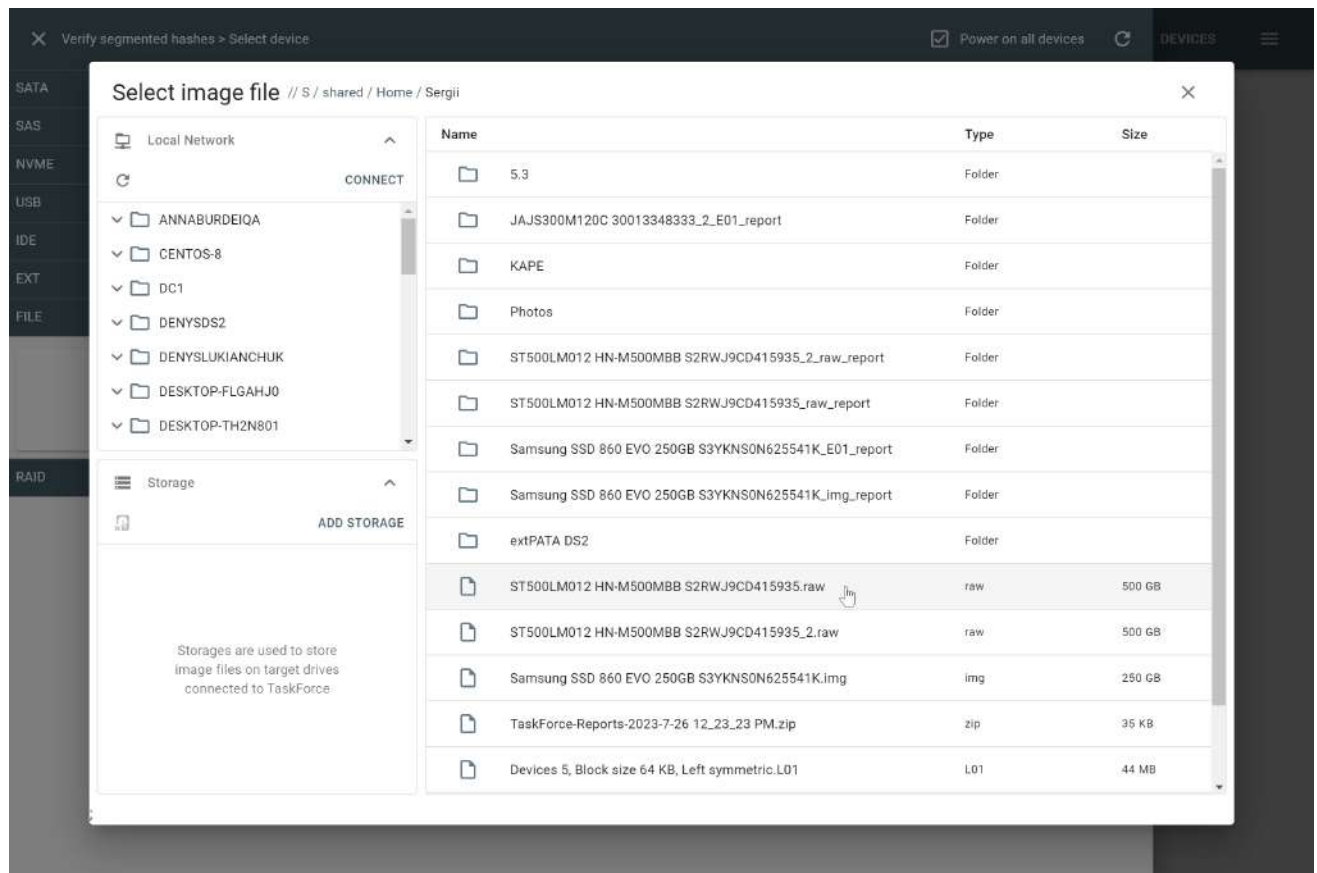
The Verify segmented hashed command on the Other page.

2. The Select device panel opens. Expand the File section and click Select file.



The Select device panel.

3. Select an image file you want to verify. The E01, AFF4 and Raw formats are supported.




Selecting an image file to verify segmented hashes.

4. Select the CSV file with segmented hashes that relates to your image file:


- Choose either **Local folder** or **Network folder** or **Storage**.
- Click **Select**.
- Find and select a CSV file with segmented hashes.

5. Click **Start**.



Selecting a CSV file with a list of segmented hashes.

6. TaskForce starts the data verification process.



The process of verifying segmented hashes.

7. If TaskForce encounters a hash value mismatch, it is reflected in the **Hash mismatches** counter and in the event log, with *Start and End LBA* of the respective segment.

1

HOME

CASES

REPORTS

Atola
TECHNOLOGY

DEVICES

≡

Diagnose

Image

Logical

RAID

Wipe

Other

ST500LM012 HN-M500MBB S2RWJ9CD415935.raw

500 GB

//S/shared/Home/Sergii

Case ID:
Not assigned

Verify segmented hashes

84%

Speed: 84 MB/s

Hash mismatches: 1

Errors: 0

Last processed LBA: 255,566,719

Hashes processed: 55

Hashes left: 10

Hashes skipped: 0

Date ↑	Message
2024/01/09 14:22:24	Verification of segmented hashes started
2024/01/09 14:45:58	Hash value mismatch at segment: 245,769,088 - 254,157,695

STOP

Hashes of one of the segments do not match.

8. Once verification of segmented hashes is completed, TaskForce generates a detailed report about its results. The report contains information about the verified image and the file with its segmented hashes, hash type, number of processed hashes and found mismatches if any.

The screenshot displays the Atola TaskForce web interface. The top navigation bar includes 'CASES' and 'REPORTS' tabs, the Atola logo, and a 'DEVICES' section. A left sidebar contains icons for 'Diagnose', 'Image', 'Logical', 'RAID', 'Wipe', and 'Other'. The main content area shows a case summary for 'ST500LM012 HN-M500MBB S2RWJ9CD415935.raw' (500 GB) with Case ID 'Not assigned'. Below this, a table lists details: Creation date (01/09/2024 2:54 PM), TaskForce serial (28882888), Location (//S/shared/Home/Sergii), Software version (2024.1), TaskForce IP (10.0.0.4), Image file (ST500LM012 HN-M500MBB S2RWJ9CD415935.raw), Size (500 GB (500,107,862,016 bytes)), Case investigator, and Case description. A section titled 'Verification of segmented hashes completed' follows, showing the selected file path, hash type (MD5), and 65 hashes in the file. The results table indicates 65 hashes processed, 0 left, 0 skipped, 1 mismatch, and 0 errors. A log section at the bottom shows three entries: 'Verification of segmented hashes started' (2024/01/09 14:22:24), 'Hash value mismatch at segment: 245,769,088 - 254,157,695' (2024/01/09 14:45:58), and 'Verification of segmented hashes completed' (2024/01/09 14:54:01). At the bottom are 'PRINT' and 'GO TO CASE' buttons, and navigation arrows.

The report about completed verification of segmented hashes.

Case management system and report types

TaskForce's case management system records every step of the data acquisition process: every operation is automatically added to the case from the moment a device is identified including date, time, imaging map and hash values. When a hard drive is imaged, its imaging map is recorded detailing all the sectors that have been skipped.

Whenever an operator connects a hard drive to the TaskForce, the system makes an automatic database lookup and retrieves all past records associated with that particular hard drive. New entries will be added seamlessly to the database. You do not need to enable case management or take any additional actions for it to start functioning; it is fully embedded into TaskForce and works at all times.

Case number can be assigned and changed at any time. The system also allows browsing through all cases and reports, without corresponding devices being connected to the unit.

Report types and formats

There are two types of reports in TaskForce:

1. Device reports are created every time an action is taken to the drive: drive identification, imaging, hashing, wiping and other operations related to the drive are documented in these reports.
2. Non-device reports are created to register any changes made to the cases: case opening, case details change, case import and export.

All reports have these key elements: a header that provides device and case details, an action summary and task details (task settings, task log, etc.).



Diagnose



Image



Logical



RAID



Wipe



Other



Samsung SSD 860 PRO 256GB S418NF0KB07067E
256 GB

SATA 5

Case ID:

Not assigned

Creation date: 07/25/2023 7:39 PM

TaskForce serial: 23166733

Location: SATA 5

Device model: Samsung SSD 860 PRO 256GB

Size: 256 GB (256,060,514,304 bytes)

Case investigator:

Software version: 2023.7

TaskForce IP: 192.168.1.113

Write protection: Off

Device serial: S418NF0KB07067E

Device firmware: RVM01B6Q

Case description:

Diagnostics completed

Results

No major hardware or firmware issues found

75% of the disk is not associated with any partition.

Estimated imaging time: 7 minutes

1. Circuit board

Device is powered on. A power cycle is needed...

Powering down the device...

Applying power and watching spin-up currents...

Current oscillogram (12V):



Current oscillogram (5V):



Device has become ready in 3 sec

Registers. Status: 01010000 Error: 00000001

Peak power consumption during spin-up: 5V line = 322.30 mA; 12V line = 0.00 mA

Integrity word of device identification data is OK.

Device identified: Samsung SSD 860 PRO 256GB SN: S418NF0KB07067E

Logical sector size: 512 bytes. Physical sector size: 512 bytes

Circuit board check passed

Elapsed: 12.3 sec.

2. Heads

Solid state drives (SSD) don't have heads.

Elapsed: 0.0 sec.

3. Media surface

Media scan is in progress...

Verifying starting sectors (LBA: 0 - 999,999)

0 error(s) found

Speed: 542 MB/s

Verifying middle sectors (LBA: 250,059,095 - 251,059,094)

0 error(s) found

Speed: 541 MB/s

Verifying ending sectors (LBA: 499,118,192 - 500,118,191)

0 error(s) found

Speed: 542 MB/s

Average speed: 541 MB/s

Media surface check passed

Elapsed: 2.9 sec.

4. Firmware

Device is not locked.

Device identification data is valid.

Max address according to device ID: 500,118,191

Native max address (ext): 500,118,191

Max Address from DCO: 500,118,191.

Reported capacity appears logically correct.

Performing SMART checks...

Power cycles: 1476 Smart status: Good Powered on: 18 days 20 hours

#	Attribute Name	Value	Worst	Threshold	RAW	Status
5	Reallocated Sector Ct	100	100	10	0	OK
9	Power On Hours	99	99	0	452	OK
12	Power Cycle Count	98	98	0	1476	OK
177	Wear Leveling Count	93	93	0	149	OK
179	Used Rsvd Blk Cnt Tot	100	100	10	0	OK
181	Program Fail Cnt Total	100	100	10	0	OK
182	Erase Fail Count Total	100	100	10	0	OK
183	Runtime Bad Block	100	100	10	0	OK
187	Reported Uncorrect	100	100	0	0	OK
190	Airflow Temperature Cel	69	41	0	31	OK
195	Hardware ECC Recovered	200	200	0	0	OK
199	UDMA CRC Error Count	98	98	0	1071	OK
235	POR Recovery Count	99	99	0	334	OK
241	Total LBAs Written	99	99	0	72409915009	OK

Reading SMART temperature...

Temperature and power cycle history

The diagram shows the device temperature history during the recent work time intervals between power cycles.



* Earliest available record. The exact time and duration is unspecified in SMART history.

White gaps represent the interval between temperature measurements. During these periods the temperature sensor was inoperative for periods of time due to power off or standby mode.

Last significant work time interval before power off or standby (in minutes): 20

Firmware check passed

Elapsed: 0.0 sec.

5. File system

Found partition at sector 0 (type exFAT). Label: allExt. Partition size: 63 GB.

75% of the disk is not associated with any partition.

File system structures check complete.

Elapsed: 0.1 sec.

PRINT

GO TO CASE

|< < > >|

Diagnostics report in Atola TaskForce.

Add a case

When you identify a device in Atola TaskForce for the first time, the system automatically creates a new case for that device and records every single operation performed with the device or with the case itself. To know more about how it works, see [Case management system and report types](#).

Until you specify case details, Case ID shows in the TaskForce interface as *Not assigned*.


To distinguish and search your cases by a case number, investigator's name, and case description, you can add a case and enter these case details. Also, you can set TaskForce to remind you to enter case details before starting any task.

There are several ways to create a new case in TaskForce:

- From the **Devices** menu.
- From the **Cases** page.
- From the **Device** page.

Add a case from the Devices menu

1. Connect a device to TaskForce.
2. In the TaskForce main window, click **Devices**.
3. In the **Select device** panel, in the port with your device, click **More** icon, and then select **Add case**.



Adding a device to a new case from the Devices menu.

4. Enter case details and click **Continue**. TaskForce creates a new case and adds your device to it. The **Case created** report appears on the Home screen and on the **Reports** page.

Atola
TECHNOLOGY

CASES

REPORTS

DEVICES

Diagnose

Image

Logical

RAID

Wipe

Other

Completed

Search...

Case created
963

a few seconds ago

Case created
213

3 minutes ago

Diagnostics completed. Device has major issues
415 – SATA 1 WDC WD5000LPVX-22V0TT0 WDWX41AC4DXSC2

19 hours ago

Diagnostics completed
413 – SATA 3 ST1000LM024 HN-M101MBB S318J9CGA16408

19 hours ago

Diagnostics completed
412 – SATA 4 Hitachi HTS722080K9A300 080302DP0B30DQGKLP1P

19 hours ago

Diagnostics completed
414 – SATA 2 WDC WD10SPZX-60Z10T0 WDWXR1E58FTPKL

19 hours ago

Case created

19 hours ago

Case created

19 hours ago

Case details changed

19 hours ago


Case created

19 hours ago

Case created

19 hours ago


The Case created report on the Home screen.



Case created report details.

Add a case from the Cases page


1. Connect a device to TaskForce.
2. In the TaskForce main window, click **Cases**.
3. On the **Cases** page, click **Add case**.
4. In the **Select device** panel, click on your device.
5. Enter case details and click **Continue**. TaskForce creates a new case and adds your device to it. The **Case created** report appears on the Home screen and on the **Reports** page.



The Enter case details dialog.

Add a case from the Device page

1. Connect a device to TaskForce.
2. In the TaskForce main window, click **Devices**.
3. In the **Select device** panel, click on the port with your device. TaskForce takes you to the **Device** page.
4. **Optional:** If TaskForce hasn't identified your device yet, then, on the **Device** page, click **Re-identify**.
5. In **Case ID** pane, open the drop-down menu, and then click **New**.



Change active case menu is opened on the Device page.

6. Enter case details and click **Continue**. TaskForce creates a new case and adds your device to it. The **Case created** report appears on the Home screen and on the **Reports** page.

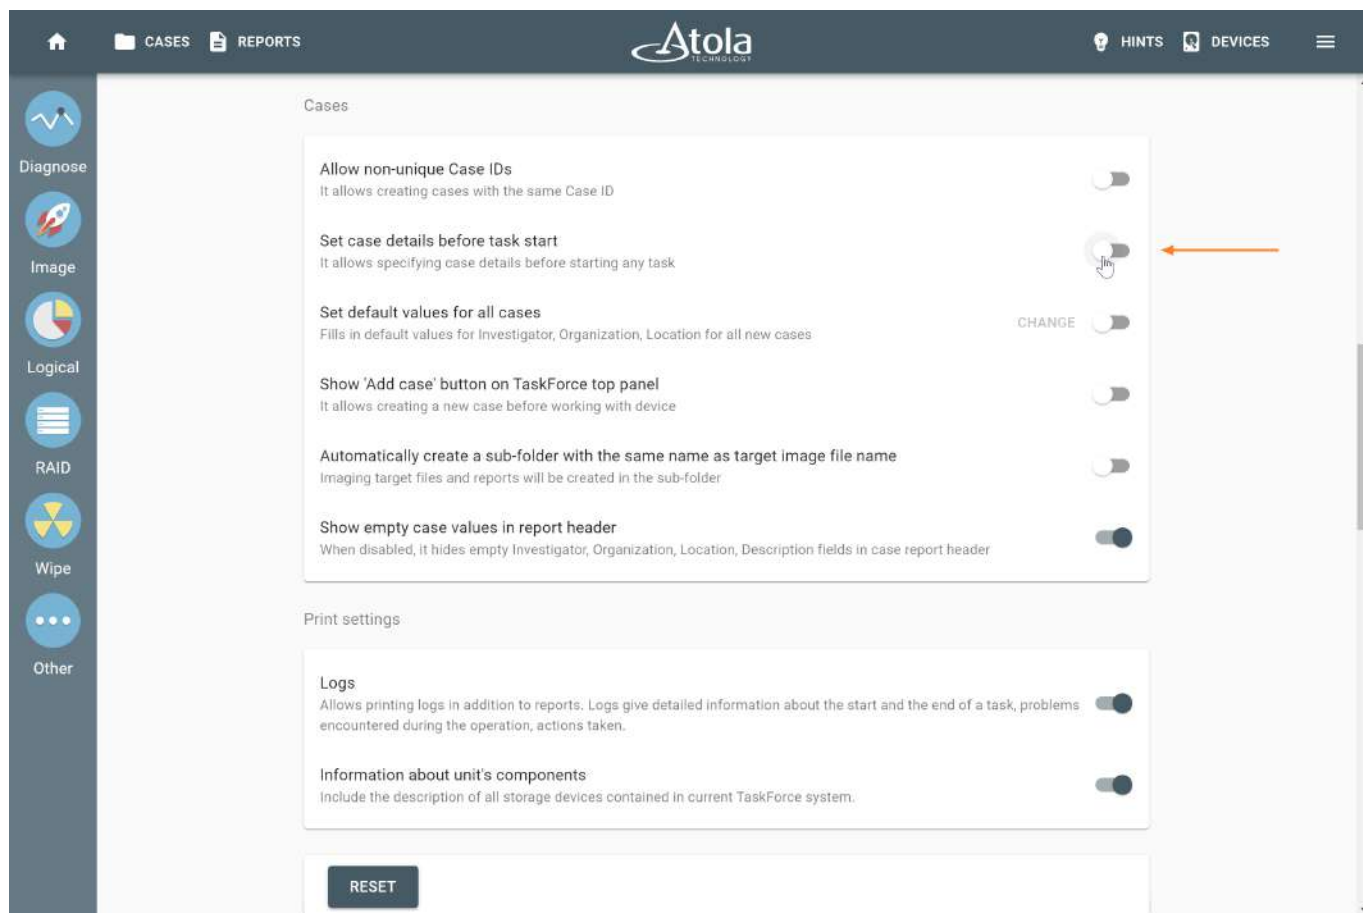
Enter case details before starting any task

Entering case details before performing any operation with a device helps you to keep your cases organized and searchable. TaskForce can be set to request that you enter case details before starting any task.


To enable this feature, do the following:

1. In the TaskForce main window, go to the **Menu > Settings**.
2. In the **Cases** section, toggle **Set case details before task start**.

Now, TaskForce will ask you to enter case details before performing any operation with a device.



The Set case details before task start toggle on the Settings page.

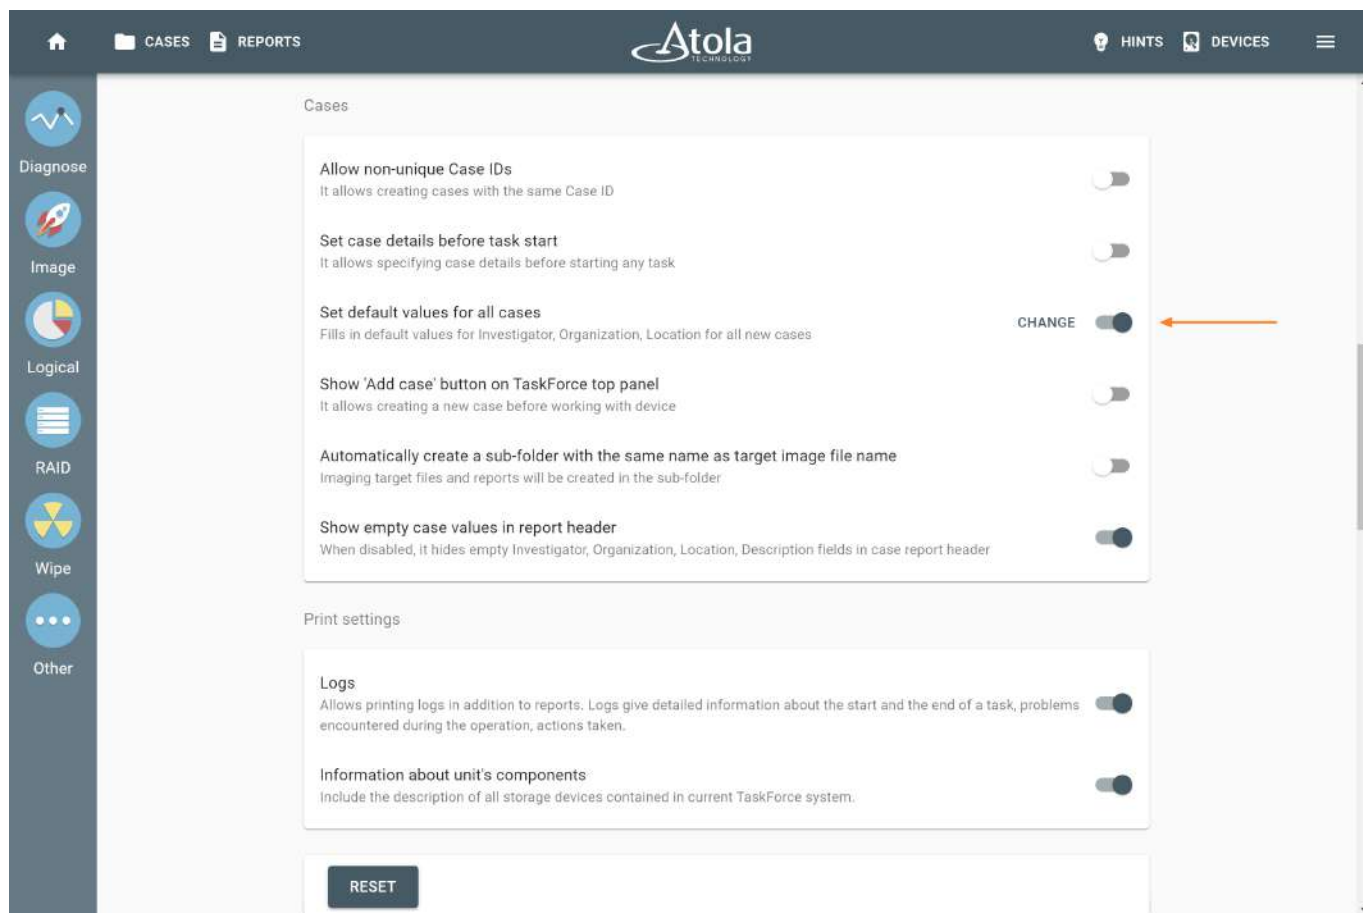


Atola TaskForce asks a user to enter case details before starting diagnostics.

Set default values for all cases


If *Investigator*, *Organization*, or *Location* are the same for all or majority of your cases, set them as default for all new cases, and TaskForce will fill those fields automatically every time you create a new case:

1. At the top right, go to **Menu > Settings**.
2. On the Settings page, scroll down to the **Cases** section.
3. Toggle on **Set default values for all cases**.



The 'Set default values for all cases' toggle on the Settings page.

- Next to the toggle, click **Change**.
- Fill in the default values for *Investigator*, *Organization*, or *Location* and click **Save**.



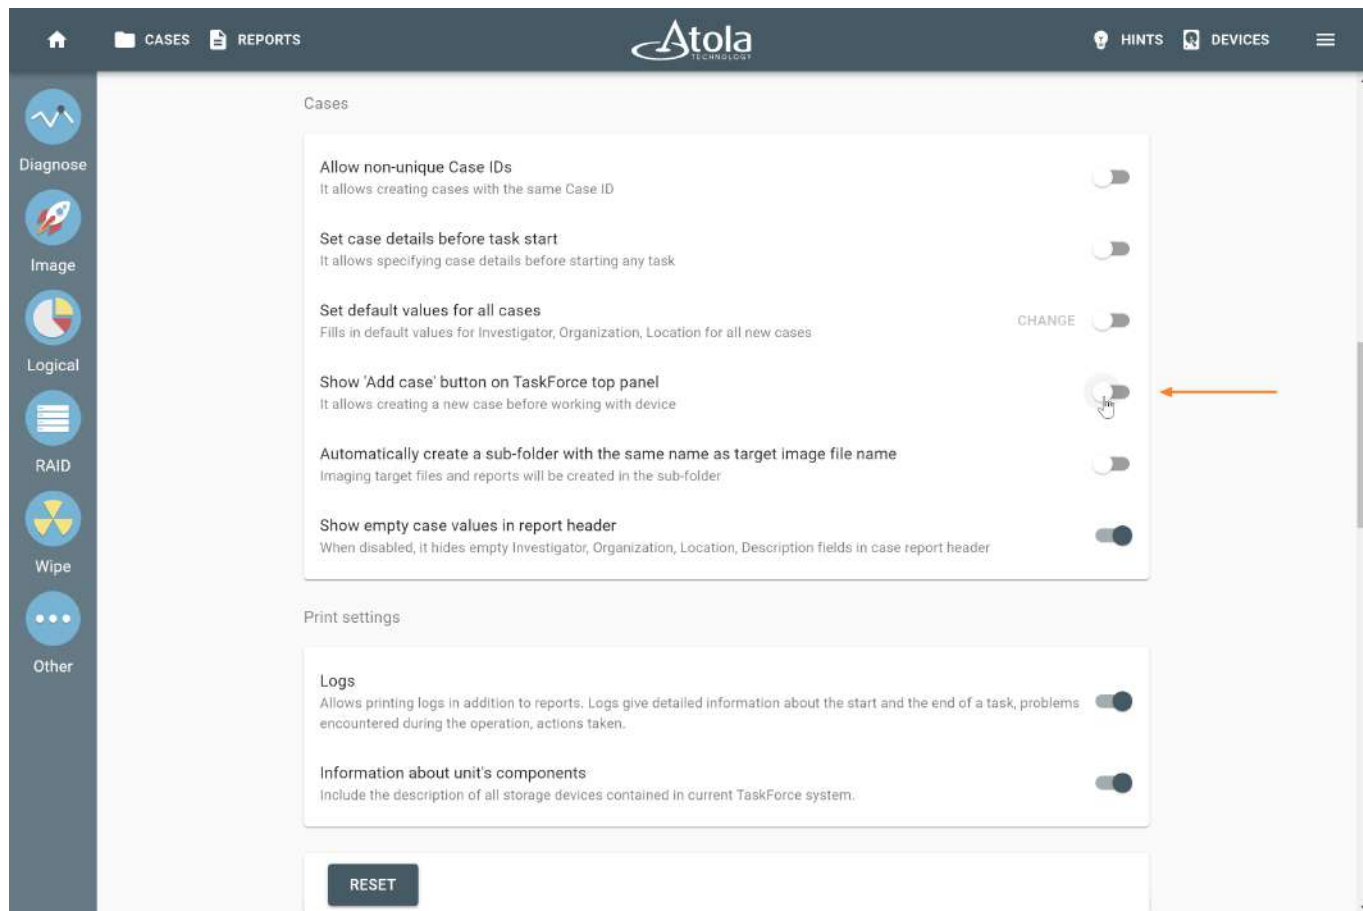
The 'Set default values' dialog.

Show 'Add case' button

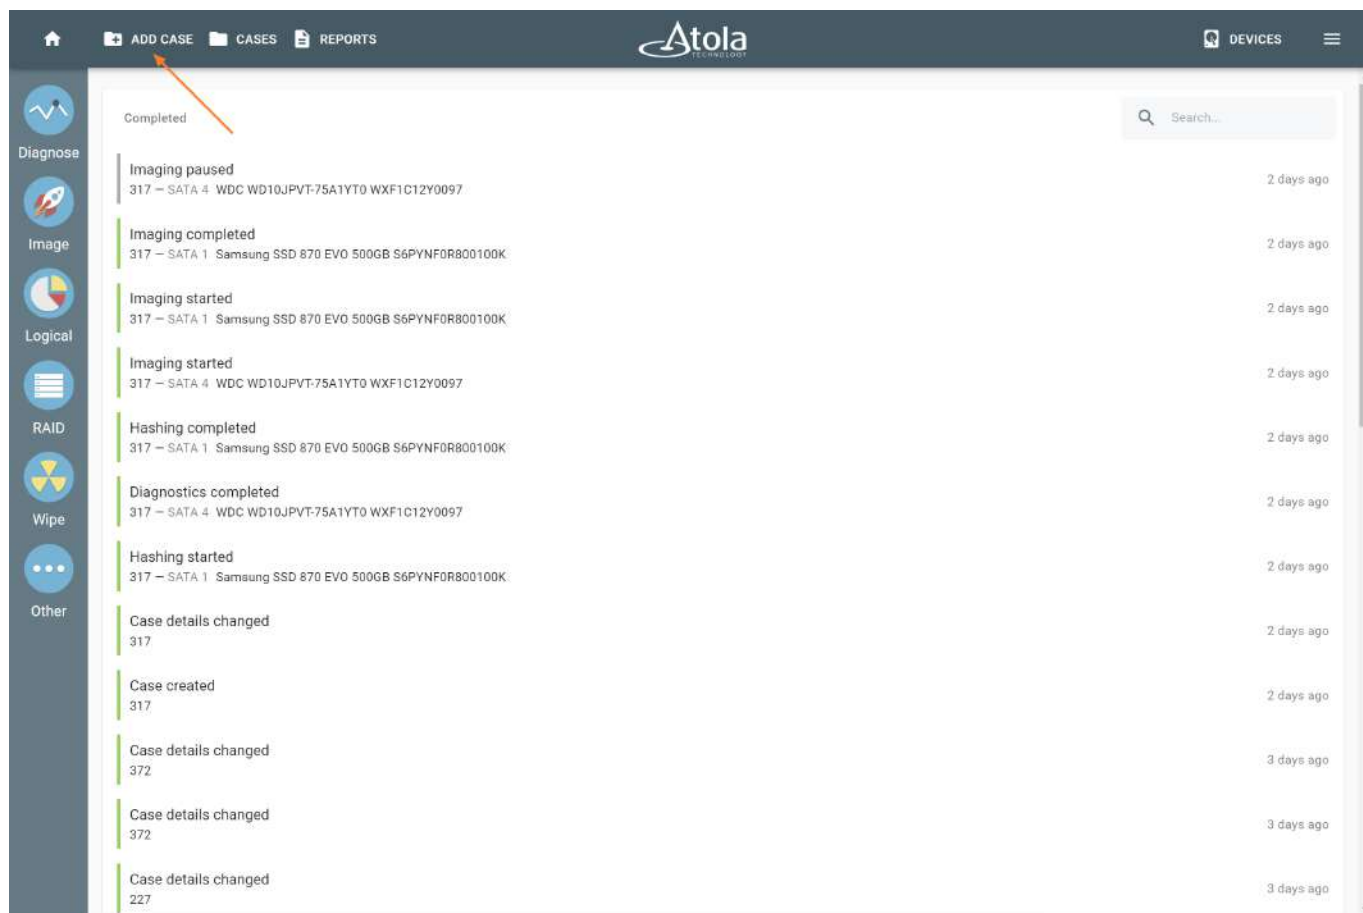
You can add cases right from the TaskForce Home screen before working with the device. To do that, set TaskForce to show **Add case** button in the top panel:

1. In the TaskForce main window, go to the **Menu > Settings**.
2. In the **Cases** section, toggle **Show 'Add case' button** on TaskForce top panel.

Now, TaskForce shows the **Add case** button in the top panel.



The toggle in the Cases section on the Settings page.



The Add case button in the Atola TaskForce main window.

Add one device to several cases

In Atola TaskForce, you can include the same device in several new or existing cases using the [Case management system](#). When a device is added to more than one case, TaskForce keeps the tasks and reports, which are associated with each case, separate.


Add a device to another new case

To add your device to several new cases, do the following:

1. Connect a device to TaskForce and identify it.
2. Add first new case and include your device to it. For guidance, see [Add a case](#).

After you have added your device to the first case, do the following steps:

3. In the TaskForce user interface, click **Devices**.
4. In the **Select device** window, click on the port with your device. TaskForce takes you to the **Device** page.
5. In **Case ID** pane, open the drop-down menu, and then click **New**.



The New button in the Change active case menu on the Device page.


6. Enter case details and click **Continue**.

TaskForce creates another new case and adds your device to it. The **Case created** report appears on the Home screen and on the **Reports** page.

Add a device to another existing case


When you already [added your device to one case](#) and want to add that device to another existing case, do the following:

1. Connect a device to TaskForce.
2. In the TaskForce user interface, click **Devices**.
3. In the **Select device** window, click on the port with your device. TaskForce takes you to the **Device** page.
4. **Optional:** If TaskForce hasn't identified your device yet, then, on the **Device** page, click **Re-identify**.
5. In **Case ID** pane, open the drop-down menu, and then click **Select**.



The Select button in the Change active case drop-down menu on the Device page.

6. In the **Select active case** window, chose your existing case. If needed, use the **Search** field to find it. TaskForce adds your device to the selected case.




List of all your cases in the Select active case window.

Switch between cases on the Device page

When a device is added to more than one case, TaskForce keeps reports, which are associated with each case, separate.

To switch between cases on the **Device** page:

1. In **Case ID** pane, open the drop-down menu.
2. Click on the case you want to make active. TaskForce shows reports, associated with selected case.



List of all cases associated with the device in the Case ID pane.


Add several devices to one case

In Atola TaskForce, you can have a forensic case that contains more than one device. TaskForce lets you add several devices either to a new or an existing forensic case.

To add several devices to a new case, you need to create a case first. For guidance, see [Add a case](#).

To add another device to an existing case, do the following:

1. Connect a device to TaskForce.
2. In the TaskForce user interface, click **Devices**.
3. In the **Select device** window, click on the port with your device. TaskForce takes you to the **Device** page.
4. **Optional:** If TaskForce hasn't identified your device yet, then, on the **Device** page, click **Re-identify**.
5. In **Case ID** pane, open the drop-down menu, and then click **Select**.



Change active case menu is opened on the Device page.

6. In the **Select active case** window, choose your existing case. If needed, use the **Search** field to find it.

2

HOME

CASES

REPORTS

Atola

DEVICES

Diagnose

Image

Logical

RAID

Wipe

Other

Case ID: 317

EDIT

Last open date: 09/05/2023 3:50 PM

Creation date: 09/05/2023 12:58 PM

Investigator: Melanie Jones

Description: WD SATA HDD 1 TB

Samsung SATA SSD 500 GB

Users: (empty)

Devices associated with case

WDC WD10JPVT-75A1YT0 WXF1C12Y0097

1 TB

SATA

Firmware: 0101A01

LBA sectors: 1 953 525 168

Security status: Not locked

LBA48 mode: Supported

Logical sector size: 512 B

Physical sector size: 4096 B

Samsung SSD 870 ... S6PYNF0R800100K

500 GB

SATA

Firmware: SVT01B6Q

LBA sectors: 976 773 168

Security status: Not locked

LBA48 mode: Supported

Logical sector size: 512 B

Physical sector size: 512 B

Active tasks

Imaging

317 - SATA 4 WDC WD10JPVT-75A1YT0 ...

3%

3 hours left

Imaging

317 - SATA 1 Samsung SSD 870 EVO 500...

33%

11 minutes left

Reports

Search...

Date

Title

2023/09/05 15:50:16

Imaging started

PRINT

You can change print setup in TaskForce settings

All devices, active tasks, and reports associated with the case.

Find and edit cases

Atola TaskForce automatically creates reports for every single action applied to each drive connected to it. Whether it is a source drive or a target drive, any action, be it imaging, wiping or physically switching write protection on or off, will be documented and stored in the system.

To find a case, click **Cases** in the top left corner, it will redirect you to the case management system.

Atola TECHNOLOGY

HOME CASES REPORTS DEVICES

Diagnose Image Logical RAID Wipe Other

Completed

Search...

Diagnostics completed	25 minutes ago
496-489-380-495 — USB 4 ADATA USB Flash Drive 21B03215801700B0	
Case details changed	31 minutes ago
Case created	32 minutes ago
Source switch of SATA 4 enabled	3 hours ago
Imaging completed	4 hours ago
WDC WD3200BVVT-63A26Y0_01.01A01_WD-WXN1A71F1539 (only MBR).img	
Imaging completed	4 hours ago
Image_C_hfs_fat32.img	
Imaging started	4 hours ago
Image_C_hfs_fat32.img	
Imaging started	5 hours ago
WDC WD3200BVVT-63A26Y0_01.01A01_WD-WXN1A71F1539 (only MBR).img	
Diagnostics completed	5 hours ago
372 — SATA 4 ST320LT014-9YK142 W0Q326G4	
Imaging completed	5 hours ago
Samsung SSD 850 PRO 256GB_EXM04B6Q_S39KNX0J743780L.img	
Imaging completed	6 hours ago

Opening case management system.

Search for a specific case or device in the Search bar (by case ID, investigator's name or device details) and sort results by any column.

Atola TECHNOLOGY

CASES REPORTS DEVICES

Diagnose Image Logical RAID Wipe Other

Cases

496

☐	Case ID	Investigator	Device	Last open date ↓	Creation date
☐	495-489-380-495	Luke Spikes	HP SSD FX900 Pro 512GB HBSE32302100021 and 4 more	09/07/2023 6:13 PM	06/30/2023 11:10 AM
☐	496	Wesley Crusher	HP SSD FX900 Pro 512GB HBSE32302100021	09/04/2023 6:03 PM	06/30/2023 11:01 AM

1-2 of 2

ADD CASE IMPORT EXPORT DELETE


Searching and sorting cases in the list.

To open a case, click the respective line in the list.

A case page contains case details, information about the devices associated with the case (name, serial number, capacity etc.), as well as reports for all tasks applied to the device.

Edit case details

To change case details, click **Edit** at the top of a case page.



Changing case details.

Finding reports

Atola TaskForce automatically creates reports for every single action applied to each drive connected to it. The system also allows browsing through all cases and reports, without corresponding devices being connected to the unit. The reports are listed and can be easily retrieved in different parts of TaskForce software.

1. Via case page

All reports related to the case are listed at the bottom of the case page. Scroll down and turn pages to view all the reports, sort them by date or by title, use the search bar to look for specific reports by their titles.

To open a report, click the respective line.

The screenshot displays the Atola Technology web interface. The top navigation bar includes 'CASES' and 'REPORTS' tabs, with 'REPORTS' being the active tab. The left sidebar contains icons for 'Diagnose', 'Image', 'Logical', 'RAID', 'Wipe', and 'Other'. The main content area is titled 'Case ID: 294' and includes an 'EDIT' button. Below this, case details are listed: 'Last open date: 09/07/2023 6:56 PM', 'Investigator: Melanie Jones', 'Creation date: 09/02/2023 7:21 PM', and 'Users: (empty)'. A section titled 'Devices associated with case' shows a single device: 'WDC WD800BEVT...DWXF1E10ZF889' (SATA, 80 GB). The device's security status is 'Not locked', and it is supported for LBA48 mode. The reports section, titled 'Reports', contains a table with columns 'Date' and 'Title'. An orange arrow points to the 'Title' column header. The table lists five reports: 'Case details changed' (2023/09/07 18:56:41), 'Diagnostics completed' (2023/09/04 18:49:38), 'Imaging completed' (2023/09/02 20:07:25), 'Imaging started' (2023/09/02 19:31:57), and 'Imaging paused' (2023/09/02 19:31:19). A 'PRINT' button and a link to 'TaskForce settings' are at the bottom of the reports section.

Case ID: 294

Last open date: 09/07/2023 6:56 PM Investigator: Melanie Jones
Creation date: 09/02/2023 7:21 PM Users: (empty)
Description:

Devices associated with case

WDC WD800BEVT...DWXF1E10ZF889 Firmware: 1101A11 LBA sectors: 156 301 488
80 GB Security status: Not locked Logical sector size: 512 B
SATA LBA48 mode: Supported Physical sector size: 512 B
Device was imaged to 3 targets [Show details](#)

Reports


Date ↓	Title
2023/09/07 18:56:41	Case details changed
2023/09/04 18:49:38	Diagnostics completed
2023/09/02 20:07:25	Imaging completed
2023/09/02 19:31:57	Imaging started
2023/09/02 19:31:19	Imaging paused

[PRINT](#) [You can change print setup in TaskForce settings](#)

Case page with reports.

2. Via Reports page

If you need to search among all existing reports, click **Reports** at the top.



The Reports button.

This will redirect you to the page with all existing report that can be filtered by date, title, case ID or device details. Search for a specific report by entering report title or drive details.

Open the report you need by clicking it in the list.

🏠 CASES REPORTS Atola 🔍 DEVICES ☰					
📊 Diagnose 🚀 Image 📊 Logical ☰ RAID ☢️ Wipe ⋮ Other					
Reports 🔍 WDC ✕					
☐	Date ↓	Title	Case ID	Device	
☐	09/04/2023 6:05 PM	Case details changed	Not assigned	Unknown	
☐	09/02/2023 8:07 PM	Imaging completed	294	WDC WD800BEVT-75ZCT2 WDWXF1E10ZF889	
☐	09/02/2023 7:31 PM	Imaging started	294	WDC WD800BEVT-75ZCT2 WDWXF1E10ZF889	
☐	09/02/2023 7:31 PM	Imaging paused	294	WDC WD800BEVT-75ZCT2 WDWXF1E10ZF889	
☐	09/02/2023 7:30 PM	Imaging started	294	WDC WD800BEVT-75ZCT2 WDWXF1E10ZF889	
☐	09/02/2023 7:29 PM	Imaging paused	294	WDC WD800BEVT-75ZCT2 WDWXF1E10ZF889	
☐	09/02/2023 7:29 PM	Imaging started	294	WDC WD800BEVT-75ZCT2 WDWXF1E10ZF889	
☐	09/02/2023 7:23 PM	Case details changed	Not assigned	Unknown	
☐	09/02/2023 7:22 PM	Diagnostics completed	294	WDC WD800BEVT-75ZCT2 WDWXF1E10ZF889	
☐	09/02/2023 5:07 PM	Imaging completed	62	WDC WD3200JD-22KLB0 WDWCAMR1837670	
☐	09/02/2023 5:02 PM	Imaging started	62	WDC WD3200JD-22KLB0 WDWCAMR1837670	


The Reports page.

3. Via Home screen

Similarly, recent reports can also be found on the **Home** screen underneath the **Active** tasks.

On the **Home** screen you can look up active and completed tasks and view reports for all completed tasks.

Quickly find specific reports by entering filters in the search field.



Finding reports on the Home screen.

Printing reports from a case

When you work on an investigation and want to have complete information about the evidence drive and all operations that have been taken to diagnose, image, calculate hash, etc., you can address Atola TaskForce's case management system to print out all reports concerning your evidence.

To do that:

1. At the bottom of the case page, click **Print**.

The screenshot displays the Atola Technology web interface for Case ID: 294. The top navigation bar includes 'CASES' and 'REPORTS' tabs, and a 'DEVICES' section. The left sidebar contains icons for 'Diagnose', 'Image', 'Logical', 'RAID', 'Wipe', and 'Other'. The main content area is divided into three sections: Case ID: 294, Devices associated with case, and Reports. The Case ID section shows the last open date (09/07/2023 6:56 PM), investigator (Melanie Jones), creation date (09/02/2023 7:21 PM), and users (empty). The Devices section lists a WDC WD800BEVT...DWXF1E10ZF889 SATA drive with 80 GB capacity, 1101A11 firmware, 156 301 488 LBA sectors, and a security status of 'Not locked'. The Reports section shows a list of reports with columns for Date and Title. A 'PRINT' button is located at the bottom of the Reports section.

Case ID: 294

Last open date: 09/07/2023 6:56 PM Investigator: Melanie Jones
 Creation date: 09/02/2023 7:21 PM Users: (empty)
 Description:

Devices associated with case

WDC WD800BEVT...DWXF1E10ZF889 Firmware: 1101A11 LBA sectors: 156 301 488
 80 GB Security status: Not locked Logical sector size: 512 B
 SATA LBA48 mode: Supported Physical sector size: 512 B
 Device was imaged to 3 targets: [Show details](#)

Reports

Date ↓	Title
2023/09/07 18:56:41	Case details changed
2023/09/04 18:49:38	Diagnostics completed
2023/09/02 20:07:25	Imaging completed
2023/09/02 19:31:57	Imaging started
2023/09/02 18:21:10	Imaging started

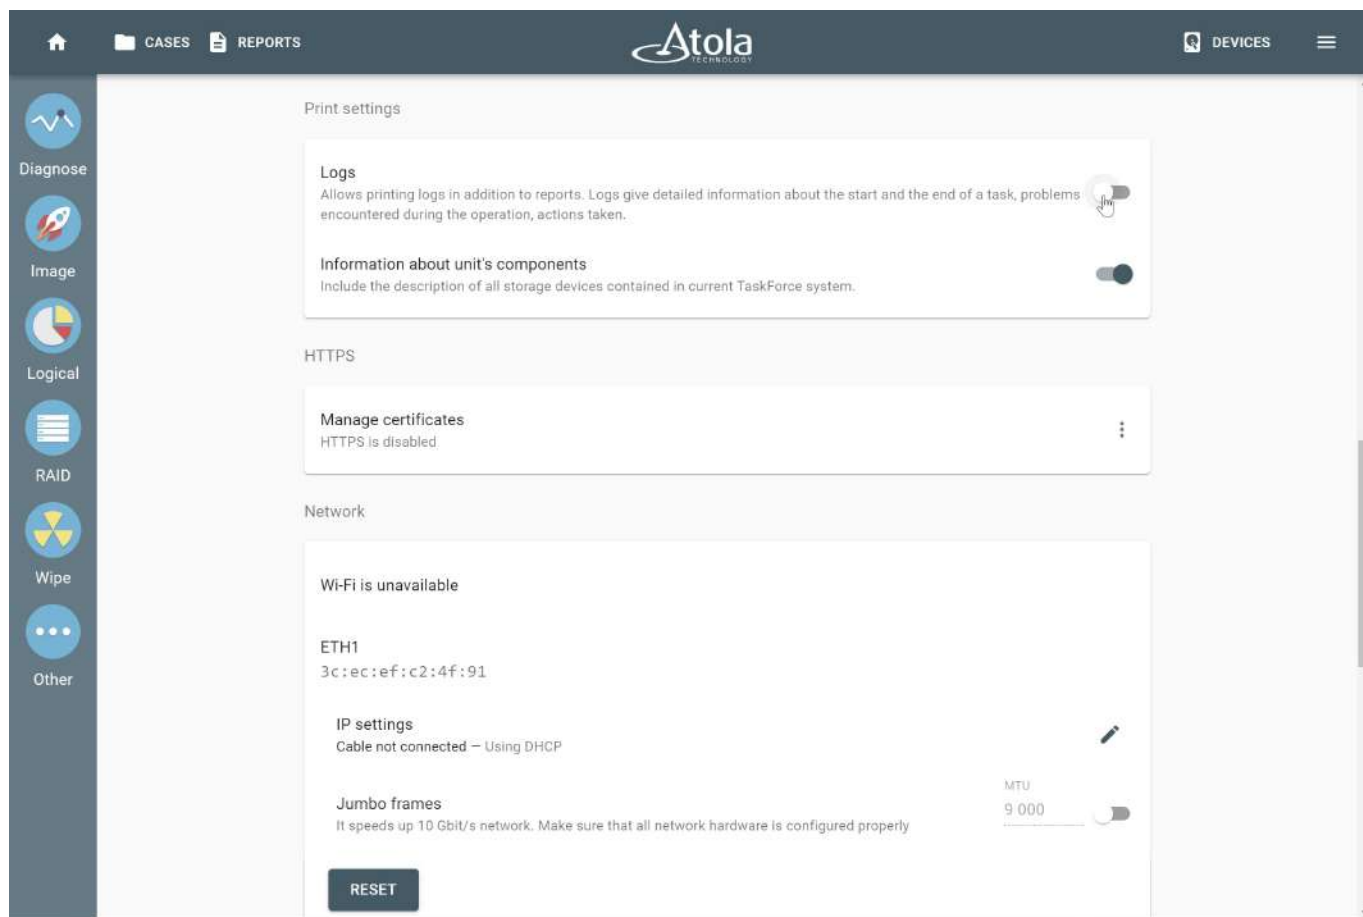
PRINT [You can change print setup in TaskForce settings](#)

The Print button on the bottom on the case page.

- Optional: If you want the printed reports to include logs, go to the the **Menu > Settings**, scroll down to the **Print settings** section and toggle **Logs**.

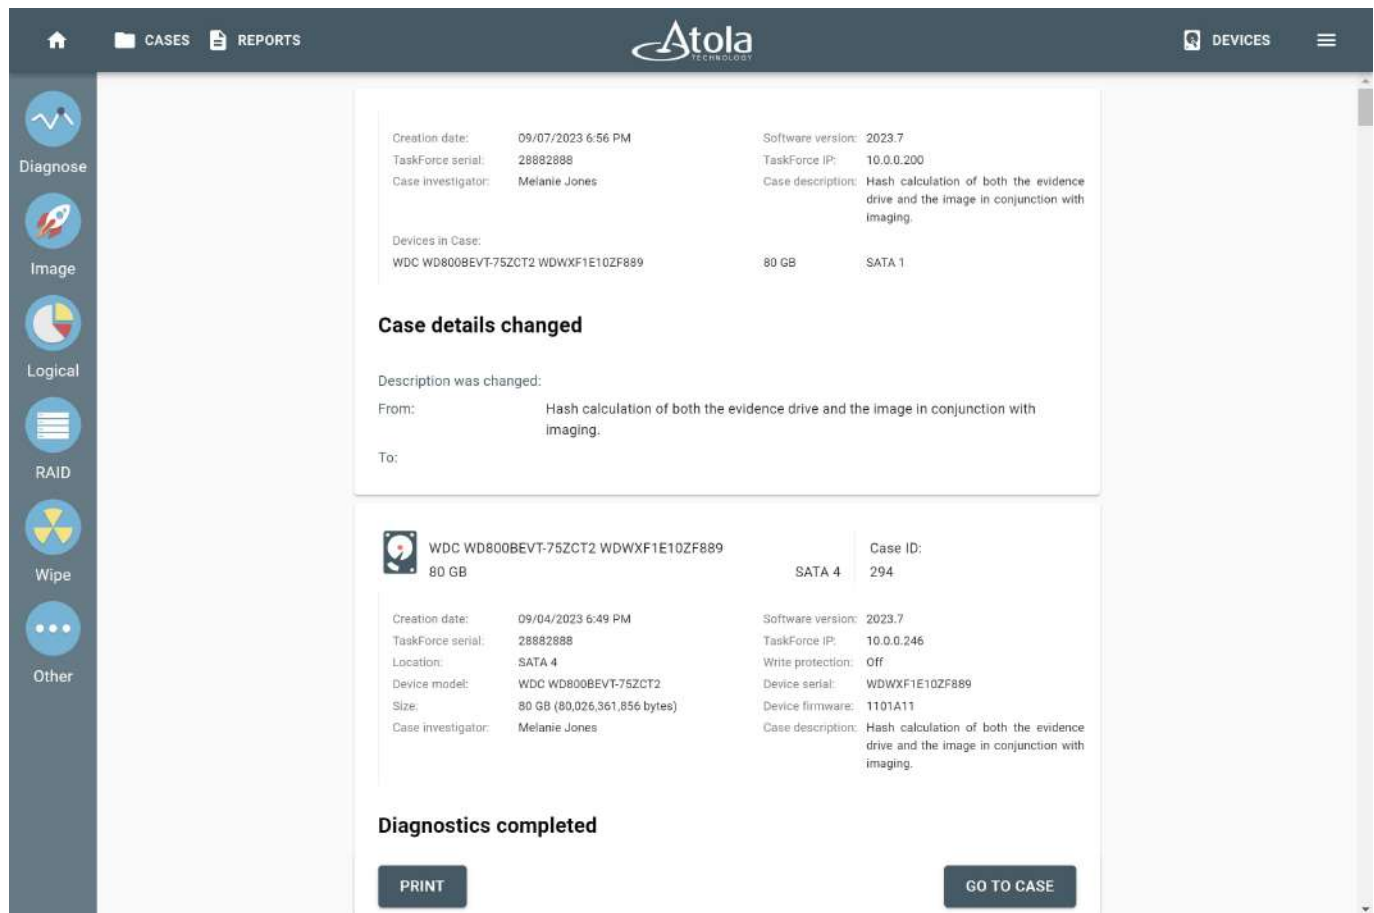
Logs are parts of the reports that give detailed information about the start and the end of a task, problems encountered during the operation (for example, inability to read a sector within the time allocated for the operation), actions taken (for example, jumps from a bad sector, completed pass of imaging, performed resets and power cycles etc.)

- Optional: If you want the printed reports to include the description of all storage devices contained in current TaskForce system, go to the the **Menu > Settings**, scroll down to the **Print settings** section and toggle **Information about unit's components**.



The Logs toggle in the Print settings section.

- After you click **Print** on the case page, TaskForce takes you to a page with full reports. There they are arranged in same the order, in which they were listed on the case page (either by date or by title).



If you selected the **Information about unit's components**, the last report generated will include the description of all storage devices contained in current TaskForce system.

CASES
 REPORTS
 DEVICES

Diagnose
 Image
 Logical
 RAID
 Wipe
 Other

List of storage devices inside of the TaskForce unit

Storage device	Description
HASP Key	Type: Persistent storage. HASP key stores 128 bytes of software licensing activation data. More specifically, it stores the TaskForce unit's serial number and the subscription expiration date. No data from hard drives (source or target) attached to the TaskForce unit is saved on this media.
Mass storage device: USB microSD card	Type: Persistent storage. TaskForce contains a microSD card in the rear System SD slot. It stores TaskForce's firmware. It is mounted in read-only mode. One sector (512 bytes) of this media is used to store TaskForce work folder path. This media is switched into the read-write mode when a firmware upgrade is performed (as initiated by the user). After the upgrade is complete, the media is instantly remounted back into read-only mode. No data from the hard drives (source or target) attached to the TaskForce is saved onto this media.
Dynamic memory device: DDR4 RAM	Type: Non-persistent storage. TaskForce unit contains a server-grade DDR4 RAM module that is used by its internal operating system to run program code. It is also used as a buffer to speed up data transfer rate. Data from the Source drive attached to the TaskForce unit is temporarily buffered onto this module before being written onto the Target. To wipe this storage, normal RAM wiping procedures can be applied. Powering off the unit and waiting for 15 minutes before re-applying power is considered to be a reliable wiping procedure under normal room temperature.
Other devices	This device utilizes a server-grade Supermicro motherboard which may contain other persistent storage, such as EEPROM (to store BIOS), NVRAM (to store BIOS settings) and other devices. None of these devices is accessed in any way by TaskForce firmware to read or save any data.

[PRINT](#)
[GO TO CASE](#)

On this page, there is another **Print** button. After clicking it you can configure printing settings.

FAQ and Troubleshooting

Imaging

How do I clone a drive and create multiple identical copies?

TaskForce is perfectly suited for cloning! Here is [how to clone into 4 SATA drives and 1 RAW file simultaneously](#).

When you need to create multiple copies, simply:

1. Go to **Image**.
2. Select a source device.
3. Select up to 5 targets (SATA, USB, SAS or RAW, E01 files) and click **Continue**.
4. Click **Start**.

What's more, TaskForce has 18 ports in total including IDE and Extension ports in combination with a huge overhead capacity to handle many imaging sessions or other operations concurrently. So if you start a 1-to-5 cloning, you still have at least 12 ports left for additional cloning/imaging sessions to run simultaneously. TaskForce is designed to [minimize the imaging time](#).

Is there a way to delete an imaging session and start imaging all over again?

When you select source and target drives previously used in an imaging session, TaskForce will indicate the progress status of the previous session, and you will be allowed to resume it.

To start a new imaging session with the same source and target, you will need to delete the previous one:

1. Go to **Image**.
2. Select the source drive.
3. On the **Select target devices** panel, select the target and click **Continue**.
4. Right-click the imaging session (or tap it with 2 fingers if you are using a touch screen).
5. In the context menu, click **Allow imaging map deletion**. A Trash icon appears in each imaging progress status tile.
6. To delete the previous imaging session, click the **Trash** icon. Now you can start a new session to the same target.

To disable imaging map deletion, click the imaging session again and select the option in the context menu.

Why does TaskForce not open the 'Imaging Settings' page when I click 'Continue' upon selecting targets?

The issue must have been caused by the browser cache.

1. Go to Chrome browser settings.
2. Open the **Clear browsing data** option.
3. Clear all the data after selecting the **All time** range.

How can I create an E01 image?

To image a source evidence drive to an E01 file, you have to create a new target file:

1. Go to **Image**.
2. Select the source evidence drive.
3. On the **Select target devices** panel, switch to the **Files** tab and click **Create file**.
4. In the file selector, find the folder to store the image and click **Create file**.
5. In the **Create image file** dialog, select the E01 file type.
6. Fill in the E01 file information, and then click **Create**.
7. On the **Select target devices** panel, click **Continue**.

How do I image to an E01 or a RAW file?

When selecting targets in the slide-out panel, click the **Select file** tile. It only works if TaskForce is connected to a network and with an accessible network folder.

How do I set up an automated launch of AXIOM after imaging?

You cannot launch third-party software from TaskForce itself. However, with Magnet AUTOMATE, TaskForce can be integrated into a workflow with Magnet AXIOM. The Atola team closely cooperates with Magnet's developers to support and enhance this integration. [See how it works](#).

Another option is a folder monitor/watchdog. Use an app that tracks when a new image file appears in a specific folder. After that, such an app can launch Magnet AXIOM against the newly created image. There are many watchdog tools on the web.

Why does a healthy drive appear to have many bad sectors during imaging?

The issue is likely caused by a loose port connection, leading to read errors and even drive identification issues.

To completely reset the physical link, unplug the cable from both the port and the storage device, wait for a full minute and plug it snugly back in.

It could also be cable wear-out. Swap cables between the "working" and "non-working" ports and see how it goes. If certain cables repeatedly cause issues, put them aside.

How to properly connect a spinning 3.5" SATA drive in a USB enclosure to TaskForce?

All USB ports of Taskforce are designed to work with USB devices conforming to the USB specification. However, some USB-enclosed HDDs consume too much current at spinup, in fact violating the specification.

Please note that TaskForce works properly only with the Chrome browser.

How do I increase the speed when I image to the network?

When imaging to the network, there are two potential bottlenecks:

1. Network performance

It can be enhanced with a 10Gb connection. Another important thing: [Jumbo frames should be enabled](#) in the TaskForce settings, on the target computer/NAS network card, and on all network switches/routers in between.

2. Write speed of the target network drive

It can be more difficult to improve, especially if you are imaging several drives to the same network location. This leads to a situation where 8 source SSDs are read at 500 MB/s speed, but the total writing capability of the target network drive limits the speed to just 400 MB/s. Distributed between 8 sessions, the speed becomes too low ($400/8=50$ MB/s per session).

The best solutions to achieve top speeds with target network locations:

- A RAID consisting of SSDs.
- A NAS combined with a 10Gb switch. For example, [Ubiquiti EdgeSwitch 16 XG](#).
- A network server with many drives and great writing performance.

Why does imaging to a NAS end up with a write error?

Possible causes and solutions:

- Maximum file size limitation by the proprietary NAS file system or OS.
- Minimum free space limitation in the NAS setting.
- Network folder permissions.
- Replace the Ethernet cable. Use Cat6 (with a length of less than 55 meters), Cat6a, Cat7, Cat7a, or Cat8 marking on it.

How do I restore a 520-byte sector-sized image to a drive?

TaskForce performs imaging and restores images in sector-to-sector mode. In the process, TaskForce does not convert logical data alignment depending on sector size, which keeps things simple and explicable.

If you have an image of a drive with a sector size of 520 bytes and you want to restore it to another drive, you must use a target SAS drive with a sector size of 520 bytes.

If the SAS drive's sector size is other than 520, you can reformat it in Linux using the [sg_format utility](https://linux.die.net/man/8/sg_format): https://linux.die.net/man/8/sg_format

Then perform imaging to the target SAS drive by selecting the image file as a source.

For such drives, we have equipped USB 3 and USB 4 ports with extra electric load capacity. Try connecting the drive to USB 3 or USB 4 ports.

How can I image a U.2 NVMe drive?

To use a drive with U.2 interface, attach the drive to TaskForce with the help of a U.2-to-M.2 adapter and cable:

- [U.2-to-M.2 adapter](#)
- [Cable for the adapter](#)

Connecting to TaskForce:

1. Plug the U.2 drive into the U.2-to-M.2 adapter using the cable.
2. Plug the U.2-to-M.2 adapter into the [Atola M.2 extension](#).
3. Plug the extension into TaskForce while the unit is powered off.

How can I image a SCSI drive?

Take advantage of USB ports and SCSI-to-USB adapters. There are a number of these on the market:

- Microtech XpressSCSI USB to SCSI adapter
- Belkin USB to SCSI Adapter
- Shuttle eUSCSI Bridge
- Castlewood ORB SCSI to USB adapter

Why does an M.2 NVMe drive not get identified?

Please note that the support of drive hotplug for M.2 PCIe NVMe drives was added in the firmware version 2021.4.

If you have an earlier firmware version, to replace an NVMe drive, TaskForce needs to be powered off:

1. Power off TaskForce. Remove the power cable and wait for 1 minute.
2. Unplug the M.2 extension with the M.2 drive installed and plug it carefully back in.
3. Power on the TaskForce unit.

What is Storage mode?

TaskForce treats a target drive as a destination for sector-to-sector imaging unless you set it to Storage mode.

When the target drive is in Storage mode, it can only be used as a container for image files (E01, AFF4, or Raw). The software protects you from accidentally using a Storage drive for sector-to-sector imaging.

Storage mode features:

- You must format a target drive into Storage mode using TaskForce beforehand.
- When formatting in Storage mode, TaskForce creates a single exFAT partition with a large cluster size for the best imaging performance.
- TaskForce will automatically mount the target device formatted as a Storage when you plug it into TaskForce the next time.
- It's safe and works simpler than it seems. Format once and use it many times.

On top of that, you can select a VeraCrypt container instead of a simple exFAT partition to store image files in an encrypted form.

RAID configuration autodetection & imaging

Can TaskForce identify a RAID's configuration if I include drives that are not RAID members (the drives are not marked properly)?

Normally, TaskForce can identify the type of an unknown RAID within a minute. TaskForce reads data from the initial 3 million sectors of each drive and searches for the correct relation between the RAID members.

With an odd drive, the module will need more time to identify the configuration.

- In the case of RAID 1 or 10, detected mirrors will be arranged into groups, the odd drive will be placed separately, and partitions will be mounted successfully.
- If it is a JBOD, the partitions will likely be identified, and the order of the drives will be correct, with the odd drive placed at the end.
- As for RAID 0 or 5, TaskForce won't reassemble such an array: data from the odd drive is taken into account when TaskForce combines data from all members.

NB: The autodetection module attempts to identify the configuration from scratch each time you remove or add a drive. Try removing a drive that seems odd. Normally, RAID members are drives of the same capacity and usually of the same type, made by the same manufacturer, etc.

How do I manually reassemble a RAID knowing configuration parameters?

If you are sure you know all the RAID parameters, you can rebuild the RAID manually, ignoring the autodetection process.

This is the only way to rebuild and image a RAID array when the TaskForce RAID module does not recognize partitions because they are either encrypted or unsupported (for example, ReiserFS).

To rebuild a RAID manually, follow the steps below:

1. Select all RAID members (drives or their images).
2. Specify the following settings manually from the drop-down menus:
 - a. RAID type (or RAID level)
 - b. Start LBA (it is an offset for all RAID members)
 - c. Block size
 - d. Block order (only for RAID 5 and RAID 6)
 - e. Device order (use the two-line button to drag and drop)
3. Click the **Go to Image** button.

Can I select E01 images of drives for RAID reassembly?

You can select any combination of image files and/or drives constituting a RAID. See examples:

- 3 drives + 2 AFF4 files
- 5 E01 image files
- 2 raw files + 1 missing member

Can I image a RAID array that contains a damaged drive?

The RAID assembly module requires good, readable drives to find a configuration. When one of the RAID's drives is damaged, it's best to image the drive first and attempt the reassembly using the image.

However, if the configuration is known and no search is required, you can try to mount the RAID even with a malfunctioning device. In this case, specify the configuration manually and proceed with imaging. TaskForce's [multipass imaging engine](#) handles RAID arrays with such issues effectively.

If you are dealing with a RAID 5, RAID 6, or RAID 1, TaskForce will use the RAID's redundancy when encountering internal read errors.

When one of the drives in a RAID 5 array is seriously damaged, you can add a virtual **Missing device** instead of the damaged device. Next, it is possible to reassemble and image the RAID 5 array from the remaining healthy drives and 1 "stub" device. TaskForce reconstructs data from such arrays on the fly during imaging.

Write error at the 16TB threshold occurs when imaging an assembled RAID. What should I do?

Check the file system of your target folder for maximum file limits. For example, these file systems are limited to the 16TB threshold:

- ext4
- NTFS
- Btrfs in some Synology NAS

To handle larger files in NTFS, you can reformat an NTFS partition to use a larger cluster size. By increasing the cluster size to 64 kilobytes (KB), you can create files up to 256 TB size-wise.

Performance

How to reach 15 TB/h?

TaskForce was designed as an ultimate multitasking tool with a huge overhead capacity. Given that most imaged devices are HDDs, whose data transfer rate is 200-220 MB/s at best, it is not frequently that TaskForce will reach the limits of its capacity.

Our QA team achieves a speed of 15 TB/hour with this specific setup:

- imaging 6 SAS ports to 6 SATA ports (all SSD drives)
- imaging 2 USB ports to 2 USB ports (all SSD drives)
- imaging the NVMe extension port to a network file (10Gb connection)
- imaging the IDE port (with an IDE drive attached) to a network file (10Gb connection)

This adds up to 10 imaging sessions at superb overall performance.

However, in most cases, you deal with HDDs as source devices, and they are half as fast as SSDs, which creates a bottleneck.

Another potential performance bottleneck is network bandwidth. To achieve maximum throughput, connect the ETH1 port to a network or NAS, and the ETH2 port to another network or NAS. It will give you 20Gb/s of throughput. It also depends on your network hardware and setup.

To learn about network throughput optimization, read the official TaskForce manual:

- [Jumbo frames for fast imaging to server](#)
- [Network setup](#)

Performance when imaging to a remote USB drive (write cache + exFAT)

There are two ways to boost imaging performance:

1. Format the target drive to exFAT. It works faster than NTFS.
2. Enable Windows write cache:
 - a. Go to **Device Manager**.
 - b. In the **Disk Drives** category, select your drive.
 - c. On the **Policies** tab, switch to **Better performance**.
 - d. Confirm the reboot of the PC.

Important: When the last imaging session completes, use **Eject Removable media** icon in the system tray to guarantee cache-flushing of the last portions of written data. Here is [an article](#) explaining how it works.

Connectivity and environment setup

How to connect TaskForce to a fiber optic network?

Option 1. Use a 10GBASE-T RJ45 Copper Ethernet to 10G SFP+ Fiber Media Converter.

This device converts signals from your TaskForce's 10 Gbit Ethernet ports to a fiber optic SFP+ connection, enabling high-speed network capabilities.

Option 2. Connect through a switch with SFP+ ports.

Integrating a network switch that has SFP+ ports can bridge TaskForce with fiber optic networks. Ensure that the switch supports 10 Gbit speeds to match TaskForce's capabilities.

Important considerations:

- Ensure that the SFP+ module is compatible with your switch brand (e.g., Cisco, Juniper).
- Use cables shorter than 55 meters with SFP+ modules; otherwise, these modules might not function properly.

How do I setup Synology DS218?

1. Go to **Control panel** > **File services** > **SMB** > **Advanced settings**.
2. Set **Maximum protocol** to **SMB3**.
3. Go to **Control panel** > **Shared folder**.
4. Click the **Create** button and specify network folder details.

If you need to get a guest account working, perform the following actions:

1. Go to **Control panel** > **User**.
2. Edit the **Guest** user.
3. Clear the **Disable this account** checkbox.

How do I set up my network shared folder?

Follow these steps to setup Windows network folder:

1. Go to **Settings** > **Network & Internet** > **Status** > **Sharing Options**.
2. Open **Private (current profile)** section.

How do I extend my server network with a DHCP-enabled switch?

Here is an example of a switch that supports static IP setup via simple web admin. To set the IP addresses for each current server network node:

[Ubiquiti EdgeSwitch 16 XG](#)

Four 10Gb Ethernet ports, twelve 10Gb SFP ports

To configure a Ubiquiti DHCP server:

1. Connect PC and TaskForce to Ubiquiti switch.
2. Set static IP address of PC to *192.168.1.4*.
3. Open a browser and enter *192.168.1.2* (default Ubiquiti switch IP).
4. Log in with default credentials: *ubnt* (both name and password).
5. Go to **System > Advanced Configuration > DHCP server > Global**.
6. Activate **Admin mode** by checking the required checkbox and clicking the **Submit** button.
7. Go to **Pool Summary** and click **Add** to make a new address pool.
8. Enter your:
 - pool name
 - network base address (192.168.1.0, for example)
 - network mask (255.255.255.0)
 - put **Default Router Address** and **DNS**

After creating your pool, you can change it via **Pool configuration** tab.

9. Click the **Save configuration** button in the upper right corner of the window and then **Save**.

Alternatively, there is a [tutorial on Youtube](#).

To enable 10Gb with jumbo frames:

1. Go to **Basic > Port summary**.
2. Select ports *0/13*, *0/14*, *0/15* and *0/16* and click **Edit**.
3. In the **Edit Port configuration** window, change **Maximum Frame Size** to *9014*.

Why does an IP address not show despite an Ethernet cable being plugged into the unit?

Usually, your network router is responsible for assigning IP addresses to other computers or devices in the same network.

Make sure your router has DHCP support and it is enabled.

There is no DHCP server on our internal lab network. How do I assign a static IP address to TaskForce 2?

1. In TaskForce, click the menu button in the top right corner.
2. Click **Settings**.
3. In the **Network** section, find IP settings for both 10Gb Ethernet cards.
4. Click the pencil button to see **Use static IP** checkbox.

3. Select **Turn on network discovery**.
4. Select **Turn on file and printer discovery**.
5. Open **Guest** or **Public** section.
6. Repeat steps 3 and 4 for **Guest** or **Public** section.
7. Open **All Networks** section.
8. Select **Turn on sharing** so anyone with network access can read and write files in the **Public** folders.
9. Select **Turn off password protected sharing**.
10. Click **Save changes**.

If the above has not fixed the issue, edit the folder share permissions:

1. Open the **Computer Management** window (press **Win** button and type in **Computer Management**).
2. Expand **System Tools > Shared Folders** and click **Shares** in the left-side tree.
3. Right-click the shared folder in the central pane and click **Properties**.
4. Select the **Share Permissions** tab.
5. Click **Add** to assign permissions to the shared folder for a user group.
6. In the dialog box, type *Everyone* and click **OK**.
7. Select **Full control** permission (**Read & Write** permissions) for the user group you have just added.
8. Click **OK**.

We strongly recommend using the latest OS versions. It is crucial for reaching high network transfer speeds.

How do I set up the Windows Server share?

In **Control** panel:

1. Enable the **Guest** account (**Administrative tools > Add users and computers > Users**).
2. Next, go to **Network and sharing center** and choose **Change advanced sharing settings**.
3. Select **Turn On network discovery** and **Turn on sharing (file and printers + public folders)**.
4. In the shared folder access options, add **Guest** or **Everyone**.

If the shared folder demands restricted access, please follow [this guide](#).

TaskForce booting does not get completed. What should I do?

To ensure TaskForce boots correctly, reset the device:

1. Power off TaskForce.
2. Detach all devices and cables from the system (including PSU cable, Extension module, SATA cables, USB devices/cables, etc.). So the TaskForce should be just by itself with nothing attached to it at all.
3. Give it 3-5 minutes to fully reset. There are a few internal circuits that need up to a minute to fully reset after powering off, but we recommend waiting at least 5 minutes to be sure.
4. Plug only the power cable back in (no network/USB/SATA cables, etc.).

Atola TaskForce demo video includes a basic explanation of the network settings. Here's a [Youtube video](#) on that.

How do I connect TaskForce to a network domain?

TaskForce does not require a preliminary setup for working in a domain. Instead, your flow would be:

1. In the left-side taskbar, click the **Image** icon.
2. In the left slideout panel, select the source.
3. To select the target, in the **File** section, click **Select file**.
4. In the file manager's **Local network** section, click the **Connect** button.
5. A new dialog appears with the **Domain** field.

The entered server will remain in the **Local Network** section list.

Additionally, in [Web API](#), there's a *targetFolderDomain* parameter of */start-image* command that may be of use.

5. Power TaskForce on and wait for 3-4 minutes.

How to see the syslog for more information about the issue I am facing?

To download the TaskForce logs:

1. Connect TaskForce to the network so that you can access it from your PC.
2. An IP address is shown on the IP screen on the front panel of the unit. Let's say, it is *10.0.0.33*.
3. On your PC, open the Chrome browser and enter *http://10.0.0.33/syslog*.
4. Click **Save**.
5. This will download the logs to your PC's *Downloads* folder.

Tablet screen

The Tablet UI has no connection or shows a blank page. How to fix this?

The issue might be caused by the browser cache. Try clearing the cache:

1. Exit full-screen mode by going to **TaskForce Menu > Toggle full screen**.
2. Go to Chrome browser settings.
3. Open the **Clear Browsing Data** option.
4. Clear all data by selecting the **All time** range.

Please note that TaskForce works fine with Chrome browser only.

If this doesn't help, follow the steps below:

1. Leave full-screen mode by going to **TaskForce Menu > Toggle full screen**.
2. On your TaskForce page in Chrome, press **F12**. This will bring up the Chrome Dev Tools panel.
3. Go to the **Applications** tab, navigate to **Local Storage** and right-click your TaskForce IP.
4. Press the **Delete** button.
5. Press **F12** again to close the panel.
6. Press **Ctrl+F5** to refresh the page.

If the problem persists, reinstall the Chrome browser.

Why when I try to turn on the screen, a battery indicator starts blinking?

The screen battery needs to be charged. Let the unit stay powered on for 5–10 minutes and avoid turning the screen on during this time. 10 minutes later, use the button on the screen to turn it on as usual.

How do I connect a mouse and a keyboard to the Surface tablet?

1. Exit full-screen mode by tapping the menu in the top right corner and selecting **Toggle Full Screen**.
2. Enable Bluetooth in the tablet's Windows 10.
3. Connect the mouse and keyboard.

Wiping

How does SSD Trim work and does it wipe a drive completely?

SSD Trim doesn't instantly wipe sectors (NAND memory cells) of a drive. It instructs SSD's firmware which sectors can be wiped by marking them as 'dirty'.

The time of erasure of 'dirty' sectors depends on the SSD manufacturer and firmware. For instance, recent Samsung SSDs have a so-called foreground garbage collection. It wipes any erased file almost immediately thanks to a TRIM command proactively executed by the operating system. In older SSDs, trimmed sectors can remain intact for minutes or even hours.

The most secure way to erase an SSD entirely is using one of the following methods:

- **Secure Erase** - for SATA drives
- **Format NVM** or **Sanitize** - for NVMe drives

The drive's internal implementation of these commands is vendor-specific. In most drives, it ensures full erasure of an SSD including non-addressable areas.

Case management & reports

How can I export reports in Standalone mode?

You can use an external USB drive for it. To proceed, it must be formatted as a TaskForce-compatible Storage device.

Here is how you do it:

1. Plug the USB drive into one of the USB ports.
2. Switch the USB port into target mode (turn off its Source switch).
3. Click **Devices** and then click the **File** tile.
4. In the **Select image file** window, click the **Add storage** button.
5. In the **Select device** panel, click the drive you want to use in the Storage mode.
6. Choose **Format device to exFAT** and click **Next**.
7. Click **Format** and confirm the action.

Once the USB device has been formatted as a Storage, you can use it as a target device to save reports on in Standalone mode:

1. Plug the USB Storage device into any USB port in target mode.
2. Go to **Cases**.
3. Search for the cases you need. Use the search field and sortable table to locate the cases.
4. Select the cases and click the **Export** button.
5. Select **Storage plugged into TaskForce** and click the **Export** button.

Can reports only be exported in PDF, can I export them in RTF format instead?

It should be possible to save a page into RTF. For example, in this [Chrome AddOn](#), you can select the reports and click Print. It will generate a single page with all selected reports, which can be exported to .RTF file.

Can all reports listed on the main page be sorted, managed, or deleted?

Yes, they can be filtered, sorted, exported, and deleted. The quickest way to perform a search is by using the **Search** field above the reports on the main page. Entering the task name (**Image**, **Diagnose**, **Wipe**) and/or device model/serial helps get a more specific output.

To sort, export, or delete reports, you can go to **Reports** in the top menu, where all case reports are shown in a table view. **Save to** helps download the selected reports in a single ZIP file. There is also a **Delete** button in the bottom right corner for your needs.

Hardware

How do I remove or swap out the internal storage from TaskForce for security reasons?

We have foreseen that such a requirement may be imposed by some organizations. That's why TaskForce allows the following:

- Removable internal SSD: locate the hatch (behind the removable touchscreen) in the lid, unscrew it to remove the SSD.
- Units assembled in 2020 onwards: the Wi-Fi module is no longer built in.
- Built-in Wi-Fi (in units assembled in 2018 - 2019) is disabled by default and can be removed if necessary (through the hatch in the lid).
- The tablet's camera and mic are disabled during the assembly.
- The touchscreen is detachable (the imager can be operated without a screen).

How can I protect my TaskForce from physical damage during transportation?

From September 2020 onwards, all newly assembled units are equipped with a [magnetic screen cover](#) that is included in the standard package with the product.

This cover helps protect the unit's screen, it is easily attached and removed when necessary. If you need a pelican case to pack the unit with all the cables and extensions, the suitable case model is [IM-2700](#). For instructions on how to obtain a fitted foam, please contact [our support](#).

Where can I purchase additional SAS and SATA cables?

eSATAp cables similar to the ones included in the package with TaskForce can be ordered on [Amazon](#) or a different online shopping site (Ebay, AliExpress, etc.) at around \$9 each. The SAS/SATA powered cables are not widely available on the market - this custom cable is made per our order.

Is there a battery to keep my TaskForce powered when there is no electric grid available?

These Cyberpower UPS systems can serve as a reliable battery for Atola imagers, TaskForce and Insight alike:

- [CyberPower PR750RT2U Smart App Sinewave UPS 750VA/750W](#)
- [CyberPower OR700LCDRM1U Smart App LCD UPS 700VA/400W](#)

The flat horizontal UPS design is convenient as it allows putting an Atola imager on top of it. The price range of CyberPower UPS systems is quite vast, and you can choose a suitable model with power ranging from 500VA/300W to 2000VA/2000W.

Subscription and firmware update

How do I activate my unit in a network-free environment?

There is a single way to perform TaskForce activation or subscription extension:

1. Enter the TaskForce serial number. It can be found at the bottom of the unit.
2. TaskForce generates a license key (Internet connection is not required).
3. You find an Internet-connected PC and visit the website: [a.atola.com](#).
4. You enter the license key and several other details.
5. The website generates a TaskForce activation code.
6. You write the code down or take a photo, and then enter it in the TaskForce activation screen.

Why does Internet Explorer rename TaskForce's .AFM firmware file to a ZIP file?

Internet Explorer identifies an AFM file as a ZIP file and automatically renames it to ZIP. Read [this article](#) about the reason Explorer does this. This issue only occurs in the Internet Explorer.

We suggest using other popular browsers, such as Chrome or Firefox.