

SOP-003 — OCT Export and Imaging Data Handling

Document	SOP-003
Title	OCT Export, Imaging Data Handling and Custody
Version	1.1 — awaiting Medical Director signature
Drafted	2026-08-18 · v1.1 2026-08-20: site references made generic
Owner	Medical Director, SpectOcular LLC
Effective	Not effective until signed by the Medical Director (\$12)
Review cycle	Annually, or on any change to device fleet, export tool version, regulatory tier, or storage architecture
Supersedes	None — first controlled imaging-export procedure
Applies to	DICOM OCT Export Tool v2.1.2.14 (Topcon), User Instructions Rev E

Read before signing. The five decisions behind this procedure were reviewed and settled on 18 August 2026 and are recorded in **Appendix C** with their reasoning. The implementation item they left open — selection and validation of the Stage 2 de-identification tool (\$7.2) — was closed on 18 August 2026 and is recorded in **VAL-001**. One confirmation against real device output remains outstanding, and until it is done no imaging leaves the covered Workspace account (\$7.2). \$5.5 (automated export) and \$6.5 (follow-up visits) were added on 18 August 2026 and are recorded in Appendix C. Nothing here is legal, regulatory or billing advice.

1. PURPOSE

To define how OCT and OCT-A examinations are exported from Topcon devices, where the resulting data is stored, how it is coded, how custody is recorded, and how long each copy is kept — so that any image used in a read or a model can be traced back to the exam that produced it.

The absence of this procedure is the gap this document closes. As of drafting, exports are produced by a tool whose output content changes with the licence installed on the machine, and there is no written rule saying which licence to use, where the output goes, or how it is verified.

2. SCOPE — AND THE DISTINCTION THAT GOVERNS EVERYTHING

The same .fda file can produce two materially different DICOM outputs depending on the licence installed. They are **not** interchangeable.

	Clinical read	Research grade / model development
Export tier	US (no licence) — <i>not deployed, see §4</i>	Research — what SpectOcular runs
Retinal layers	5 — ILM, RNFL/GCL, IPL/INL, OS/RPE, BM	10 — adds GCL/IPL, OPL, ELM, IS/OS, CSI
OCTA metrics (vessel density, FAZ)	Not available	Available — <i>but see §3.1: not applicable to the current fleet</i>
Segmentation SOP Instance UID	...7.1	...7.3
Regulatory status	FDA-cleared subset	Not FDA-cleared

A Research-tier export must never be used to produce a clinical read, and must never be filed in a patient's clinical record. The additional layers and the OCTA metrics are outside the FDA clearance. Using them for care is a regulatory problem, not a preference. This mirrors SOP-001 §2: the same images and the same people, two activities that must not share a workflow.

The tier is verifiable after the fact — see §8.3. There is no need to trust anyone's memory about which licence was active.

3. WHAT THE TOOL PRODUCES

Input is a .fda file from a Topcon device (3D OCT-1, Maestro2, DRI OCT Triton). Output is written to a directory of your choosing.

Principal DICOM objects:

Object	Filename suffix	Contents
OP	...2.1.dcm	Fundus photograph
OPT	...1.1.dcm	B-scan tomography
OPTBSV	...3.1.dcm	Unprocessed OCTA flow volume
OPTENF	...6.y.dcm	En-face slabs
SEG (HMSEG)	...7.y.dcm	Layer segmentation, heightmap — y encodes the tier

Derived, non-DICOM outputs used for analysis: oct.json, layer_001...010.json, grid.json and faz.json (OCTA metrics), thickness binaries (Retina.dat, RNFL.dat, GCL+.dat, GCL++.dat), and rendered images (bscan_*.jpg, color_001.jpg, ir_001.jpg, shadowgram.jpg, projection.jpg).

3.1 The current fleet is structural only — read this before §5

The current pilot fleet operates **Topcon colour fundus and structural OCT**. There is no OCT-A.

Consequences, so nobody spends time chasing outputs that cannot exist:

- **Expected objects:** OP (fundus), OPT (B-scan), SEG (layer segmentation)
- **Will not be produced:** OPTBSV (...3.1.dcm, OCTA flow volume), OPTENF (...6.y.dcm, en-face slabs)
- **Will not be produced:** grid.json, densityMapGrid*.png, faz.json, fazEnface.png — the OCTA metrics in Appendix B
- Their absence is **correct behaviour on a structural scan, not a failed export**

The -octa and -enfaceSlabs switches inside -allDcm are inert on a non-angio .fda, so the standard parameter set in §5.2 needs no change. Leaving them in means the same command keeps working unchanged if an angio-capable device is added later.

This does not weaken the case for the Research tier. On a structural-only fleet the tier still governs **layer segmentation — 10 layers versus 5** — which is the scientific payload for retinal biomarker work. GCL/IPL, OPL, ELM, IS/OS and CSI are available only above the US tier and are exactly the boundaries a layer-thickness model depends on. §2 and §4 stand as written.

Record the exact Topcon model and software version for each site in the equipment log. Model determines which outputs are possible at all (Appendix B), and it is not safe to infer it from the fact that a scan succeeded.

3.2 Device scope — this SOP covers Topcon only

The export tool reads .fda, Topcon's proprietary format. **This procedure therefore applies to Topcon devices only.** Zeiss, Heidelberg, Nidek and Optovue each export DICOM by their own route and are outside its scope.

Two consequences:

1. **Procedural.** Admitting a non-Topcon device requires its own export procedure — an amendment to this SOP or a sibling document — before any imaging from it enters the corpus.
2. **Scientific, and the more important of the two.** Layer segmentation algorithms differ between manufacturers, so thickness values are **not directly comparable across devices**. Pooling them into one biomarker dataset without modelling the device effect is a recognised confound. While N is small, a homogeneous fleet is worth more than a larger heterogeneous one.

Device eligibility is governed by the qualified device list referenced in SOP-001 §5. A device not on that list is quarantined, not exported.

4. REGULATORY TIER AND LICENSING

SpectOcular runs the Research tier only. There is no clinical installation of this tool.

Clinical imaging work is done in IMAGENet 6, Topcon's clinical software. This utility exists in the SpectOcular workflow solely to produce study data, so no site has any reason to hold a US-tier installation.

Installation	Licence	Where
Research	Research tier, licence key from Topcon	Enrolling sites and SpectOcular
Clinical (US, no licence)	—	Not deployed. Add only if a documented need for clinical DICOM from this tool arises, and record it here first

A site holding only one installation cannot select the wrong one. The §8.3 check then confirms intent rather than catching a mistake.

The licence is a 32-byte hex key in DicomOCTExport.exe.license, tied to the machine's Hardware ID. Obtain it by running the tool once with no arguments to display the Hardware ID, then requesting the key from your Topcon representative.

The consequence of the Research tier is that its output is not FDA-cleared (§2). No export from this tool may be filed in a patient's clinical record or used for a care decision. If a clinician needs imaging for care, it comes from IMAGENet 6.

5. EXPORT PROCEDURE

Performed either by the **SpectOcular site agent** (§5.5, the normal case once a site is provisioned) or by trained site staff. The manual procedure below is what the agent automates, and remains the documented method for any site where the agent is not installed.

5.1 Before exporting

3. Confirm the subject is enrolled and has a **Subject ID** (format NN-XXXXX) recorded on the signed consent packet.
4. Confirm patient demographics on the device match the clinic record. **Demographics edited in IMAGENet 6 after acquisition do not propagate back into the .fda file** (Tool Instructions, Appendix B). If they were corrected downstream, the export will carry the stale values unless overridden per §7.2.
5. Confirm which installation you are running — clinical or research (§4).

5.2 Running the export

Either interface is acceptable. The GUI (ProcessFDABatchFiles.exe) is preferred for routine work; the command line is preferred for batch and for reproducibility.

Standard research export — default parameter set:

```
DicomOctExport.exe "<path>\<exam>.fda" "<output_dir>" -allDcm -json -debug
```

- -allDcm — equivalent to -dcm -segDcm -octa -enfaceSlabs; produces OP, OPT, OPTBSV, OPTENF and SEG
- -thicknessMap — retinal thickness binaries (Retina.dat, RNFL.dat, GCL+.dat, GCL++.dat). The study's endpoints are layer-thickness biomarkers, so these are captured directly rather than derived from the segmentation object afterwards
- -json — layer and scan metadata used for analysis
- -debug — writes a detailed log to the output directory; this is the export's audit record and is retained

Verify -thicknessMap on the first real scan. The tool's parameter compatibility matrix (Tool Instructions §3) lists exclusions that could not be read reliably from the published PDF. If it proves incompatible with -allDcm, run it as a second pass into the same output folder rather than dropping it:

```
DicomOctExport.exe "<path>\<exam>.fda" "<output_dir>" -thicknessMap -debug
```

Record which form was used in the equipment log, and amend this SOP to match.

In the GUI, tick **New Folder for Each FDA File**. Without it, batch output from multiple exams is written into one directory and cannot be separated afterwards.

Do not use -HarmonyExport true unless the Glaucoma Module specifically requires it; it substantially increases output volume.

OCTA metrics (-densityMapGrid, -fazData) emit only for the device, scan size and resolution combinations in **Appendix B**. Their absence for other scans is expected behaviour, not a failure.

5.3 Output folder naming

Name the output folder with the **Subject ID and the exam date only**:

03-HN7JZ_2026-08-18

Never put patient name, MRN or date of birth in a folder or file name. Folder names propagate into backups, sync logs and screenshots, none of which are covered by the controls applied to the contents.

A subject scanned twice in one day — both eyes acquired as separate exams, or a repeat after a blink artefact — produces two identically named bundles, and two same-named files in one Drive folder cannot be told apart afterwards. A six-character **scan discriminator** is therefore appended:

03-HN7JZ_2026-08-18_a1b2c3.zip

The discriminator is a digest of the exam's SOP Instance UIDs. It carries no patient information and is deterministic, so re-exporting the same exam yields the same name and overwrites rather than duplicates. The date stays ahead of it so a plain alphabetical sort of a subject's imaging folder is a visit timeline.

5.4 What may be exported and uploaded — bounded by consent and by equipment

Two independent limits apply to every upload. Neither is optional and neither substitutes for the other.

1. The subject's consent. A subject consents to SpectOcular Patient Services, and separately and optionally to Retina IQ. Those are different permissions.

Subject consented to	May be exported and uploaded	May be routed to Retina IQ
Research only	Yes — study corpus	No
Research + Retina IQ	Yes	Yes
Declined	Nothing. Do not export	No

The enrolment record carries RetinalQRequested (Yes/No) against the subject. **That field, not the operator's recollection, governs routing.** If it reads No, the imaging supports the study and nothing else, however clinically interesting the scan.

2. The equipment at the site. A site uploads what its devices actually produce, and nothing is expected of it beyond that. A structural-only site (§3.1) uploads OP, OPT and SEG; the absence of OCTA objects is not an incomplete submission and must not be queried as one. Site capability is recorded in the equipment log and is the reference for what a complete upload looks like *at that site*.

Where the two limits disagree, the narrower governs.

5.5 Automated export and upload — the SpectOcular site agent

Where installed, the export and upload steps in §5.2–§5.3 and §6.1 are performed unattended by the **SpectOcular site agent**, a scheduled task on the IMAGEnet PC that runs nightly after clinic hours. It runs the same tool with the same parameters and produces the same output; it removes the operator from the export and upload, not the requirements from the procedure.

It reaches out; nothing reaches in. The agent makes outbound HTTPS connections only. There is no inbound firewall rule, no VPN into the practice network, no standing remote-access account for SpectOcular on clinic hardware, and no Google credential on site hardware. This is a scope control as much as a network one: remote access into an IMAGEnet PC would expose the practice's entire record, whereas the agent can only ever act on exams belonging to subjects the enrolment record says have consented.

Default deny. An exam is uploaded only if it matches **exactly one consented subject**. Because the export tool operates on .fda files rather than on patients, and no IMAGEnet path can be relied on to say whose scan it is, attribution is made from the identity inside the produced DICOM, after export. Anything else is quarantined — the staged copy destroyed on the clinic PC, and a reason code carrying no patient identifier reported centrally:

Situation	Outcome
Matches one consented subject	Upload
Patient not on the consented worklist	Quarantine — no-consented-subject-match
Two consented subjects share a clinic patient ID	Quarantine — ambiguous-subject-match
Two patients present in one export folder	Quarantine — mixed-identity-within-export
Wrong regulatory tier (§4)	Quarantine — wrong-regulatory-tier
Backend unreachable	No export runs at all

The last row is deliberate. An unattended machine must never decide for itself what it may upload while the consent record is unreachable.

The agent does not de-identify. It uploads Copy 1, identified, as produced (§6.1). SpectOcular continues to build the coded and de-identified copies centrally (§6.2, §7). Automating the site's upload does not move the de-identification duty onto sites.

Sites without the agent are not second-class. Some practices will not permit third-party software on a clinical PC. The manual procedure in §5.2–§5.3 remains fully supported for them.

Installation, preflight and operating procedure: [tools/site-agent/docs/INSTALL_WINDOWS.md](#). Design and decisions: [tools/site-agent/docs/ARCHITECTURE.md](#).

6. WHERE THE DATA GOES — TWO COPIES, DELIBERATELY DIFFERENT

This is the section to read carefully. The answer is two copies, and the value depends entirely on them not being identical.

A straight duplicate is worse than a single copy: it doubles the PHI footprint and buys nothing. The reason for a second copy is that the two serve different purposes with **different retention schedules under Protocol §14**.

6.1 Copy 1 — the subject folder (identified, source)

Destination: the enrolled subject's existing Drive folder,

01_Sites / Site NN - <Practice> / Subjects / <SubjectID> / 05_Retinal_Imaging_ZIP /

- Exported DICOM **as produced**, demographics intact
- Zipped per exam, named as §5.3
- This is the source record. It is the copy that supports a clinical question, a re-read, a query about what was acquired

- Governed by the identifiable-data retention schedule (§9)

6.2 Copy 2 — the SpectOcular research corpus (coded)

Destination: a SpectOcular-controlled corpus outside the site folder tree, organised by cohort rather than by site.

- Demographics **replaced with the Subject ID**, and identifying DICOM tags removed, per §7
- Contains no patient name, MRN or date of birth
- This is the copy that feeds analysis, grading and model development
- Governed by the de-identified-dataset retention schedule (§9)

Made by SpectOcular, not by the site. The site performs one export (§6.1) and uploads it. SpectOcular produces the coded and de-identified copies centrally from that upload. Site staff carry no second command and no de-identification responsibility — every extra step at the point of care is a step that gets skipped on a busy clinic day, and de-identification is not a task to distribute across sites.

6.3 Why two, stated plainly

Retention forces it. Protocol §14 says identifiers and linkage keys are destroyed when no longer required, while *"de-identified datasets used for model development may be retained beyond this period."* Those are different clocks. If the imaging exists only inside a subject folder keyed to the linkage system, you cannot honour the first without destroying the second. The separation has to exist before it is needed, because it cannot be created retrospectively once the key is gone.

Access forces it. Subject folders sit in a site's folder tree. If a site is ever granted access to its own subjects, the research corpus must already be somewhere else.

Analysis forces it. A model cannot be trained on data scattered across per-subject folders. The corpus needs its own organisation, and building it later by walking the subject tree re-introduces the identifiers you were trying to leave behind.

6.4 Recording custody

Every export is recorded in the master database Document_Uploads sheet with Subject ID, document type, file name, upload timestamp and uploader — written by the site agent where installed, and by hand otherwise. Agent uploads additionally record the SHA-256 digest, object count and exam date in Imaging_Uploads. The agent does not delete its local copy until the custody record exists, so a failure at this step costs a retry rather than the data. An export that exists in Drive but not in that sheet has no custody record and is not considered filed.

6.5 Repeat imaging and follow-up visits

A subject is enrolled once and keeps one Subject ID for the life of the study. A follow-up visit is not a new enrolment.

For a repeat scan the technician does nothing beyond scanning the patient as usual. There is no form to complete, no data to re-enter, and **no new Subject ID to generate**. The agent attributes the exam to the existing Subject ID by matching the clinic patient identifier against the consented worklist, and the visit lands in the same subject folder as the baseline:

Subjects / 03-HN7JZ / 05_Retinal_Imaging_ZIP /

03-HN7JZ_2026-08-18_a1b2c3.zip baseline

03-HN7JZ_2026-11-20_7f9e21.zip 3-month

03-HN7JZ_2027-02-14_c40d8b.zip 6-month

Re-enrolling a returning subject is an error, and a costly one. The enrolment form generates a fresh random code each time it is used. A subject enrolled twice acquires a second Subject ID and a second folder, their follow-up becomes unlinkable to their baseline, and the longitudinal change the study exists to measure is lost. It cannot be repaired afterwards without re-identifying both records to establish that they were the same person. Where a visit must be recorded manually, **look the subject up first and reuse their existing ID**.

The intervals survive de-identification. Stage 2 shifts dates by a fixed per-subject offset rather than deleting them (§7.2, VAL-001), so the spacing between a subject's visits is preserved into the research corpus while the calendar dates are destroyed. That is what makes the series usable for longitudinal analysis and model development after the identifiers are gone.

7. CODING AND DE-IDENTIFICATION — TWO STAGES

Coding and de-identification are different things and are performed at different points. Conflating them is the most common way a dataset ends up described as something it is not.

	Stage 1 — coded	Stage 2 — de-identified
Patient name, MRN, DOB	Replaced with Subject ID	Absent
Key linking ID to person	Exists — Patient_Linkage, MFA-restricted	Destroyed, or withheld from the recipient
Device serial, institution, exact scan dates, DICOM UIDs	Still present	Removed or shifted
Status	Still PHI. BAA applies	Outside HIPAA if done completely
May be disclosed externally	Only under BAA or DUA	Yes

7.1 Stage 1 — coded, at intake

Applied when SpectOcular builds Copy 2:

-patientData: ^<SubjectID>^^^^^^

The switch replaces PatientID, name fields, gender and date of birth. Parameter order is fixed; see Tool Instructions §1.1.

This is where most people stop, and stopping here is the mistake. After the switch, the DICOM still carries Study/Series/SOP Instance UIDs, device manufacturer, model and serial number, institution name, and acquisition date and time. Dates finer than a year are themselves Safe Harbor identifiers — a scan date plus a clinic's appointment book can re-identify a subject without any name at all. Burned-in annotation, where present, is untouched.

A Stage 1 corpus is **coded data under the BAA**. It is not de-identified and must not be described as such.

7.2 Stage 2 — de-identified, before any external release

Required before the corpus, or any subset of it, leaves the covered Workspace account — to a collaborator, a statistician, a contract reader, or a regulatory submission.

The step is a DICOM de-identification pass against **DICOM PS3.15 Annex E**, applying at minimum the Basic Application Level Confidentiality Profile with:

- **Retain Longitudinal Temporal Information** — **modified dates**, so intervals between visits survive while absolute dates do not. Longitudinal change is the scientific signal; the calendar date is not
- **Retain Device Identity: off** — manufacturer, model and serial removed
- **Retain UIDs: off** — UIDs regenerated, with the mapping held only inside the covered account
- **Clean Pixel Data: on** where any burned-in annotation is possible

Tool selected and validated — 18 August 2026. The record is **VAL-001**

(03_Quality_SOPs/VAL-001_Deidentification_Tool_Validation_2026-08-18), which carries the selection rationale, the profile, the method and the full result. In summary:

Tool	pydicom 2.4.4 driven by a written SpectOcular profile module
Profile	SpectOcular Basic + ModifiedDates, v1.0 — PS3.15 codes 113100, 113107, 113108
Source	tools/deid/spectocular_deid.py, SHA-256 d3dbf130a13babefb90ca16085f5efbe29a693f895da94c9 fe5fde617fcc3813
Validated	59 checks, 59 passed, 0 failed — including a residual scan confirming twelve known source values absent from every output element

pydicom with a written profile was chosen over dcm4che and the RSNA Anonymizer on auditability rather than capability: the profile is source code, so it is version-controlled, diffable and hash-pinned, and what a past run did is recoverable from the repository rather than from shell history or a screenshot of a settings pane. Any change to the profile invalidates the cited hash and requires re-validation (VAL-001 §5).

One confirmation remains open. VAL-001 validates the profile implementation against a synthetic file built to carry every identifier a Topcon export is known to write; no real .fda export exists yet. **On the first real export, before that scan or any other leaves the covered account,** dump its full tag set and confirm no identifying tag is present that the profile does not handle (VAL-001 §6). Until that confirmation, **no imaging leaves the covered account,** and the corpus is described as coded.

7.3 Why build Stage 2 now

Protocol §13 permits de-identification *"after initial data collection and quality assurance processes are completed"*, so deferring it is compliant. It is nonetheless cheaper now: retrofitting de-identification across a grown corpus means reprocessing everything, re-verifying it, and reconciling two generations of UIDs. Defining it while the corpus is empty costs a decision and a validation run.

8. VERIFICATION — DO NOT SKIP

Performed by the operator for every export, before the source .fda is considered archived.

8.1 Completeness

Confirm the expected DICOM objects are present (§3). An .fda with no B-scans yields OP only; this is expected. Anything else missing is a failed export.

8.2 Log review

Open the result_*.log written by -debug. Any error or warning is investigated before the export is filed.

8.3 Tier confirmation — the check that cannot be faked

Inspect the segmentation object filename. The final numeric component encodes the regulatory tier:

- ...7.1.dcm → US tier
- ...7.2.dcm → OUS tier
- ...7.3.dcm → **Research tier**

A study export must read 7.3. A clinical export must read 7.1. **If the suffix does not match the intended use, the export is quarantined and re-run on the correct installation** — it is evidence that the wrong licence was active, and any downstream use is suspect.

8.4 Identifier check on the research copy

Confirm the coded copy carries the Subject ID and no patient name, MRN or date of birth. Spot-check one DICOM header per batch.

This confirms **Stage 1 only**. It does not confirm de-identification — device serial, institution and acquisition dates survive Stage 1 by design (§7.1). Stage 2 has its own validation, performed once against the chosen tool rather than per batch (§7.2).

9. RETENTION AND DESTRUCTION

Per Protocol §14:

	Retention
Source .fda files	With the device/site per clinical record practice
Copy 1 — subject folder, identified	Minimum 6 years after study completion; destroyed with identifiers and linkage keys when no longer required
Copy 2 — research corpus, coded (Stage 1)	Retained with the identifiable-data schedule while the linkage key exists — it is still PHI
Copy 2 — de-identified (Stage 2, §7.2)	May be retained beyond that period for model development and validation, once de-identified and the key destroyed or withheld
Export logs	With the copy they document

The condition on the second row is the whole point of the two-copy model. Retaining a coded corpus indefinitely while the key still exists retains identifiable data by another name.

10. FAILURE MODES TO WATCH

Symptom	Likely cause	Action
Segmentation suffix 7.1 on a study export	Research licence not installed or wrong installation used	Quarantine, re-run on research installation
Batch output all in one folder	New Folder for Each FDA File not ticked	Re-run; do not attempt to separate by hand
Demographics wrong in DICOM but right in IMAGENet 6	Downstream edit did not propagate to .fda (Appendix B)	Re-export with -patientData:
No OCTA metrics, no OPTBSV/OPTENF	Structural-only device (current fleet, §3.1), or scan configuration outside Appendix B	Expected; not a failure
-bscanAVG produced DICOM but no smoothing	Relative path used for the .fda	Re-run with a full absolute path
Export in Drive, absent from Document_Uploads	Custody step skipped	File the record before the export is used
Imaging routed to Retina IQ for a research-only subject	RetinalIQRequested not checked before routing (§5.4)	Stop the routing; the enrolment record governs, not recollection
Site queried for OCTA objects it cannot produce	Site capability not consulted (§3.1, §5.4)	No action at the site; correct the expectation
Every agent export quarantines as wrong-regulatory-tier	Clinical (US/OUS) installation in use, not Research (§4)	Reinstall the research installation; no data is lost, exports re-run
Every agent export quarantines as no-consented-subject-match	Nobody enrolled at that site yet, or the clinic patient ID on the device differs from the one captured at enrolment	Correct the enrolment record; do not relax the matching
Agent reports ambiguous-subject-match	Two enrolled subjects carry the same clinic patient ID	A data problem at the site. Fix the enrolment; never resolve it by picking one
A site stops appearing in the fleet report	Agent stopped, PC replaced, token revoked, or endpoint protection quarantined the executable	Investigate within one working week — a quiet site is indistinguishable from a site with no patients
A returning subject appears under a second Subject ID	Re-enrolled instead of looked up (§6.5)	Stop before more visits accrue; merging requires re-identifying both records
Request to send imaging outside the Workspace account	Stage 2 profile validated (VAL-001), but not yet confirmed against real device output (§7.2)	Refuse until the first real export has been tag-dumped and confirmed against the profile per VAL-001 §6

11. TRAINING AND DELEGATION

Only personnel trained on this SOP perform exports. Training is recorded in the individual's study file and includes: tier distinction (§2), the two-copy model (§6), and the verification steps (§8). Delegation is recorded on the site delegation log per SOP-002 §10.

12. APPROVAL

This SOP is not effective until signed.

Medical Director _____

Date _____

APPENDIX A — COMMAND REFERENCE

Purpose	Command
Study export, standard (site)	DicomOctExport.exe "<in>.fda" "<out>" -allDcm -thicknessMap -json -debug
Thickness as a second pass, if incompatible	DicomOctExport.exe "<in>.fda" "<out>" -thicknessMap -debug
Coding for the corpus (SpectOcular, §7.1)	add -patientData:^^<SubjectID>^^^^^^
Discover Hardware ID	run DicomOctExport.exe with no arguments
Reusable parameter set	-config "<path>\parser.config", editing OptionSwitch in <appSettings>

No clinical export command is listed: SpectOcular does not deploy the US-tier installation (§4).

Mutually exclusive: -dcm / -straightDcm; -dcm / -bscan. Full matrix in Tool Instructions §3.

APPENDIX B — WHEN OCTA METRICS ARE AVAILABLE

Vessel density — Maestro2 and 3D OCT-1, macula, superficial, ETDRS grid, 3.0/4.5/6.0 mm at 320×320; DRI OCT Triton additionally 6.0 mm at 512×512; Triton disc RPC RNFL 4.5 mm at 320×320.

FAZ — DRI OCT Triton only, macula, superficial, 3.0×3.0 mm, 320×320.

Not available on the US tier at any configuration.

Not applicable to the current device fleet at all — the pilot fleet is structural-only (§3.1). This appendix applies from the point an angio-capable device (e.g. DRI OCT Triton, or an angio-enabled Maestro2 / 3D OCT-1) enters the study.

APPENDIX C — DECISIONS SET ON 2026-08-18, FOR REVIEW

All five were reviewed and settled on 18 August 2026. They are recorded here so the reasoning survives the people who made it.

6. **Research tier only; no clinical installation (\$4).** Confirmed. Clinical imaging is done in IMAGENet 6, so this tool serves the study alone. A site holding one installation cannot select the wrong one, which turns \$8.3 from error-detection into confirmation. Revisit only if a documented need for clinical DICOM from this tool arises.
7. **Two copies, identified and coded (\$6).** Confirmed. Driven by Protocol §14's differing retention clocks — a single copy cannot satisfy both, and identical duplicates satisfy neither. Refined on review: **the site performs one export and SpectOcular builds the coded copy centrally**, so no de-identification duty is distributed to sites.
8. **Stage 2 de-identification defined now, not deferred (\$7.2).** Changed on review. The original default was to describe the corpus as coded and add de-identification before first external transfer. Building it now was judged cheaper than retrofitting across a grown corpus. Tool selected and validated on 18 August 2026 — pydicom with a written, hash-pinned profile, 59/59 checks passed (VAL-001). The corpus remains coded and nothing leaves the covered account until the profile is also confirmed against a real Topcon export.
9. **-allDcm -thicknessMap -json -debug (\$5.2).** Amended on review to add -thicknessMap: the study's endpoints are layer-thickness biomarkers, so the thickness binaries are captured directly rather than derived later. Compatibility with -allDcm is to be verified on the first real scan, with a second-pass fallback documented.
10. **OCTA switches retained despite a structural-only fleet (\$3.1).** Confirmed. Inert on non-angio files, and keeping them means the command does not change when an angio device joins. Scope of an upload is bounded instead by consent and site equipment (§5.4).
11. **Automated export and upload at the site (\$5.5), added 18 August 2026.** The manual procedure was written first and remains the fallback. Automation was added because the export tool is a command-line binary and therefore schedulable, and because the alternative — remote access into each clinic PC — does not scale past a handful of sites and exposes the practice's whole record rather than the consented subjects' exams. The division of labour in §6.2 is unchanged: the site still performs one export of Copy 1, and SpectOcular still builds the coded and de-identified copies centrally.
12. **One Subject ID for the life of the study (\$6.5), stated explicitly 18 August 2026.** Raised on review: the enrolment form mints a fresh random code on every use and has no returning-subject lookup, so a technician re-enrolling a follow-up visit would create a second Subject ID and silently sever the subject's time series. Under the agent a repeat scan requires no enrolment at all. A backend guard that returns the existing Subject ID instead of minting a new one is drafted and awaiting deployment.

Implementation item carried out of this review, now closed: selection and validation of the Stage 2 de-identification tool (§7.2) — completed 18 August 2026, recorded in VAL-001. The one task it leaves behind is confirmation of the profile against the first real Topcon export (VAL-001 §6). Nothing may be disclosed externally until that confirmation is done.

APPENDIX D — REFERENCES

- DICOM OCT Export Tool, User Instructions Rev E, v2.1.2 (Topcon), 14 November 2025
- DICOM OCT Export Tool Description v2.1.2 Rev L
- SpectOcular Study Protocol v1.4, §13 Data Handling and Confidentiality, §14 Data Retention and Destruction
- SOP-001 Retinal Image Grading, §2 (clinical vs research separation)
- SOP-002 Reader Qualification and Certification, §10 (delegation logs)
- DICOM PS3.15 Annex E — de-identification profiles
- VAL-001 — De-identification tool selection and validation record (§7.2)
- tools/site-agent/docs/ARCHITECTURE.md — site agent design (§5.5)
- tools/site-agent/docs/INSTALL_WINDOWS.md — site agent installation and operation (§5.5)