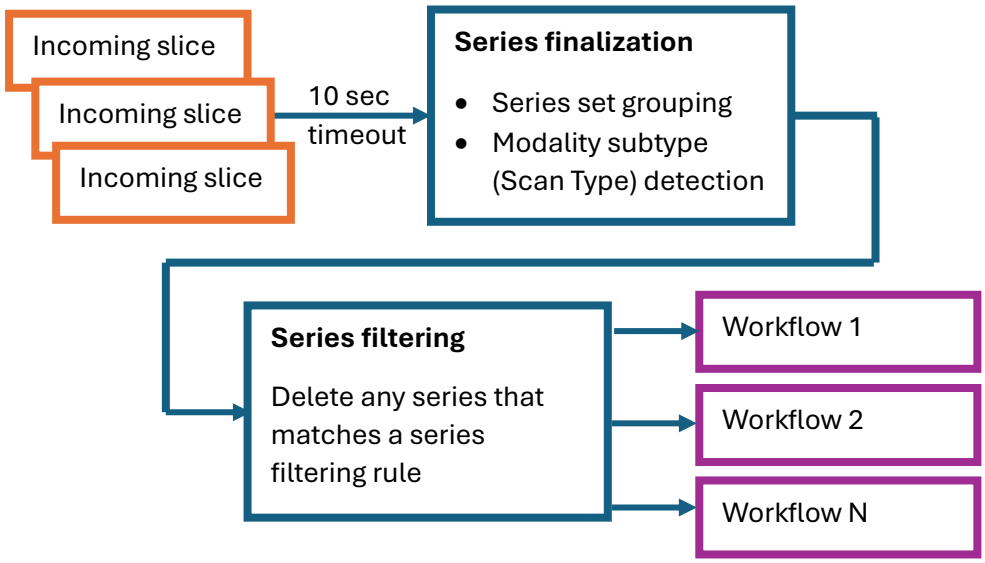
How DICOM Series Routing Works on Next Gen

Hello, intrepid Brainomix system wrangler! It's Eric here, with a half-baked explanation of how this stuff works. The basics:



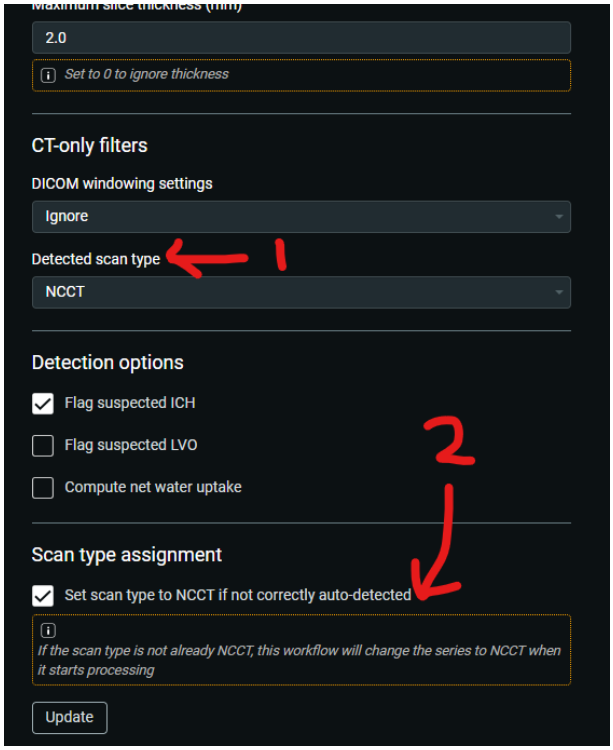
The Next Gen server acts more like a PACS than in Zygoma.

- Individual DICOM slices are stored as they arrive. The series (as determined by Series Instance UID tag) displays as "Uploading" in the patient list.
- Then, after a 10 second timeout when no additional slices have arrived, the series enters Series Finalization. This has a couple of sub-stages:
- The system checks for possible **Series Sets**. A Series Set is a concept we use for scans that arrive as separate series (different Series Instance UID) but actually need to be processed as a single dataset. See Appendix 1 for definition of the rules used.
    - The typical example is a Canon Aquilion One CTP scan, when configured in a mode that generates e.g. 19 separate series (one for each volume) but we need to process as a single CTP.
    - Other examples include the one-slice-per-series Aquilion Prime CTP protocol, and some multi-series DWI MRI datasets.
    - Note that we don't have any rules that try to group multi-phase CTA datasets sent as separate 3D series. HOWEVER – if a series set grouping results in fewer than 6 distinct 3D volumes, it will be flagged as "mCTA" subtype by the next rule (not just yet though)! Auto-mCTA-grouping like this is almost always an indication of something going wrong.
    - If this series is part of a series set, then maybe another series in the series set is still uploading (not at the 10 second timeout yet) in which case the series set will not be finalized yet – but this series within it will be.

- The system applies a **Modality Subtype**, which is what is displayed in the workflow routing rules as "Scan Type". DICOM files only indicate a "Modality" of "CT" or "MR" (or whatever). We need to know more: if it's a CT, is it NCCT, CTA, or CTP? If it's MR, is it DWI, DSC, FLAIR or something else? That's why we try to autodetect the Modality Subtype. We try to make this detection as foolproof as possible, because if it's correct, it's the only information you need to correctly route a series! See Appendix 2 for definition of the rules used.
- The system emits a SeriesCreated event (and a SeriesSetCreated event if the series is the last one in a Series Set). The series will now display as "Uploaded".
- **Series Filters** (configured in DICOM/Series Filters) then get applied. Any configured Series Filter works like any other normal workflow, getting matched with the new series with its configured rules. If any Series Filter matches, the series is deleted (and purged) immediately, and no further processing will happen.
- Remaining series then get picked up "simultaneously" by the configured processing workflows. **Importantly, if one workflow picks up a series and processes it, it doesn't stop other workflows also picking it up and processing it!** This is the major change from Zygoma. (I put "simultaneously" in scare-quotes because actually the events get handled sequentially by a single process, so they won't execute truly in parallel). "Picking up" an event means that the workflow looks it the new series and applies its matching rules.
- Processing jobs queued by the matching workflows then get processed sequentially. Each processing job can either add a "Failed result" or a normal Result to the series. If a series has been processed by multiple workflows, it will have multiple results – this often indicates a configuration problem (we don't usually have a good reason to process the same series multiple times). The series will show as "Queued" or "Processing" while this happens, then "Complete" when all processing is complete.
    - If the "Scan type assignment" checkbox is selected in a workflow, when processing starts, the system will assign the workflow's hardcoded Scan Type to the scan. That might mean no change (e.g. if an NCCT is processed with e-ASPECTS, setting it to NCCT won't do anything) but it might change the Scan Type to something else (e.g. if an NCCT is processed with e-CTA due to your workflow rules, it could change the scan type to CTA if this box is checked). More on this soon.

Scan type (modality subtype)

OK, so now your head is spinning and you feel suffused with a sense of omniscience and power, now that you know how stuff works. But wait, that Scan Type stuff is still a bit confusing. It's worth reiterating a couple of things, especially in context of where it appears in a workflow config admin page:



The "**Detected scan type**" (item 1) is a filter. Only scans with this type assigned by the initial modality subtype auto-detection logic will be processed by this workflow. In this case (an e-ASPECTS default workflow) it's set to NCCT, so anything detected as NCCT will match, and all other types (CTA, CTP, Unknown etc) will not.

If you don't care what the scan type is, you can set this to "Ignore", but this is generally a bad idea. Even if scan type is incorrectly detected by the automatic rules, setting it to a specific "wrong" type (to match the incorrectly-detected scans) is better than choosing Ignore.

The "**Scan type assignment**" section (item 2) is only triggered if this workflow has matched a scan and started processing it. If the box is checked, and the series starts processing on this workflow, then the scan type will be forced to change to the named type (NCCT in this case, because this is an e-ASPECTS workflow).

Example

Therefore, if you're closely watching the patient list as a scan comes in, you might see the following:

- While scan is "Uploading", it shows as "CT" (we haven't run modality subtype detection yet)
- 10 seconds after the last DICOM slice was uploaded, it changes to "Uploaded" state and shows the detected modality subtype (e.g. "NCCT").

- Now imagine you've configured e-CTA to match "Detected scan type" of "NCCT" (because maybe you have a badly configured scanner sending CTA scans with no ContrastBolusAgent tag) along with some other rules that this scan matches. It would then show as "Queued" and then "Processing", at which point its scan type would have changed from "NCCT" to "CTA".

Important: Note that even though it now has "CTA" type, any workflows looking for "CTA" scans will **not** pick up this scan for processing, because when it was finalized, it was "NCCT" type. **However**, if you choose "Reprocess scan", it will **not** process on the original workflow (because that workflow was looking for NCCT scans) and **might** process on a different workflow that's been configured to look for scan type of "CTA" (if it matches all other filters on that workflow too).

This is obviously confusing, and therefore we prefer to avoid incorrect modality subtype detection wherever possible, to avoid having to configure workflows that will have to change the modality subtype and thus cause this confusion. Unfortunately, it can't always be avoided, so if you have scenarios where the modality subtype detection fails, please let support / product know, so we can improve the detection in future. Our long term aim is to have perfect modality subtype auto-detection, which would then mean the default workflows will always work.

Recommended configuration

Great, so now you know everything, what should you do? That's easy!

Rule 1: Always try to only use default workflows

This should work for most sites. The default workflows are pretty simple:

- There is a Series Filter that deletes any DICOM series with LOCALIZER in (0008, 0008).
- e-ASPECTS looks for scan type "NCCT", no ContrastBolusAgent, 50-1000mm size in Z, max 2.0mm slice thickness. Note that the "no ContrastBolusAgent" rule is unnecessary, and we could set to Ignore quite safely, but whatever.
- e-CTA looks for scan type "CTA", must have ContrastBolusAgent, 50-1000mm size in Z, max 2.0mm slice thickness. Note that the "must have ContrastBolusAgent" rule is unnecessary, and we could again set to Ignore quite safely, but whatever.
    - There is a default 4D mCTA workflow looking for scan type "mCTA", 50-1000mm size in Z, max 2.0mm slice thickness. This should work for any 4D mCTA.
- e-CTP looks for scan type "CTP", 0-1000mm size in Z, max 10.0mm slice thickness.

At this point, with a site with well-configured scanners only sending the series you need, you might be done – no extra magic needed. However, there are often some extra things you need to think about, I've tried to list these next steps in order of how frequently I think they might be needed.

Rule 2: delete unnecessary stuff

Lots of sites will be sending you stuff you don't want: dose reports, coronal / sagittal recons, 5mm recons, bone windows, etc etc. If you followed Rule 1 above, at least the localizers should be getting deleted already. For all the other stuff:

- Decide whether people might want to view them. If so, they should be fine as-is (the default rules will hopefully not try to process any of these things – if they do, try constraining the filters on the relevant workflows further).
- Add Series Filter rules for deleting anything you don't want to be kept. Try to avoid using series description regexes – instead, look for a specific DICOM tag that describes the thing you want to delete.
    - If it's a bone window, maybe using Scan Type = "NCCT" and "Scans must not have a brain tissue window" would be enough.
    - If it's a dose report, maybe you can look at tag (0008, 0008) ImageType and see if it has a unique value that you can fully or partially match on

Once you've sorted your filter rules, obviously keep an eye on incoming studies and make sure everything is getting applied correctly. It won't be obvious what is being deleted (because it's gone!) unless you look carefully at the DICOM log, but even that won't give you much data about the stuff that's been auto-deleted! So be careful that you're not auto-deleting real scans that should be processed, because it won't be obvious if you are doing that.

If you think there are additional DICOM Series Filter rules that should be applied globally, let us know, and we can add them to the new installation defaults in future software releases.

If you want to delete all data EXCEPT things matching a filter, you can do this by implementing a Series Filter with an inverse-match regex, for example if you only want to keep scans with a Series Description containing “good”: something along the lines of (?!good) should work. (Ask a developer if you need help developing a regex).

Rule 3: process CTA missing ContrastBolusAgent tag

Some CTA scans are missing the ContrastBolusAgent tag because something in the scanner / auto-injector / workstation / PACS hasn’t been configured correctly. These will be detected as “NCCT” by the auto-detection logic. Therefore:

- You need to ensure the e-ASPECTS workflow doesn’t pick up these scans. A safe way to do this might be to add the “Must have a brain tissue window” filter to the default e-ASPECTS workflow, which will hopefully continue to match true NCCT head scans and ignore these CTA-looking-like-NCCT scans. If that doesn’t work, consider changing the “Acceptable range of size along z-axis” max to something like 350mm.
- You then need a new e-CTA workflow that can pick up these scans. Use:
    - Acceptable range of size: min at least 350mm (or whatever is needed for your particular scanner’s typical acquisition) [this is optional: the below rules might be enough]
    - DICOM windowing settings: Scan must not have a brain tissue window
    - Detected scan type: NCCT
    - Scan type assignment: set scan type to CTA if not correctly auto-detected

Then you can try sending some scans for processing and test all is fine. Remember, don’t try to re-process previously incorrectly routed scans if their Scan Type has been changed by a previously incorrectly configured workflow, because it won’t work!

In the future, we might change the auto-detection logic to use the window and / or size of the scan and / or body part tags) to correctly identify these CTAs (ideally with no false positives on true NCCTs) to avoid you having to do this manually.

Rule 4: process 3D multi-phase CTA scans

For sites performing mCTA and not generating a single 4D series, you’ll have to create a new mCTA workflow for 3D series, because this is not set up by default. You’ll also have to stop the series being processed individually by e-CTA, because they will each probably look like a valid CTA scan.

If you are expecting ONLY mCTA and no single-phase CTA, then:

- Disable the default e-CTA workflow
- Create a new multi-phase CTA workflow matching “Acute” and “Study” queues and with all defaults otherwise (Detected scan type = CTA, expected number of series = 3, single-phase fallback = 120 secs). This should just work in most cases.

If you are expecting BOTH mCTA and some single-phase CTA, then:

- Create a new multi-phase CTA workflow matching “Acute” and “Study” queues and with all defaults otherwise (Detected scan type = CTA, expected number of series = 3, single-phase fallback = 120 secs), except adding a Series Description regex for a known string that all your multi-phase scans will have in their series description (e.g. MULTIPHASE).

Make sure the radiology department knows that mCTA will only be processed correctly if all mCTA scans have this word in their series description. Test that it works.

- Add a Series Description regex to the default e-CTA workflow to exclude scans with that string (?!MULTIPHASE) or whatever. Test that it works.

This is pretty much the only case where I would recommend using a Series Description regex to figure out what is mCTA vs single-phase CTA (if you're expecting both), because there often isn't an obviously more reliable way of differentiating the two. If we find one, we can consider adding a new default 3D mCTA rule, but that might be difficult.

New installs of Next Gen vs upgrades from Zygoma

A key difference between new installs and upgrades from Zygoma is that you will **not** get default workflows or default series filters when you upgrade – instead, the upgrade script will try to move existing workflows as best it can. **This is often not really the best option!** It means that you won't get the default "Delete Localizers" series filter, and the workflows will be using "old" Zygoma logic instead of just using the Scan Type (modality subtype) for routing. For most sites, it should be possible to change all the workflows to just use the Scan Type and nothing else for routing, and it should Just Work™! Of course, this should take into account the known exceptions mentioned above, and it won't work for all sites.

Appendix 1: Series Set rules

These are subject to change (by developers – you can't change these in any way), but currently look like this:

Non-DWI case

A non-DWI series is a candidate for being a multi-3D series with another series if they have (in terms of DICOM tags):

- the same **Modality** (e.g. "CT" or "MR")
- the same **SeriesDate** and **SeriesTime**
- the same **FrameOfReferenceUID**
- the same **ContrastBolusAgent** tag (or same lack of such a tag)
- the same **ConvolutionKernel**
- the **Acquisition Date + Time** is within 300 seconds of the other series

and:

- (a) there is a **TemporalPositionIndex** tag AND only one volume per series, **OR**
- (b) the series is a single-slice, multi-volume series AND the number of volumes matches

Rule (a) is to avoid grouping of different recons of the same 3D series with different timestamps, and to avoid grouping of different recons of 4D datasets with one series per 4D sequence.

Rule (b) is to ensure grouping of Aquilion Prime (or other) one-slice-multi-timepoint-per-series CTP scans.

DWI case

There is a separate set of rules (from 14.4 onwards) for DWI. A series set is created for DWI with another series if:

- The new series is detected as DWI by the subtype detection logic (even though this hasn't officially run yet! We run it specially for DWI if we've got here and noticed we have an incoming scan with **Modality** == MR)
- Both are single-volume series (not multiple timepoints)
- The other existing series has ModalitySubtype UNKNOWN or DWI
- Both series have the same **StudyInstanceUID**
- the same **SeriesDate**
- the same **FrameOfReferenceUID**
- The B-values are complementary (i.e. the B-values in series 1 don't overlap with the B-values in series 2)
- The slice counts of both series are the same

Once a series set has been created, you can't ungroup it again – if you think series-set logic has incorrectly grouped some series that should not be treated as a single dataset, please report it to support or a developer so we can fix the bug!

Appendix 2: Modality subtype detection rules

The modality subtype (Scan type) assignment logic is as follows. We start by looking at the first DICOM instance within the series (or series set) linked to a series image and check its Modality tag.

Case 1: Modality is CT

We start by looking for any axial images, based on our instance-grouping logic and parsing of image orientation tags. This is complicated, and looks at lots of stuff in the DICOM files. Then:

- If we have mixed orientations, discard any non-axial images
- If the first remaining image is non-axial, or there are no images, return **UNKNOWN**
- If it's in a series set, or is Enhanced CT, and there are >= 6 volumes, type is **CTP**
- If it's in a series set, or is Enhanced CT, and there are > 1 and < 6 volumes, type is **mCTA**
- If it is Enhanced CT and has a **ContrastBolusAgentSequence** tag with at least one entry, type is **CTA**
- If it is Enhanced CT and the above rules haven't matched, type is **NCCT**

Otherwise, use our standard loaders to try loading all instances and split into volumes based on the **ImagePositionPatient** trend changing when sorted by instance number (another relatively complicated bit of logic). Then:

- If there are >= 6 volumes, type is **CTP**
- If there are > 1 and < 6 volumes, type is **mCTA**
- If there is a **ContrastBolusAgent** tag and it is set to a non-empty value that is not "NO" (case-insensitive) then type is **CTA**
- Otherwise, type is **NCCT**

Case 2: Modality is MR

We start with logic for series-sets that may have done some DWI grouping:

- If it's already part of a series set, and B-values are found, type is **DWI**

Otherwise, the dataset could be 3D or 4D, from classic or Enhanced MR, or a set of MOSAICs. Could be DSC, DWI, FLAIR but also MRA-TOF, T1, T2 … We only have subtypes for the first three (the ones we can process). Rules are applied in this order:

- If it has valid diffusion values (computed by some magic in a complicated function), and it has a **TE** > 50, type is **DWI**
- If it has > 15 volumes, and **TE** < 70, type is **DSC**
- If it has one volume, and either: (1400 < **TI** < 3000) OR ("FLAIR" [case-insensitive] is in **SeriesDescription**), type is **FLAIR**

Case 3: Otherwise

In any other case, the modality subtype will be set to **UNKNOWN**.