J Digit Imaging (2015) 28:264–271
DOI 10.1007/s10278-014-9748-2

Losing Images in Digital Radiology: More than You Think

Catherine Oglevee · Oleg Pianykh

Published online: 17 December 2014

Abstract It is a common belief that the shift to digital imaging some 20 years ago helped medical image exchange and got rid of any potential image loss that was happening with printed image films. Unfortunately, this is not the case: despite the most recent advances in digital imaging, most hospitals still keep losing their imaging data, with these losses going completely unnoticed. As a result, not only does image loss affect the faith in digital imaging but it also affects patient diagnosis and daily quality of clinical work. This paper identifies the origins of invisible image losses, provides methods and procedures to detect image loss, and demonstrates modes of action that can be taken to stop the problem from happening.

Keywords Digital imaging · Digital Imaging and Communications in Medicine (DICOM) · Digital imaging management

Background

One of the principal expectations of digital medicine is the ease and reliability of digital data exchange. Before PACS and Digital Imaging and Communications in Medicine (DICOM) were introduced in the early 1990s, medical imaging was plagued with data transfer problems: printed image films were expensive to print, hard to move around, but fairly easy to lose. Moreover, printed images typically existed in single copies, so any loss would entail redoing the entire patient exam—a high price to pay for poor image management technology.

Digital imaging was meant to solve all of these problems at once, making imaging data easy to copy, transmit, and display [8, 9]. DICOM standard indeed revolutionized the entire radiology workflow [4, 5]. However, we seem to forget that what sounded so good in theory could run into new practical obstacles of its own. And as we have recently discovered in our work at Massachusetts General Hospital (MGH), the old problem of "lost film" has not gone away, but has mutated into a new, "digital" form. In this paper, we would like to share our experience and approaches to identifying and eliminating digital image losses.

Image Loss in Digital Era

Despite decades of DICOM standardization, cross-vendor connectathons,[1] and perpetual workflow tweaking [2], the reliability of digital image transfer still cannot be taken for granted. The digital pipes in a hospital can be leaking without anyone even knowing the scope and consequences of these leaks. Despite the gravity of this problem, there have been very few publications on this subject [3, 7, 10]. This implies that most hospitals are still completely unaware of their image leaks [1].

Dealing with this problem on a daily basis, we made a list of the most common image loss problems and their origins:

1. *Vendor Incompatibility.* Although all vendors claim DICOM conformance, they still provide their own implementations (and often interpretations) of how the standard transactions should work. For example, when PET images started disappearing in our radiology

[1] http://www.ihe.net/connectathon/

C. Oglevee · O. Pianykh (✉)
Massachusetts General Hospital, #470, 25 New Chardon Street, Boston 02114, MA, USA
e-mail: opianykh@mgh.harvard.edu

C. Oglevee
e-mail: coglevee@hamilton.edu

Fig. 1 Example of original DICOM data formats (top VR values in each line) being replaced by different (bottom values in each line) in cross-vendor communications. VRs are responsible for correct DICOM data processing, and their modification can easily make this data unreadable. The screenshot was provided by the vendor of the workstation, where the corrupted images were failing

DICOM tag	VR	Data	
(0029, 0013) Private Creator	LO	1	MITRA PRESENTATION 1.0
(0029, 1031) PMTF Information 1	**LO**	1	202.0.8317148
	OB	*<Binary Data>*	
(0029, 1032) PMTF Information 2	**UL**	1	131072
	OB		
(0029, 1033) PMTF Information 3	**UL**	1	0
	OB		
(0029, 1034) PMTF Information 4	**CS**	1	DB TO DICOM
	OB		
(0029, 1160) Not in Dictionary	**LO**	1	Com
	OB		
(0029, 1208) Platform OOG Type	**CS**	1	MEDCOM OOG 2
	OB		
(0029, 1209) Platform OOG Version	**CS**	1	VD30C
	OB		
(0029, 1210) Platform OOG Info	OB	1	**<Binary Data>**

department, it took several months to realize that the digital archive from one major vendor was redefining data format tags (DICOM VRs) in the PET images acquired with another vendor (Fig. 1). This modification was a clear violation of DICOM, but it would result in failures only for a particular cross-vendor, cross-product combination. This is why it remained unnoticed for years, with images being lost, and with us being unaware of these losses.

2. *Unsupported Options*. Even when two image-exchanging devices (such as scanner and PACS archive) are fully DICOM-compliant, they may implement different standard and non-standard options, which may fail when mismatched. The most classic example is the use of image compression. There are 18 flavors of JPEG compression alone in DICOM standard, and checking "use compression" on the sending device by no means guarantees that the same compression algorithm will be supported by the receiving application. In theory, when two DICOM devices cannot agree on advanced transfer options, they should fall back to the default, most basic DICOM transfer syntax. In reality, the devices may just fail and stop communicating (Fig. 2). Avoid using any advanced options or test them completely even if they look obvious.
3. *Overload*. Even the most perfect imaging pipeline has its own bandwidth. When its limit is exceeded, image failures and losses become inevitable. The first example to come to mind might be a bottlenecked network. This, however, is not the only one; any hardware or software overload might result in medical image loss. As an example, we recently discovered that one of our DICOM routers was failing during the busy transmission times simply because its database could not index more than a certain amount of images per second. When this amount was surpassed, the database would start skipping some of the images. Ironically, the images were still transmitted fully and correctly, but when dropped from indexing, they would not be listed anywhere and not included in subsequent transmission or display. In this way, losing image records would still result in losing the images from the diagnostic pipeline, even if the "forgotten" image files can still be found somewhere.

   Image compression, mentioned earlier, often results in similar overload problems. Medical device administrators tend to turn the compression on, in hopes to reduce network traffic and storage. This makes sense, but often only in theory. Image compression algorithms require significant processing [6] which in turn consumes time and computer resources. As a result, naïve exercise of image compression may slow down your image transfers, and overload your devices, especially when you deal with high image volumes and old computer hardware (a sadly typical clinical setup). Test the compression with a stopwatch to make sure it delivers what is expected, and disable it if it does not.

   Note that overload problems usually happen outside the realm of standardized image transactions: DICOM does not manage your resources' utilization. As a result, DICOM image transfer may complete without any errors (such as in our indexing example above) and may not signal any problems.[2] Yet, the system will fail right after, without any notification of the failure.
4. *Collateral Image Loss*. Collateral loss is the most difficult and the most random case. The images are lost due to some other (often completely unrelated) problems. What may come to mind are

[2] This also explains why more detailed DICOM protocols, such as Storage Commitment, may not be able to catch all image leaks.

```
23.11.13 18:23:50.312 CsaStuImageTransformer::checkConversionResults -> JPEG lossless decompr.: OK <-2131
23.11.13 18:23:50.484 CsaStuImageTransformer::checkConversionResults -> JPEG lossless decompr.: OK <-2131
23.11.13 18:23:50.640 CsaStuImageTransformer::checkConversionResults -> JPEG lossless decompr.: OK <-2131
23.11.13 18:23:50.812 CsaStuImageTransformer::checkConversionResults -> JPEG lossless decompr.: OK <-2131
23.11.13 18:23:51.671 CsaStuImageTransformer::checkConversionResults -> JPEG lossless decompr.: OK <-2131
****** CsaStatus: item at index :0
  nlsMessageId   : CSA_TF_E_146
  fileName       : \MedCom\src\PATEX\MBUconv\src/CsaCvrtJis2Uni.cpp
  lineNumber     : 1439
  handlingFlag   : 16777216
  freetext       : Illegal JIS code!
  nr. of params  : 0
  host           : TRUEDNUCMED
  process        : CsaGenericMain
  timeStamp      : 2013 11 23  18:23:51.765

CsaStuDicomToDbHandler::svc -> An error was detected in this request, handling error now806
23.11.13 18:23:51.780 CsaDicomCStorePerformer::abortAllAssoc() aborting assocs
```

Fig. 2 Radiology workstation log file fragment: example of image decompression process, quietly failing in the background. As a result, the workstation aborts all other ongoing image transfers ("associations"), leading to even more substantial data loss. No alerts are shown to the users

blackouts and server failures, but in reality, collateral image loss may also occur under less dramatic circumstances and much more frequently. A good example would be the PET vendor-incompatibility problem described earlier. When we were looking for the causes of that failure, we discovered that the failure would occur only for specific studies, when DICOM VRs from the original vendor were corrupted by the other vendor device. But despite this clear failure pattern, each attempt to re-transmit the same case was resulting in a different number of failed images. Moreover, even the images with non-corrupted tags were failing in some instances and transmitting fine in the others.

This was a really puzzling case, and only the analysis of the vendor log files helped to explain what was going on. It turned out that when the receiving workstation was failing on a corrupted image, it would abort all other concurrent image transfers (Fig. 2, second box). Thus, in our case, a single corrupted image would randomly abort 40–50 perfectly normal images, being received at the same instant. Due to this domino effect, there was loss of much more data than was originally corrupted. You may have significant collateral image losses on a daily basis, with no red flags raised, and with completely random patterns.

These four types of digital image losses could be very persistent and substantial and must not be ignored. Moreover, in most cases, these losses occur well before the images reach the radiologists, meaning that diagnostic information is lost and patients' lives are put at risk. In most cases, these losses happen in the background, without alerting the users. Finally, in most cases, the users (radiologists included) do not even expect anything to get lost—we all thought we had solved this problem with PACS a long time ago, had we not?

Unfortunately, the problem is still there, and this teaches us an important lesson: what was once perceived as a pure "uptime," "six sigma" data flow could in fact be interleaved with frequent and hidden downtime incidents. Therefore, the only way to fight image loss is to become aware of it, to identify and eliminate its origins and to design "loss-resistant" processes.

Detecting Image Loss

In order to identify and locate the image loss problem, we developed a few different tactics for checking image counts on the sending and receiving devices:

Method 1 *Site visits*. We visited a particular scanner within the hospital and recorded patient name, study date, accession number, series names, and total image count within each series. About 1 week's worth of this information was collected from the scanner and then was compared to information in PACS.

A spreadsheet was made for each scanner and it logged total image count on the sending device (such as a scanner) and total image count on the receiving device (such as PACS, Table 1.) Whenever there was a discrepancy in these two numbers, we investigated the case further to identify the cause

Table 1 Sample spreadsheet to identify image losses

Accession no.	Scanner image count	PACS image count	Affected series
12345678	100	100	None
99999999	1000	999	Series X (p1)

for the differing numbers. The column entitled "Affected Series" is where we recorded any series that were partially or completely not sent to PACS. The notation "p1" means that "Series X" was a partially sent series and one image is missing on the receiving device.

Detailed recording of patient cases caused lengthy visits to the scanners. Method 1 also required much coordination with the technologists; finding a time when workflow would not be interrupted was not easy. Thus, gathering the amount of data we needed from each patient file was a challenging task. After using this method for a couple days, we realized how inefficient it was and decided to change our procedure.

Method 2 *DICOM Query*. We decided to identify any and all scanners that could be connected to remotely via DICOM Query so that physical visits to the scanners were minimized (Fig. 3). Remote connection allowed for easier and faster comparisons between scanner and PACS information. The same spreadsheets from method 1 were used, but with method 2, they were completed much more efficiently; we were able to crosscheck series names and images counts at the same computer. Moreover, connecting to scanners remotely does not interrupt the workflow and is therefore the most efficient way to collect sending device information.

Scanners that could not be connected to remotely were visited physically. Instead of using method 1 procedures to collect case information, we decided to simplify the process (Fig. 3). This time, every piece of information in a case was not recorded; instead, only case accession numbers were recorded. We then took these accession numbers and searched them in PACS, printing out a list of their series names and image counts. We took these printouts back to the scanners to do the crosscheck between the scanner image counts and PACS image counts. While this method requires two trips to a scanner, the total time spent at the scanner was greatly minimized compared to method 1. More importantly, workflow was not interrupted; it was easier to coordinate short blocks of time with the technologists as opposed to one long block of time.

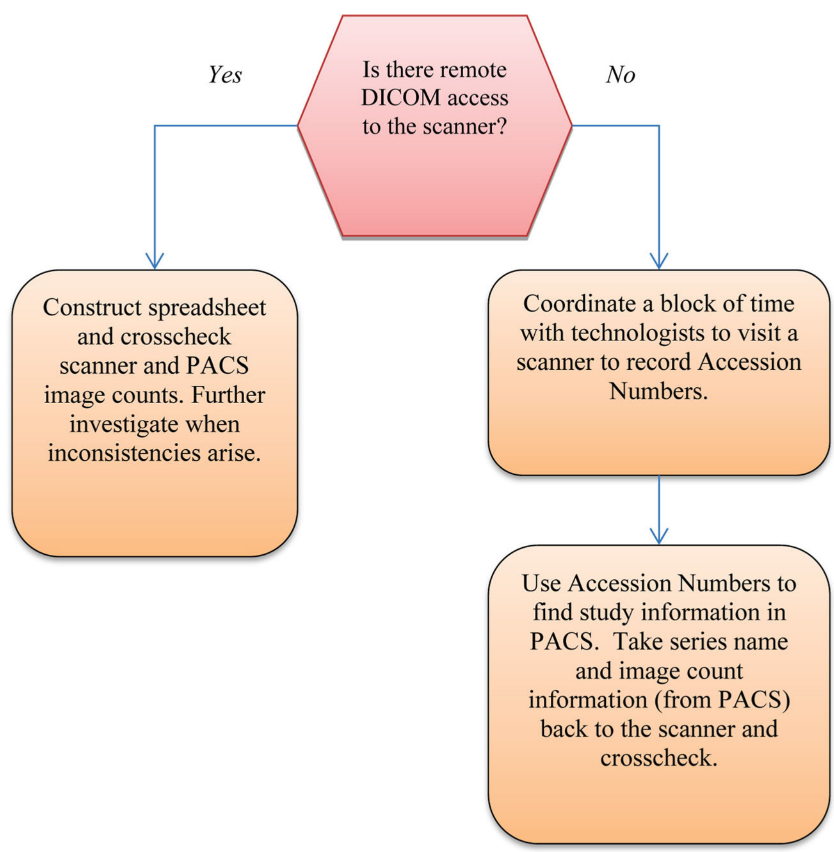

Fig. 3 Procedures to take when gathering information from the sending device

Fig. 4 Example of the sending device (scanner) dropping an entire series when sent to the receiving device (PACS)

Scanner		PACS	
Descr.	*Images*	*Descr.*	*Images*
Dose Report	1	Dose Report	1
COUTS	2	COUTS	2
COR CH	112	COR CH	112
CHEST ABD PELVIS	128	CHEST ABD PELVIS	128
MIP CH	129	MIP CH	129
SAG CH	134	SAG CH	134
SUPER D/THINS	318	SUPER D/THINS	318
SAG ABP	109		

Results

After following these methods for many weeks, we observed three recurring themes:

Theme 1 *Complete Series Missing*. Sometimes series (in their entirety) are not found on the receiving device. Most frequently, these "losses" are benign: some series are often purposefully excluded from sending (errors/artifacts, localizers, thin slices). It is important to speak with the technologists to understand which types of series are never sent so that these "series drops" are expected. For example, in Fig. 4, the series "SAG ABP" was present on the scanner but was not present in PACS.

Certain series, such as "PhoenixZIPReport," were never sent to PACS due to vendor incompatibility (lack of DICOM SR support, Fig. 5). Many times the total image count on the two devices would only differ by the number of images in a series that is not compatible with the receiving device. Although these instances were not detrimental for our workflow, we documented all of them to address separately, as data transfer policies.

Theme 2 *Added Series on the Receiving Device*. Similarly, new series (not present on the sending device) may be found on the receiving device: post-processed 3D reconstructions being the most obvious example (Fig. 6). This type of occurrence explains how the receiving device can sometimes have a higher image count than the sending device.

Theme 3 *Partial loss*. Finally, we ran into the case where series are only partially sent to PACS [7, 10]. For example, a series on the sending device would have a total image count of 50 images, but only 49 of these 50 images were found on the receiving device. Technologists confirmed that while they may intentionally drop an entire series, they would never remove a subset; either a series is sent in its entirety or it is not sent at all. Thus, it was a partial transfer pattern that flagged non-planned, genuine image failures. Figure 7 illustrates this pattern in our audit: the series "AAScout" had 257 images on the sending device, but only 256 images were on the receiving device. Since the images were disappearing randomly and at random locations, partial losses are virtually impossible to detect on the receiving end; only the comparison with the sending device image counts would make these losses visible.

As a result, the partial loss pattern was deemed the most dangerous and the most characteristic for all types of technology-driven failures described earlier. It has therefore become the center of our attention (Table 1, bottom line). To investigate these losses at our facility, we conducted a 2-

Accession #	Scanner Image Count	PACS Image Count	Affected Series
Confidential	1538	1526	PhoenixZIPReport(12)
	244	238	PhoenixZIPReport(6)
	232	222	PhoenixZIPReport(10)
	1199	1188	PhoenixZIPReport(11)
	187	181	PhoenixZIPReport(6)
	230	221	PhoenixZIPReport(9)
	1541	1523	PhoenixZIPReport(18)
	248	241	PhoenixZIPReport(7)

Fig. 5 Example of series not sent to receiving device due to vendor incompatibility

Scanner		PACS	
Descr.	*Images*	*Descr.*	*Images*
PosDisp: [4] SWI_Images	1	PosDisp: [4] SWI_Images	1
PosDisp: [12] SWI_Images	1	PosDisp: [12] SWI_Images	1
PosDisp: [2] Sag T1 SE flip angles	1	PosDisp: [2] Sag T1 SE flip angles	1
PosDisp: [18] TOF 3D Cow_MIP_TRA	1	PosDisp: [18] TOF 3D Cow_MIP_TRA	1
PosDisp: [15] TOF 3D COW	1	PosDisp: [15] TOF 3D COW	1
localizer	12	localizer	12
Sag T1 SE flip angles 75-170	30	Sag T1 SE flip angles 75-170	30
Dif_ADC	30	Dif_ADC	30
Dif_TRACEW	30	Dif_TRACEW	30
Dif_EXP	30	Dif_EXP	30
Dif_FA	30	Dif_FA	30
Dif_ColFA	30	Dif_ColFA	30
AX FLAIR T2	30	AX FLAIR T2	30
AX T2	30	AX T2	30
		3D COW	37
		3D COW	37
		3D ANT CIRC	37
		3D R ANT	37
		3D L ANT	37
		3D POST HORIZ	37
		3D POST VERT	37

Fig. 6 Example of added series in PACS

month image count audit, using our CT and MRI scanners. The losses we identified lead us to a few DICOM routers, failing at peak times (overload scenario discussed earlier). Our PACS team had to work with the router vendor to pinpoint the problem and to deploy a fix. Rechecking image counts after the fix was installed confirmed that all partial losses were eliminated.

Discussion

Our “image leaks” audit described above took two full months. All image count comparisons during this time were done in collaboration with modality technologists and managers to investigate all cases with confirmed image count discrepancies. Any cases that needed elevated attention and action, such as partial losses, were

Scanner		PACS	
Descr.	*Images*	*Descr.*	*Images*
localizer	3	localizer	3
T1 Sag	23	T1 Sag	23
AX FLAIR T2	25	AX FLAIR T2	25
T1 AX	26	T1 AX	26
ADC	26	ADC	26
DWI	26	DWI	26
LOWB	26	LOWB	26
EXP	26	EXP	26
FA	26	FA	26
FA color	26	FA color	26
Susc Ax	26	Susc Ax	26
AX T2 POST	26	AX T2 POST	26
T1 AX POST GAD	26	T1 AX POST GAD	26
T1 AX POST GAD_SUB	26	T1 AX POST GAD_SUB	26
AX DIFFUSION	28	AX DIFFUSION	28
AX POST MPRAGE RFMT	186	AX POST MPRAGE RFMT	186
Sag MPRAGE POST	192	Sag MPRAGE POST	192
AAScout	257	AAScout	256
COR POST MPRAGE RFMT	442	COR POST MPRAGE RFMT	442

Fig. 7 Example of a partial image loss

highlighted in red for easy identification (Table 1, bottom row).

Although this effort did take a considerable amount of time, it was definitely worth doing:

- We did identify a number of hidden image leaks due to faulty router software, vendor incompatibilities, and uses of poorly supported options (such as compression, leading to collateral image leaks). The scope and variety of these problems came as a total surprise to most of us—we all used to put more trust into "digital" than we should have. But once the issues were identified, they were promptly addressed and eliminated. Solid, reliable, and loss-proof image flow was the main outcome of our project.

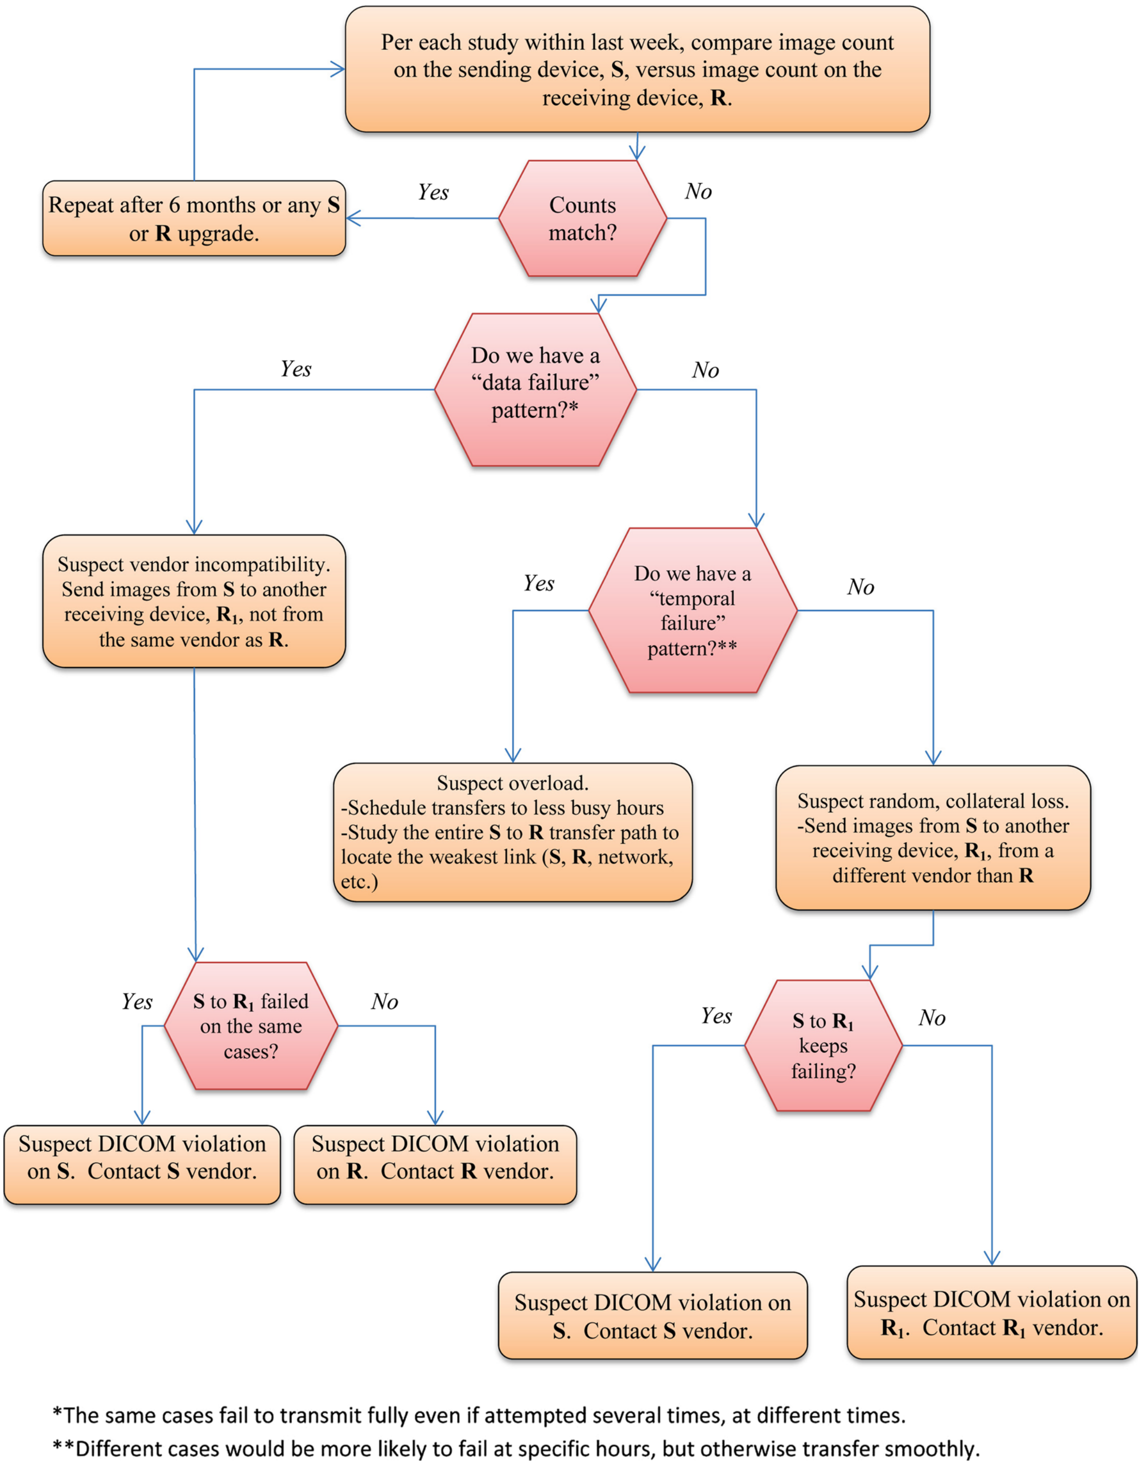

Fig. 8 Lost image audit chart

- In a similar way, we identified and documented a multitude of intentionally avoided image transmission, done for diverse reasons, and often in rather inconsistent ways. This resulted in new discussions and policies on what data should and should not be sent, how canceled series should be identified, and how they should be labeled.
- We enhanced our imaging routers with additional count-checking logic, to automatically detect any partial series transfers and to alarm the users. Manual image audits are too laborious to be run all the time. Modern imaging systems should be able to monitor their data integrity automatically, alerting the users of any problems in real time.
- We reported a number of cross-compatibility problems to DICOM vendors. The fact that many popular DICOM vendors still cannot get 100 % compatible image transfers raises the concern that too many hospitals have invisibly leaking cross-vendor imaging.

Our leak-auditing experience is summarized in Fig. 8, and we strongly recommend this approach to all hospitals and digital imaging facilities. Taking these appropriate measures will not only shed light on the fact that there is a problem occurring but it will also lead you towards a solution.

Finally, our main recommendation to all the hospitals would be to not assume that "digital" or "modern" means "error-free." Error-resistance should not be assumed—it should be built into your practices and into your devices. Lost images are a serious matter and need to be handled professionally and systematically.

Conclusions

It is clear that the migration from film to digital imaging has not completely terminated image loss issues. Imaging devices need to be checked to identify any image loss problems happening under the surface. Hospitals that have evidence of an image loss problem need to identify the source of the problem so that a fix can be implemented as quickly as possible. Despite any fixes that are installed to remedy the problem, devices should still be checked every 6 months or after any major upgrade to ensure that nothing is going awry. Lost images can affect patient diagnoses and it is critical to stop this problem from happening.

It is also worthwhile to solidify all image transfer policies to ensure their consistency. Images should not just be removed intentionally or not; there should be a set of well-defined rules to control image transactions and to ensure data integrity. Without any doubt, medical image loss should not be taken lightly.

References

1. Bassett M (2012): GE PACS system losing medical images. Retrieved June 2014, from Fierce Medical Imaging: http://www.fiercemedicalimaging.com/story/ge-pacs-system-losing-medical-images/2012-11-12
2. Huang HK: PACS and Imaging Informatics: Basic Principles and Applications. John Wiley & Sons, Hoboken, 2010
3. Mike McBiles AK (2000): Coping with PACS downtime in digital radiology. J Digit Imaging, 136–142
4. Oosterwijk H (2004) The DICOM standard, overview and characteristics . Retrieved July 2014, from Ringholm Whitepaper: http://www.ringholm.de/docs/02010_en.htm
5. Pianykh OS (2013 (second edition)). Digital Imaging and Communications in Medicine (DICOM). Springer Verlag
6. Pianykh O S (2014): Digital Image Quality in Medicine . Springer Verlag
7. Ridley EL (2012): PACS viewing software may not show all image slices. Retrieved June 2014, from Aunt Minnie: http://www.auntminnie.com/index.aspx?sec=ser&sub=def&pag=dis&ItemID=100897
8. Shaw G (2001): A clinician's guide to digital X-ray systems. J R Soc Med, 391–395
9. Strickland, N. H. (2000). PACS (picture archiving and communication systems): filmless radiology. Arch Dis Child, 82–86
10. Wallis JW (2004) Nuclear PACS: Problems, Solutions. Decis Imaging Technol, 21–25