VISAGE IMAGING®

Can you? Visage can.

Volume 1: Speed

12625 High Bluff Drive, Suite 205, San Diego, CA 92130
(888) 998-4724 | www.visageimaging.com

Introduction

Visage's General Manager, North America, Brad Levin, has recently published a series of new blog posts, *Can you? Visage can.*, with the exciting first topic focused on Speed. We've aggregated the individual blog posts to create this whitepaper: **Can you? Visage can. | Volume 1: Speed**.

It's all about differentiation and value, and Visage offers distinction in a sea of PACS sameness. While Visage customers appreciate the unique value of Visage 7, we're trying to make it easier for outsiders to get insights into why Visage is truly a generation ahead. Get prepared, we're going to give you a peek into Visage 7 that goes beyond features lists and RFP responses, cutting to the chase, while filtering out jargon, buzz and hype.

Speed is a topic that is near and dear to Visage. I think Dr. David Clunie said it best a few years ago at SIIM, when he commented that in Imaging, the top three requirements for any software application are, "Make it fast. Make it fast. Make it fast." He was and remains spot on. Speed is central to Visage 7, but it's actually much more than you're thinking. More than just the rapid display of images, speed manifests itself in *Image Display, 3D/Advanced Visualization, Digital Breast Tomosynthesis (DBT), In-Viewer Workflow, At Home/Remote-Reading, Regional Enterprise Imaging, Thin-Slice Reading, Priors, Study Ingest Performance, Multiple-Workflows and Interoperability*. While other vendor solutions use trickery, restrictions and workarounds to 'simulate' speed to the end user, Visage 7 transcends speed expectations across the entire platform.

Let me explain.

12625 High Bluff Drive, Suite 205, San Diego, CA 92130
(888) 998-4724 | www.visageimaging.com

VISAGE IMAGING®

Chapter 1: Image Display

The most prominent, visceral indication of speed is image display. For years in PACS, there has been an expectation that the time "to first image" needed to be 2-3 seconds. Some vendors met that requirement, most did not. And even for those that did, was the "first image" sufficient? What about the entire dataset? If you wanted to view a specific image or series of images, users often had to wait for the entire dataset(s) to load. Some solutions implemented local/auto-caching to the client/workstation, to trick the user into believing that the viewer was fast. But if the current/prior dataset was not local, you still had to wait for the data to download. And for double-digit gigabyte sized datasets, that's an impossibly long wait. Given these limitations, time to first image is a false metric of speed.

Visage 7 never caches images to the local disk, and yet, is incredibly fast. How fast? Visage tracks a key performance indicator (KPI) for our largest customer in the world, with a single production Visage 7 Backend Server (also replicated in a remote datacenter for High Availability), currently processing 2.8M annual new studies (growing to 3.8M across the system) as well as millions of priors, across dozens of distributed, regional hospitals. Their median display time to view the entire study, across all imaging modalities, across all study sizes, is 1.1s. While there are outlier studies that may take longer than 1.1s, this KPI is based on production Visage 7 customer data, at massive scale, spanning all diagnostic and clinical workflows, across all sized studies. Let me repeat: 1.1s.

This is no myth, but verifiable truth.

Chapter 2: 3D/Advanced Visualization

In Chapter 1 of Speed, we discussed why Image Display is the immediate reflex most people think about when considering the speed of enterprise imaging. While it may begin with Image Display, it does not end there. In this post, we'll continue the conversation into Chapter 2: 3D/Advanced Visualization.

To this day, many legacy PACS cannot display advanced visualization reconstructions or VR images, within their core viewer(s). The user must launch a third-party application to view additional multi-dimensional images. If the images are not on the third-party server, the user has to wait for the data to be transferred, be processed, and displayed. In most cases this delay is intolerable, so much so that radiologists ask that static secondary captures be generated at the modality console or workstation. And even for those rare vendors that do support advanced visualization in their viewers, often the advanced visualization tools are a plugin, with a different workflow and look and feel, slowing the user down.

Visage 7 pre-processes cross-sectional image data at the server (CT, MR, PET/CT, PET/MR, DBT, etc.), and when the user requests current and prior(s) for display on-demand, all of the images are immediately available for multi-dimensional display in the protocol.

12625 High Bluff Drive, Suite 205, San Diego, CA 92130
(888) 998-4724 | www.visageimaging.com

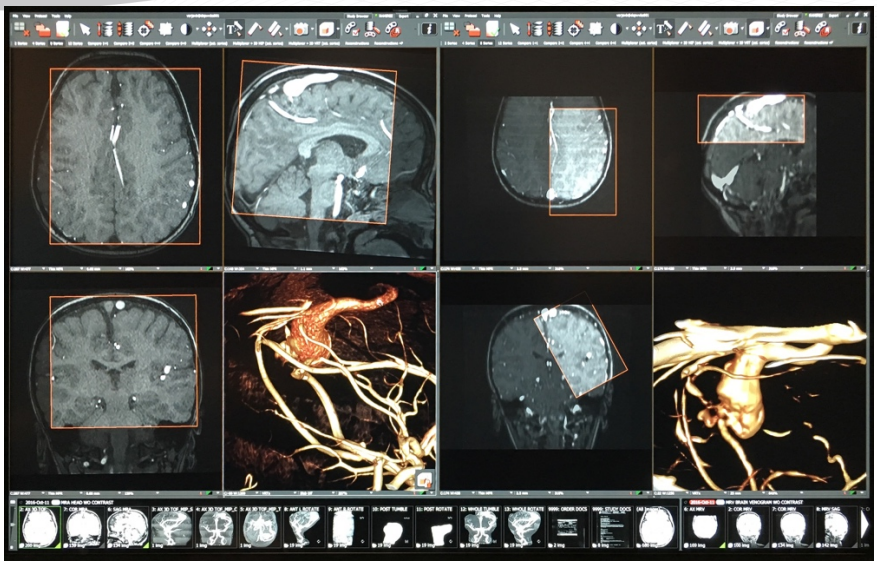

Visage 7: MRA and MRV

100% of Visage's advanced visualization features are native Visage code. With Visage, the haves and the have-nots are a thing of the past. Referring physicians who require radiologist caliber tools, including 3D/advanced visualization, have access without additional cost, limitation or any performance impact whatsoever.

Chapter 3: DBT

In the first two chapters of Speed, we discussed the topics of Image Display and 3D/Advanced Visualization. Many may think that speed of image display and having access to 3D views are most important; however, the size of imaging studies continues to grow and grow. A few years ago, cross sectional imaging made its way to mammography and far too many imaging organizations continue to struggle with their infrastructure, workflow, and interpretation.

How important is speed? Not too long ago we were at a demo at a large outpatient imaging chain, when the next radiologist on the schedule sat down for their demo. She immediately said, "Is it ok if I drive?" She clicked on the study and immediately the DBT painted the viewports of the multi-headed display. She froze with delight, literally. Without saying a word to anyone, she 'dropped' the mouse and shuffled out of the room in absolute amazement. When she came back a few moments later, all of us were puzzled by her reaction. She paused and said, "You have to excuse me, as I've been struggling with DBT for a long, long time. What I just saw impressed me so much I just had to leave and compose myself. Now, let me do that again!"

With that experience as the context for what is possible, in this post, we'll continue the conversation with Chapter 3: DBT (Digital Breast Tomosynthesis, DBT or "tomo").

12625 High Bluff Drive, Suite 205, San Diego, CA 92130
(888) 998-4724 | www.visageimaging.com

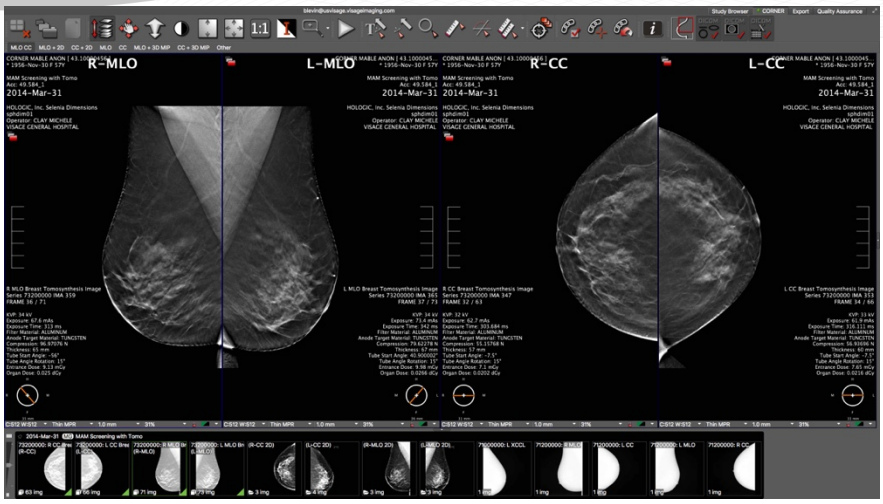

With Visage 7, DBT is just another modality, no separate module or workstation required

Even today, many legacy PACS do not support DBT at all. And those legacy PACS that do support tomo do so with significant limitations. How so? Speed of image display, restricted access to priors, the need for a dedicated DBT archive, the need for dedicated DBT workstations, the need for significant network augmentation, the need for workflow modification (e.g., perform procedures in the AM, move priors all other hours in advance of the next day) are all limitations we frequently hear about regarding legacy PACS support of tomo.

For Visage 7, DBT is simply another modality, requiring no workarounds or dedicated workflows whatsoever. Visage 7 supports the display of DBT along with DBT priors and multimodality presentation of all breast imaging studies (MG, DBT, Breast MR, Breast US, AWBUS), or any modality for that matter, from both Windows or macOS clients. Due to these capabilities, Visage 7 is able to eliminate the need for expensive, restrictive, dedicated breast imaging workstations. Despite moving upwards of tens of gigabytes of data for current and prior requirements, Visage 7 is able to display the required data incredibly fast for interpretation, local or remote. How many enterprise imaging platforms are able to support remote DBT reading at scale and eliminate dedicated workstations? We're not aware of any system.....other than Visage 7.

Chapter 4: In-Viewer Workflow

In our last chapter, we discussed how DBT ("tomo") has paralyzed many imaging organizations due to its massive files, but for Visage, tomo is just another modality. What we have learned in practice is that speed is relative. In this post, we'll explain what we mean, by continuing the discussion with Chapter 4: In-Viewer Workflow.

Many legacy PACS do not have reliable, usable hanging protocols. So much so, many institutions have given up trying to improve them. We know that historically the majority of unproductive time spent during legacy PACS radiologist interpretation is to setup a case in order to begin "reading". Finding 'relevant' priors, dragging and dropping images/series/studies, panning, zooming, linking, W/L, etc. This is incredibly time consuming, and yet, for many PACS, this is required prior to reading every new study. What does it matter if the image viewer is able to display the first image in 2-3 seconds, if the reading radiologist has to spend tens of seconds to minutes, preparing the studies prior to starting interpretation? What does it matter if the image viewer is able to display the first image in 2-3 seconds, if the user has to travel to a specific workstation, which may be local (via walking distance) or remote (travel by vehicle), to use the appropriate clinical tools or access the relevant priors they need? As mentioned, speed is relative.

12625 High Bluff Drive, Suite 205, San Diego, CA 92130
(888) 998-4724 | www.visageimaging.com

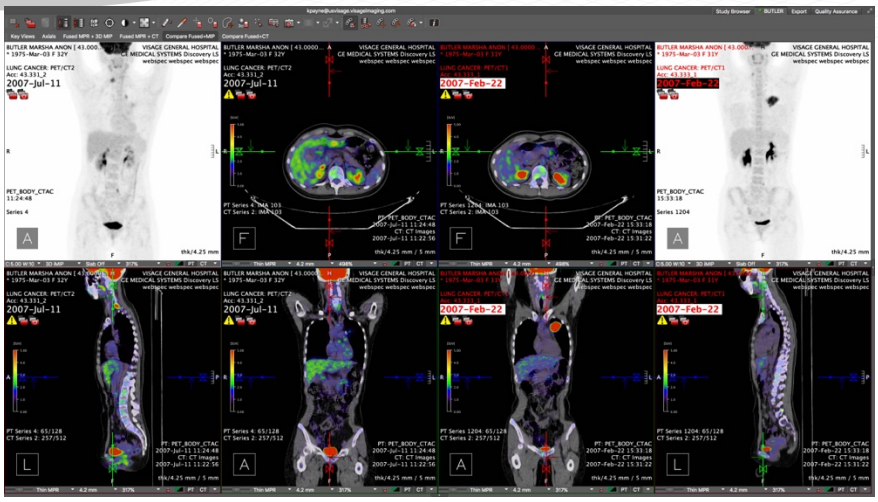

Visage 7: PET/CT Lesion Tracking | Fused, Registered, Current/Prior with Linked Rotating MIP

Visage 7 supports the industry's most powerful, most sophisticated protocols, for all modalities (including medical multimedia objects and non-DICOM) including advanced visualization/multidimensional content for current/priors. Utilizing sophisticated matching algorithms and intelligent layouts, Visage 7 automatically displays any combination of current/prior studies, regardless of size or modality, using Auto Prior rules and native thin-client advanced visualization tools. This does away with unnecessary clicks, drag and drop, and switching between tools, thereby significantly streamlining interpretation workflow for radiologists.

These differentiated Visage 7 capabilities are described as In-Viewer Workflow, and allow interpretation of more cases, in less time, with the utmost quality. This practically means that when radiologists read with Visage 7, they don't need to "setup" the stud(ies) prior to interpretation, and can simply launch and read, over and over and over again. Optimal In-Viewer Workflow eliminates unproductive interpretation time, eliminates distraction/frustration, and allows radiologists to get into a "Zen-like" interpretation state for highly productive, high quality reading. Now that is truly fast.

Chapter 5: At-Home/Remote Reading

In our last chapter, we discussed how In-Viewer Workflow differentiates truly fast viewers from viewer pretenders, exploring the elusive "start of interpretation" that many lessor viewers struggle overcoming. These are topics that most don't consider when initially thinking about speed. Speed is an emotional topic, you know it when you see it, but there's even more to it than "surface" speed. And that brings us to our next topic in Chapter 5: At-Home/Remote Reading.

Legacy PACS performance for users outside of the LAN has historically been dreadfully poor. We've encountered customers who have told us flat out their PACS is unusable remotely. Sometimes that performance is attributed to needing to run their viewer(s) over a VPN. Other times we are told that broadband speeds are simply unsupported. One major PACS vendor has gone so far as to place a disclaimer on their website requiring 1 Gigabit/s connections between client and server for adequate performance. I'll also share a real world story where a customer told Visage, "I know you're great at advanced visualization remotely, but can your viewer support batch reading of plain films with priors, on-demand, back to back to back? We can't do that with our current viewer at all." Visage proved the remote batch reading performance and more, and is now live across their multi-hospital enterprise.

Visage 7 performs nearly identically, whether local or remote, over VPN, using Windows or macOS, even over broadband speeds. For primary reading, Visage recommends ≥6 Mbps bandwidth, with <80 ms of latency. For other clinical use, Visage recommends ≥2 Mbps bandwidth, with <200 ms of latency. Visage customers tell us that there's no practical difference in

12625 High Bluff Drive, Suite 205, San Diego, CA 92130
(888) 998-4724 | www.visageimaging.com

reading remotely. One large health system who could not read remotely on their prior PACS, first piloted at-home reading on Visage 7, and then quickly expanded the initiative. Why? They found that their radiologists reading at home were even more productive than in the hospital setting with higher quality, and the quality of life benefits were having a big positive impact on their mental health. With all of the stresses of radiology, enabling more of their radiologists to read from home more frequently was a win-win-win for the institution, for the radiologists and for the patients they care for.

Yes, speed is an enabler that cannot be underestimated.

Chapter 6: Regional Enterprise Imaging

In our last chapter, we shared how even in 2017, most diagnostic viewers are not ambidextrous. Meaning, unlike Visage 7, most diagnostic viewers do not perform as well at-home for remote reading as they do local at the hospital or imaging center. Given how mobile today's healthcare professionals are, that's an unconscionable limitation and should not be tolerated for today's standard of care.

Over the past several years, consolidation has truly transformed the healthcare landscape. Community hospitals have been acquired and merged to create large regional networks, many with academic medical centers as their flagships. Outpatient imaging chains continue to grow, many via merger and acquisition. Additionally, many radiology groups have been acquired into regional and national networks, creating massive radiology organizations of hundreds of radiologists, performing millions of new studies with hybrid models of outsourced, onsite and remote teleradiology reading services. But when we speak with many of these growing regional and national organizations (inpatient/outpatient/hybrid) their technology stack is all too frequently a heterogeneous mix of legacy solutions. It's a mishmash of disparate technologies that offer no consistency, no foundation for streamlined operations and economies of scale you'd otherwise expect to see in a large network. When we talk to these groups, many don't think that a single, regional or national enterprise imaging platform is even remotely possible.....until they learn about the capabilities of Visage 7.

Since the early days of PACS, it has usually been a foregone conclusion that there was a 1:1 relationship between the # of PACS "server instances" and the # of locations needing to be served. Six hospitals? Six different local instances of the PACS. Much like Goldilocks, if the study mix was 'just right', or if the local volume was 'not too big', or if the distance between facilities was 'not too far', or if the WAN connection was 'stable with adequate capacity', perhaps you could get away with supporting 2 or slightly more locations with one PACS instance. Of course there would be other questions, such as cross-site reading, enterprise access to priors, performance implications, and other limitations. As an aside, it's quite common for Visage to visit multi-location health systems that have standardized on the same vendor PACS, with each location having their own PACS servers. Almost universally the local hospitals are islands of PACS, with zero or restricted access to the other PACS, even though they are all in the same system. Sometimes 'manual' access is a workaround provided through swivel chair workflow (e.g., multiple dedicated workstations for each facility) or through a manual KVM switch. This is not a regional enterprise imaging system. This is a tinkertoy implementation.

With Visage 7, regional or even national implementations are reality, and at significant scale.

We often discuss a regional customer example as a speed comparison to multiple legacy PACS. Our Visage 7 customer reads for several large hospitals, using two of the most widely used legacy PACS in the hospitals, but remotely their radiologists read using Visage 7. Both of the legacy PACS were LAN-connected at 1 Gbps, while Visage 7 was WAN-connected at 5 Mbps. With every study type, even though the network was 200X slower, Visage 7 opened the studies faster than either of the other legacy PACS. And... in each of the legacy PACS, the data was already cached on the local workstation! Similarly, another one of our enterprise customers performed interpretation tests from one of their hospitals on the east coast of the US, connected to their datacenter in the upper midwest, a mere 1,350 miles away. It was consistently faster, in fact dramatically faster to display images in Visage 7, streamed from 1,350 miles away, then it was to display the same images in the local legacy

12625 High Bluff Drive, Suite 205, San Diego, CA 92130
(888) 998-4724 | www.visageimaging.com

PACS. For the first time in the history of this nationally distributed enterprise, their radiologists will soon be able to interpret imaging studies and have access to priors from anywhere their patients present in the system.

Chapter 7: Thin-Slice Reading

In our last chapter, we shared how regional enterprise imaging is redefining reading relationships. Consolidation is transforming the logo landscape, but behind the scenes the technology is more often a mismatched quilt of this and that. You'll get your images read, but not without a lot of moving parts, manual intervention, and sweating administrators. We elevate "mediocore" to "high-quality" across the imaging landscape and truly represent breakthrough, enterprise-class solutions. Everyone is trying their best to provide a consistent, quality product for referring physicians, but with variable underlying technologies that's easier said than done. With Visage, we've proven and delivered what others have said isn't possible. There's lots of vendors that do plenty of talking, Visage prefers getting things done….with differentiation.

The historical and current state of legacy PACS does not support thin-slice reading due to a number of factors, including performance and speed. This lack of capability has largely been industry-wide, despite massive improvements in slice count, speed and capability of thin-slice modalities, such as CT and MR. Yes, despite the slice wars in CT, legacy PACS has never caught up. As a workaround, technologists send thick-slices to PACS and secondary capture reconstructions from the console/workstations. This process is time consuming for technologists and if the radiologist doesn't like the view that was captured, new secondary captures are often requested. This delays the radiologist interpretation, delays the technologist from scanning patients, and negatively impacts patient care.

Other than Visage 7, Visage is unaware of any commercial PACS or enterprise viewing platform that supports thin-slice reading at scale.

In stark contrast, Visage 7 natively supports thin-slice diagnostic reading and clinical access at massive scale, without any performance degradation or client-side dependencies, for all users. Visage 7 customers explain that reading thin-slice image data allows them to read faster, with greater clinical accuracy. Why? Because they can dynamically see clinical abnormalities in the thin-slice data that are otherwise obscured in thick-slice and secondary capture reconstructions.

We've worked with a range of Visage 7 users around the globe that have tackled thin-slice reading with different approaches. Some have gone thin-slice right out of the gate for all cross-sectional imaging. Others have targeted thin-slice reading just for CT, or select studies. Another group of Visage 7 users have not gone thin-slice initially, but have moved to thin-slice reading shortly after go-live. These users often tell us that once they have settled in reading on Visage 7, their reading habits that were influenced by lessor performing viewers tend to shift for the better. Their days of only scrolling through axial stacks and secondary captures created by others rapidly come to an end.

12625 High Bluff Drive, Suite 205, San Diego, CA 92130
(888) 998-4724 | www.visageimaging.com

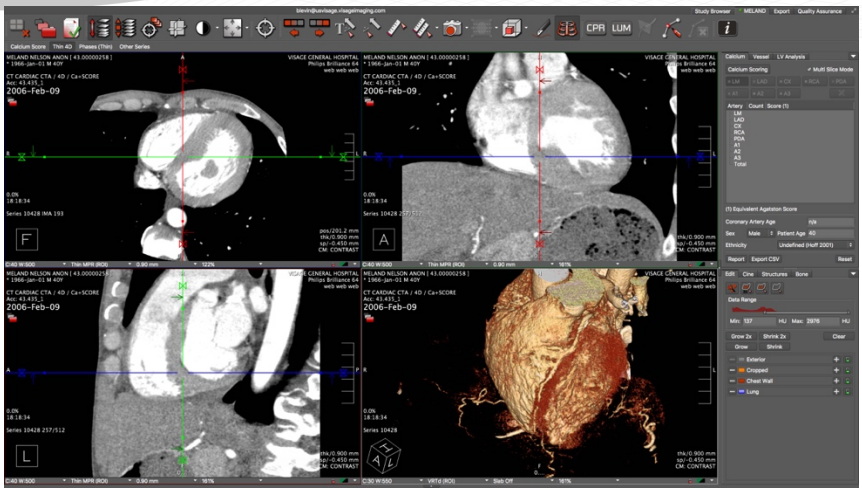

Visage 7: 10-Phase 4,000 Slice Cardiac CTA | Thin 4D Layout

With dynamic reconstructions at their fingertips and the ability to change slice thickness on the fly from the original axial data, radiologists quickly realize the power of being in full control, with no waiting. There's no hoping that the technologist picked the right views. There's no delay while new reconstructions are being captured and resent. And fundamentally, there's newfound confidence that they are seeing everything that the scanner can provide. Radiologists tell us that patient care is significantly improved because reading off of the thin-slice data with Visage 7, they can identify lesions and abnormalities that they simply could not see previously. This begs repeating - Visage 7 customers explain that reading thin-slice image data allows them to read faster, with greater clinical accuracy.

Patient care is significantly improved. Visage 7 uniquely allows customers to take full advantage of the thin-slice images generated from today's most advanced thin-slice scanners, with incredible speed. This is irrefutable - all other PACS do not.

Chapter 8: Priors

In our last chapter, we discussed how thin-slice reading using Visage 7 enables radiologists to read faster and to provide the best possible patient care. While in recent years thousands of institutions have spent tens of millions of dollars each on the latest thin-slice scanners, their viewers and diagnostic workstations only support reading "thicks" and secondary capture reconstructions created in the hope they are acceptable to the reading radiologist. These legacy viewers don't support thin-slice reading and choke on the data. And by choke, I mean they just won't work.

No one wants to put a speed limiter on a sports car, so why have healthcare institutions put a "care-limiter" on their significant scanner investments? It is because traditional PACS technology has not kept pace, but there is a better way…with Visage.

Only Visage 7 enables institutions to get the greatest value from their scanners, enabling radiologists to read faster, with greater clinical accuracy. If one considers thin-slice reading to be cutting edge for primary interpretation, how about access to prior imaging? Pretty basic, wouldn't you think? Think again.

Most legacy PACS have performance limitations that extend to prior exams.

Access to prior imaging is a more widespread issue than is commonly discussed, with priors frequently not available at the time of interpretation. In fact, at SIIM 2017 there was a presentation I observed by one of healthcare's most prominent institutions that discussed how access to priors remains a systemic enterprise challenge. They presented analytics research

12625 High Bluff Drive, Suite 205, San Diego, CA 92130
(888) 998-4724 | www.visageimaging.com

that "in general, comparison exams used were under 5 days old for the average exam interpretation". On most days, this institution speaks about their innovative thinking and broad vision, but on this morning, they spoke about their radiologists having difficulty accessing priors during interpretation. It was refreshing to hear their challenges openly presented, because we hear about these challenges when we speak with institutions that use a single PACS or multiple PACS, across one location or multiple locations.

Multiple PACS certainly complicates prior availability. This is probably not surprising, but a friend who has worked at numerous high profile institutions shared with me that in his experience radiologists reading in one organizational PACS will not search for prior images in a different organizational PACS, even if they know priors are available, because the effort is too disruptive to their reading workflow. When prior imaging is not readily available, radiologists may seek out prior reports if those are easily accessible, may delay the interpretation until prior images/reports are made available by other support staff, or may dictate the case without comparison.

What is well known is that even if prior images are available, many legacy PACS cannot support the rapid display of current and (all required) prior exams for reading. For example, there is no denying that the display of today's multi-thousand slice CT, compared to multiple recent multi-thousand slice CTs cause challenges for most legacy PACS. It's a widespread problem. And as noted, if performance or availability is an issue, radiologists are relegated to looking at the current exam and referring to prior reports (if available), instead of prior images.

Sometimes even if prior images are available in legacy systems, inadequate relevancy rules get applied and the result is a set of displayed priors that leave the radiologist wanting more. What if prior(s) are known to exist by the radiologist and unavailable for display because they didn't match the rule? Will the radiologist wait for the prior(s) to be retrieved for review? Possibly depending on the perceived importance of the prior(s), but from what we hear from both hospital and ambulatory based radiologists, waiting is unlikely, which results in a poorer quality report. This is highly disruptive and frustrating for the reading radiologist who knows prior image(s) exist, but can't view them in a timely fashion, and report quality ends up being sacrificed under the weight of unread studies to be reported.

Another pertinent example about legacy PACS performance of priors is related to digital mammography and tomo (MG/DBT). We've been to numerous imaging organizations who struggle under the performance implications of the delivery and display of multiple gigabyte MG/DBT priors. Speaking of limiters, I've met with imaging organizations that have capped the number of supported MG/DBT priors artificially due to performance problems related to display of current/priors, pre-fetch of priors, study routing and all-of-the-above. One of those organizations I met with had a standing rule that limited the maximum number of priors delivered to two priors for every current exam. Two. This is contrary to the practice of breast imagers who expect to be able to view and compare numerous priors (if available) during interpretation. For some organizations, the network impact is so significant, patient scans can only be performed in the morning, and the balance of the day is spent managing the data burden of interpretation and *in anticipation of the next day's schedule*. Quality is sacrificed. Radiologists are slowed down, interrupted and patient care is negatively impacted.

In stark contrast, Visage 7 has an all priors philosophy.

Quality and speed are never sacrificed with Visage 7, because in addition to Auto Prior rules for current/prior display, all priors are always immediately available for display. The reality is relevancy rules are adequate only up until a radiologist runs into a case where they know a prior is required and available, that didn't match the relevancy rule. Visage is of the opinion that the ultimate arbiter of relevancy is the mind of the reading radiologist, for the specific study being read. Each and every study is different in the context of the given patient, and fundamentally, rules are fallible. Radiologists deserve to have immediate access to all priors, whenever and wherever required, in order to deliver optimal patient care. That is the standard Visage 7 provides, for each and every interpretation.

12625 High Bluff Drive, Suite 205, San Diego, CA 92130
(888) 998-4724 | www.visageimaging.com

Chapter 9: Study Ingest Performance

In our last chapter, we discussed how access to prior exams at the point of interpretation is a continuing challenge for many institutions. There's numerous reasons for that struggle, from the inability to access the priors from their originating institutional archive(s), to the massive aggregate size of multiple priors (particularly for multi-slice CT and DBT studies). In contrast, Visage 7 has an all priors philosophy, without any sacrifice of quality or speed. In addition to Auto Prior rules for current/prior display, all priors in Visage 7 are always immediately available for display. Radiologists deserve to have immediate access to all priors, whenever and wherever required, in order to deliver optimal patient care. That is the standard Visage 7 provides, for each and every interpretation. Just last week Visage 7 went live at an esteemed regional institution, and for the first time in their history, radiologists have immediate, routine access to all patient priors at the time of interpretation. We'll explore this story in a future blog post, but once again speed is the enabler.

Legacy PACS routinely struggle with data volumes. In fact, our customers (prior to moving to Visage) frequently tell us that they are drowning in their data, struggling to keep their "heads above water". Daily, one large outpatient imaging center chain was ~3,000 studies behind with their legacy PACS. Their PACS server(s) could not keep up with their volume. In a few short weeks after moving to Visage, their daily backlog was eliminated for good, and radiologists were reading in real-time for the first time in the history of the 1M+ new annual study practice.

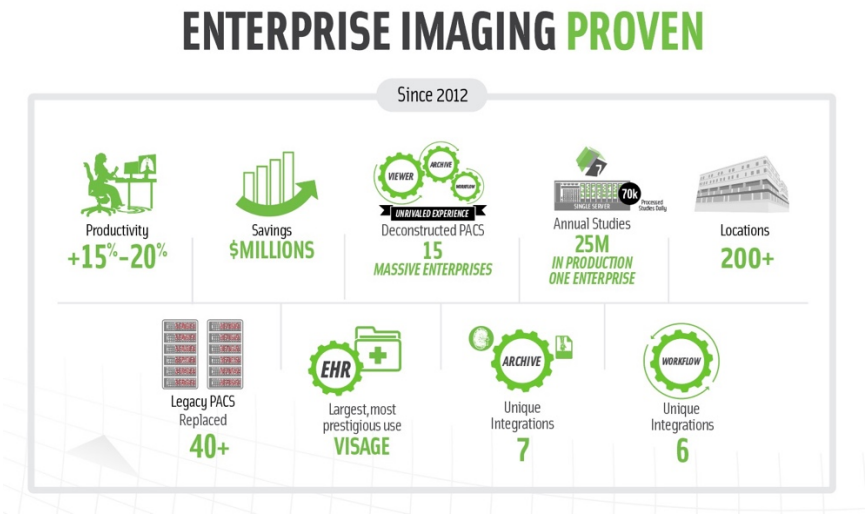

Visage 7 is scaled to support massive ingest performance. Yet another production Visage 7 customer is processing 25M current and prior studies through a single Visage 7 Backend Server, with a daily run rate of 75,000 studies. Visage has also experienced production performance rates in excess of hundreds of millions of studies/year during peak load times, without missing a beat.

Additionally, Visage has customers that virtualize their Visage 7 Backend Server using VMware. One large customer processes 12M current and prior studies annually across their virtualized Visage 7 system. Remember, Visage 7 processes currents and all prior studies for patients, so they are all available for users when required for immediate display.

Visage has never encountered a PACS, viewer, archive or VNA that can even come close to the speed, ingest performance and scale of Visage 7!

12625 High Bluff Drive, Suite 205, San Diego, CA 92130
(888) 998-4724 | www.visageimaging.com

Chapter 10: Multiple Workflows

In our last chapter, we discussed how most people think about speed in terms of the front end (e.g., the viewer(s)), but don't adequately consider the speed of the backend (e.g., the server(s)).

A good way to look at speed is that they are two sides to the same imaging coin. If either is slow, it's going to slow your imaging organization down. If you're an organization with significant incoming volume, the slower the backend, the likelihood your radiologists will be working from behind is high. If both sides of the Imaging coin are slow, then you'll likely have a slew of problems organizationally that only get compounded with scale. One measure of backend speed is study ingest performance, or how quickly can the server(s) process incoming DICOM studies. Visage 7 is scaled to support massive ingest performance, supporting tens of millions of studies annually. It's also important to repeat that Visage has never encountered a PACS, viewer, archive or VNA that can even come close to the speed, ingest performance and scale of Visage 7.

Most legacy PACS architectures do not have streamlined architectures, meaning they require multiple servers, with a variety of functions, networked together in order to deliver the total requirements of the PACS. When institutions want to add (for example) additional locations, add multiple new modalities, add mobile imaging access, and support new -ologies, additional servers are typically required. These add to the architectural complexity that must be supported, but how does it impact speed?

In the late 1990s and early 2000s, legacy PACS were designed to support one workflow and one workflow only: interpretation (and archiving) of diagnostic images. The truth is that most of today's legacy PACS are still based on these architectures. Over time, additional capabilities were added such as web access for clinicians and 3D, but only through the addition of separate servers and licensed technology. And while since then legacy PACS has been upgraded to support even more workflows, most require their own servers, different modules, additional software, isolated workstations and third-party applications to fulfill the needs at hand. Each of these servers needs to be configured, managed, integrated, licensed and maintained at great expense and complexity to the institution.

The resulting system architecture is more often one of disorder than order, one of numerous servers and systems that are tied together with a variety of standard and proprietary integrations. How can one easily assess the speed of an individual workflow? It's not straightforward without doing a Proof of Concept, because each silo of functionality may demonstrate one "speed" in isolation, but as soon as it is tested with other required system components the speed slows down. Integrated testing must be done in order to provide an accurate reflection of actual production performance.

Here's four representative examples related to viewers:

1. **3D/Advanced Visualization.** If 3D is delivered via desktop integration to a legacy diagnostic workstation, or even a separate server-based system, it needs to get access to images. If the images aren't sent in advance, or if the 3D server is in a remote datacenter over a limited bandwidth connection, the end result is the radiologist is going to wait. What images are sent to the 3D system? Just the current study? How about priors that also require 3D? Having a separate workflow for 3D slows down primary interpretation.
2. **Mobile Access.** Very few, if any legacy PACS offer mobile capabilities as a core part of the PACS. It almost always is a separate server or stack of servers required for use. These server(s) need to be synchronized with the PACS in order to have the latest studies required for viewing, along with the latest changes to studies (edits, deletes, ADT updates, etc.). Does it have all studies? Likely not. What happens when the studies required for viewing are not on the server(s)? The end user will have to wait, or the stud(ies) may not be available for access.
3. **Universal Viewers.** With most legacy PACS, universal viewers provided by the PACS vendor have either been built in house, or more than likely are a separate third party viewer, or are a re-branded third party viewer. Even if they were built in house, they typically have been built using a separate technology stack than the legacy PACS. That means separate server(s) to support the universal viewer, even if it was built by the PACS vendor. Again, separate

12625 High Bluff Drive, Suite 205, San Diego, CA 92130
(888) 998-4724 | www.visageimaging.com

servers mean additional management overhead, and separate integration points to the PACS. Often times the universal viewer works great in the demo, but once it goes live in production and is integrated to the local or multiple PACS, performance takes a hit.

4. **Breast Imaging.** The vast majority of legacy PACS have inadequate solutions for breast imaging, plain and simple. The result is separate islands of breast imaging workstations, keypads and archives all dedicated to breast imaging. They take up desktop real estate in the reading rooms, generate a ton of heat, and cause radiologists to pivot from one workstation to another all day long. These workflow disruptions are an island from the rest of the PACS and cause headaches for imaging administrators. One of our customers estimated their lifecycle cost for a single breast imaging workstation at $250K. Once they went live with Visage 7, they were able to retire dozens and dozens of breast imaging workstations (and eliminate the need to purchase additional dedicated breast imaging workstations) saving the institution millions in cost savings and cost avoidance.

In contrast to these legacy solutions, Visage 7 performs with incredible speed while supporting a One Viewer philosophy.

Using a single Visage 7 Backend Server, all of the following workflows are supported at incredible scale: Diagnostic, Clinical, EHR, 3D, Mobile, Diagnostic Mobile, Cardiology, Breast Imaging, Specialties, Telerad, At Home/Remote Reading, Non-DICOM, Visible Light, Multimedia, HD Video, Audio, QA and DICOM Modality Worklist.

When all workflows are supported on the same ultrafast architecture, they are all incredibly fast and straightforward to support. Instead of an architecture of disorder, Visage 7 offers a streamlined architecture built for scale, availability and performance.

12625 High Bluff Drive, Suite 205, San Diego, CA 92130
(888) 998-4724 | www.visageimaging.com

Chapter 11: Interoperability

In our last chapter, we discussed how legacy solutions haven't handled multiple workflows very well, requiring third party software, additional licensing and servers required to support the additional requirements. In contrast, Visage 7's elegant architecture eliminates redundant dataflows, streamlining the enterprise imaging infrastructure.

In today's Imaging environment, interoperability reigns supreme. Radiologists don't read off of a series of displays in a vacuum, they read in an integrated cockpit. Surrounded by multiple large format, high resolution displays, multiple applications perform an orchestrated ballet in patient context with each and every new study launched. Within a 180 degree glance, the radiologist typically has visual access to their reading worklist, their reporting application, their diagnostic viewer, the EHR, as well as other diagnostic and clinical software applications.

Interoperability is usually perceived much broader than just the visual fitting of all of the different parts of the radiologist cockpit. In this post I'll be discussing both the desktop user experience and systems interoperability. This is really about the systems speaking "the same language". Whether it is for visual coordination or information exchange of workflow orchestration, it's all about knowing the common language and being able to learn the dialects quickly. Concerning interoperability and speed, it is about how long it takes an application to understand what it is being asked to do in the context of the cockpit and how long it takes for the application to react.

The speed of the diagnostic viewer has been addressed in detail over the prior 10 chapters of this series, but it's also important to recognize the speed and completeness of visual interoperability within the integrated reading cockpit. Fast access to irrelevant information is meaningless. Visage 7's ultrafast speed enables delivery of the information the user needs, when they need it. What I am speaking about is the speed of the first launch and subsequent launches of the diagnostic viewer; speed of closing; the speed of changes in patient context in the diagnostic viewer as directed by other integrated software applications; and the overall speed of the roundtrip from start to finish of an individual interpretation. Based on my travels visiting with numerous institutions, I've personally observed lengthy delays when the radiologist cockpit first boots up. Sometimes it takes only a few minutes, other times it shockingly takes 5-10 minutes, or sometimes even more. Due to these delays, radiologists develop interesting habits while they wait!

12625 High Bluff Drive, Suite 205, San Diego, CA 92130
(888) 998-4724 | www.visageimaging.com

In addition to being an ultrafast viewer, from our direct observations integrating to most of the major third-party applications in use today, Visage 7 is hands down one of the fastest, if not the fastest software application available. When observing the integrated reading cockpit, you can quickly assess which applications are more lethargic than others, which ones are more responsive, and which ones are downright slow.

As the Dalai Lama of PACS stated 10-years ago in The Dalai Laws of PACS, *VI. PACS should not get in your way*. For far too many organizations, their legacy viewers and software applications do not easily retain patient context, nor do they all support bi-directional integration. Some applications, including one major EHR manufacturer, only supports a fire-and-forget integration. Others have poor interrupted workflow. And when it comes to speed and productivity, if interoperability is sacrificed, workflow and productivity are bound to be negatively impacted –searching for information; opening this window and that window; verifying information to ensure they are on the right patient. I'm sorry, Dr. Dalai, these subpar solutions definitely get in the way of the reading radiologist. Not good.

In contrast, Visage 7 supports tight bi-directional integration, including to the most widely used EHR, which powers interoperability with both speed and completeness. Radiologists are able to maintain bi-directional, tab-based workflow, easily moving from patient to patient without having to hunt for the clinical information they require or worry about patient context getting out of sync. Robust interoperability improves radiologist confidence, gets them into a reading "zone", speeds interpretation and improves the quality of interpretation.

Another front end measure of interoperability relates to quality assurance (QA). Many enterprise viewers and VNAs do not support image QA at all, and require yet another third-party solution for QA. And for those that do support image QA, they don't do it well. Visage 7 supports image QA with the same viewer that also supports all major workflows. In contrast to many legacy PACS that require performing QA on an isolated PC (resulting in technologists that have to complete exams on one PC, and then jump to another PC to perform QA), Visage 7 seamlessly runs on the same PC with multiple applications, including workflow applications and the EHR. This simplifies and speeds technologist QA with an easy workflow that enables technologists to quickly complete their work and confidently move on, so they can spend more time with patients.

From a backend perspective, standards-based integration is absolutely key, leveraging the latest DICOM and HL7 standards, and IHE guidelines. The use of web services and emerging integration approaches such as FHIR are also important to support cutting edge initiatives. One of the most important standards, but only recently beginning to be implemented over the last few years is Imaging Object Change Management (IOCM). IOCM ensures that VNAs and PACS archives communicate automatically with enterprise viewers, such that QA changes automatically are sent to the VNA to stay synchronized. Without IOCM, proprietary approaches are also available. What is shocking is that far too many institutions have invested millions in VNAs, but are still manually making changes on the PACS/enterprise viewer and on their VNA. Really? Really. Talk about an impact to speed and quality, duplicate manual edits on the viewer and VNA are unacceptable in 2017, and a huge drain on enterprise imaging initiatives. Visage 7 fully supports a standards-based approach (HL7, DICOM, IHE, FHIR, IOCM, Web Services), for the most efficient, fastest, most complete enterprise imaging solutions.

Too many legacy vendors pay lip service to interoperability, but usually don't apply it, because they want to sell more products across their portfolio. For example, many PACS companies say they endorse standards, have open APIs and promote interoperability, but the reality is there are many companies that use terms like "single platform, etc." because they want to sell more stuff, or their products don't integrate well at all with other solutions. We've seen this time and again from legacy vendors that will say one thing pre-sales, but do another during implementation. Visage can confidently say that we support standards-based interoperability by way of our large Deconstructed PACS install base. Visage has proven the model, it works.

12625 High Bluff Drive, Suite 205, San Diego, CA 92130
(888) 998-4724 | www.visageimaging.com

Conclusion

And with that, we have now come to the close of the speed series of Visage Blog posts. We've repeatedly been posing the question, "Can you? Visage can." Hopefully you now have a better, more comprehensive understanding of speed in Imaging, and can make a detailed assessment about how your PACS/enterprise viewer(s) are performing. How do your systems stack up?

Speed is critically important across the entire Visage 7 Enterprise Imaging Platform: speed across the dataflows, speed launching Visage 7 independently, or launched from third-party applications in an orchestrated bi-directional interpretation cockpit, speed launching current/prior(s) local or remote, speed for any stud(ies) regardless of size or workflow need. Speed enables access to information, and information is what matters. Fundamentally, Visage 7 provides differentiated value on a scope and scale that competing solutions cannot approach.

Can you? Visage can.

Brad Levin
General Manager, North America
Visage Imaging

12625 High Bluff Drive, Suite 205, San Diego, CA 92130
(888) 998-4724 | www.visageimaging.com