



Integrated high-speed photo-image workflows with XMP metadata — Interview with Aaron Holm of Industrial Color

Aaron Holm

began his career in engineering audio systems for animated television productions at Nelvana Ltd. His experience spans over 15 years of creative production and technology innovation for the entertainment, media and technology industries. In the mid-1990s, Aaron worked with Getty, modeling digital video ingestion and distribution systems for stock footage. In 1998, Aaron joined Headline Media Group's national broadcasting division to direct IT operations and develop interaction television properties with partners such as Major League Baseball and the National Football League. In 2004, Aaron founded Markham Street Media, a company focused on the development and deployment of advanced digital asset management systems. In 2006, Aaron joined Industrial Color to manage business and software development for its GLOBALedit initiative. GLOBALedit has established the leadership position in digital workflow software, and has led transformative initiatives for clients such as Victoria's Secret, Showtime Networks, NBC Universal and McGraw-Hill Publishing.

Keywords: *enterprise creative workflow, cloud computing, high-speed file transfer, digital photography, XMP metadata, collaboration*

Abstract How has XMP metadata contributed to the integration of image asset workflows for creative and broadcast workflows? The real benefit of XMP is that it removes the need to tightly couple systems, but bridges systems using XMP as a language — to exchange information.

Journal of Digital Asset Management (2008) 4, 331–340. doi:10.1057/dam.2008.44

MM: We're here with Aaron Holm of Industrial Color. Aaron — if you would — give us a little background in terms of your role there at Industrial Color, and maybe a little bit of your history.

AH: I'm the Vice President of Development and Integration at Industrial Color. My job is to oversee GLOBALedit — both from a software development and a business development perspective. I come from a media background. At this point, I view myself as very fortunate to have been involved with several different media platforms as they've gone from analog to digital.

I started out working in audio engineering, specifically for animation. I worked on Saturday morning cartoons, doing audio engineering — building mix theatres, editing suites, doing atmospherics and working with sound libraries. This was in the first days that audio went digital.

From there, I worked in video with Getty — with their first stock footage acquisition — designing digital image ingestion and distribution systems. From there, I moved to in broadcast. That's probably the best training ground out there for learning technologies,

because of the wide variety of technology there in terms of satellite feeds, video systems audio systems, and on-air tickets and CGI.

Then I got more into the business side of things. I opened up the headquarters for a Canadian software company. Then I got into digital asset management. I guess it was in 2004 that I started working with clients in New York on the photo side. I realized that there were no systems in place for digital photography. The entire industry had just embarked on a fairly massive shift in embracing digital photography. That was going to be an increasing trend, and there didn't seem to be systems out there for this new marketplace.

I got involved with Industrial Color to build the GLOBALedit platform. Steve Kalalian, who is the president of Industrial Color, understood very early on — in 2001 or 2002 — *the need for people to collaborate in a web application environment and to manage large volumes of photography — millions of images — in a workflow context.* We both saw this and decided to get very serious about building a platform for that. So we started working on GLOBALedit.

Aaron Holm
VP of Development and
Integration,
Industrial Color,
32 Avenue of the
Americas,
22nd Floor,
New York, NY 10013, USA
Tel: +1 212 334 3253
E-mail: aholm@
industrialcolor.com

MM: In terms of the overall history of Industrial Color, how has GLOBALedit been the natural evolution of the company's overall business development?

AH: Steve started a company called Impact Digital in 1991 — a graphic design retouching company. He built an excellent firm with a very good reputation for their high-end retouching services. It was very much grounded in design and fashion and the graphic arts.

In 2001, Steve saw the market need for a company to provide digital capture services. The retouching company had started to receive digital files. Previously, everything had been provided in an analog format for all their jobs. Then, all of a sudden, these files were coming in digitally. There was no consistency and no color management.

Steve saw the market need for there to be a company that could provide this type of service for large photo producers — for people that shoot a lot of imagery — to go there on set, with all the equipment and all the techs and all the knowledge — to do two things.

First, to help them make the jump to digital so that their photographers could go from a film/photo/digital workflow seamlessly. Second, to be able to provide the type of support, infrastructure and scale needed to support companies like NBC and Victoria's Secret — people that shoot a lot of content.

MM: So Industrial Color was formed almost as a movie production firm? But this would be — in fact — a production logistics firm for digital capture, high-quality color, digital photography. Correct?

AH: Right.

You have these high-demand photo shoots, where they're going on-location or on-studio. They're going to shoot for three or four days and capture thousands and thousands of images. They're on a tight production schedule, and work needs to be turned around very quickly. That's their professional working environment.

We've shot out of the side of a helicopter up at Mount Rushmore. We've done shoots in the jungle in Belize. We've done shoots in the desert in Morocco. It requires a real investment in people and in technology, to be able to go for eight days to Iceland and shoot on the North Atlantic. You have to really have your act together.

Providing excellent service is imperative. You have to invest, to train, and you have to have infrastructure.

That was the direct foray into GLOBALedit. As soon as all these files were coming in, the next natural step is that people want *access* to them. They want to be able to work with them. There's a time requirement.

At that point, there needed to be a system that would allow this. The first step was to look out there at the marketplace and hopefully just buy something. The original intent was not to become a software development company at all. The original intent was to provide a service for clients. It just so happened that there was no software out there.

After a couple of DAM solutions, we realized that we needed to get pretty serious about building something. That began, in a real way, our global operations.

The decision was made to build and own the infrastructure, to build the data center, to bring development in-house, to have software development become part of the culture and really utilize the vertical integration of the different companies.

Our goal was to create a software development company that was very much focused on creative production and viewed it from a creative workflow perspective — understanding how this industry works, and what's needed for people to be able to do their jobs at the highest level.

MM: Perhaps you can take us through the typical set of activities from one of these high-demand photo shoots — specifically, how GLOBALedit supports these various activities and workflows.

AH: The first step is photo assignment. At some point, the decision is made that there will be a photo shoot — either because it has to support a television or film production, or to support the production of a catalog or a news or sporting event.

Then the various people that are needed to handle that shoot are assembled. It's important to consider that step in the workflow. There is the point of origin called the job assignment. When we get into the concept of metadata, that's really where metadata collection begins. At the photo assignment stage, information starts — information about a shoot starts

coming in. Even though there's no content yet, that information has started to roll in.

MM: In some respects, Aaron, you're really talking about creating a metadata container that will subsequently get filled with raw files and then edited photography, so on and so on?

AH: We call this a "Job Token." Essentially, in our world, we're working with XMP metadata. That's where we're really starting to push things — and where a lot of our clients are getting benefits. Just getting that really early, and getting the basics of the job together, when the people have the information, when they've negotiated the rights, and when they've made the schedule — when they're produced the call sheet. That's when they have all this information. That's when the metadata for a creative production process begins.

The next step is the actual shoot. The shoot is a really magical situation where all the different elements come together. You have the talent, the creative director, the art director, photographers, technicians, assistants, makeup, hair and wardrobe. All of these people come together to produce the shoot.

It's a very unique time in photography, where concepts become real, and the photography is captured. It's important to understand that process. From that point on, you don't want to lose any of the value that was created on set.

Factors such as the color and the information and what the intent of the shoot was. How is that going to be carried through the product process in such a way that that value's not lost?

The photo shoot happens. It can be one day or multiple days. If it's news, it can be ongoing. If it's sports, it can be a series shot over a period of hours. There's a large amount of photography that's produced.

People who are shooting digitally tend to shoot quite a few files. A typical day might produce 2,000 frames. Using medium-format camera backs now, where you're capturing at 39 megapixels — you're talking about unprocessed raw files that are 80mg each. They're then processed out to 125 megabytes. Multiply that by 2,000 files and you have a *lot* of data.

So there's a challenge there in the workflow to be able to get that data moving to the people that need to see it next in the production process. As soon as the shoot's over, there's a big demand to then go from 2,000 frames down to

the ones that are needed. The team goes through them and gets rid of all the ones that aren't going to work. They start to play with different concepts, and they want to view them in layouts. They want to review the color and check the focus.

MM: At that point, the creative director works with a photo editor?

AH: Typically, you're going to have several steps within this process. One is what they call a *broad edit* or a *wide edit*. The intent of the broad edit is to get rid of the captures that just don't work — all the frames that seem awkward or are lit incorrectly. The broad edit involves getting rid of those shots and editing down to a manageable chunk of content.

If you go from 2,000 frames down to 30 or 50, that might be your broad edit. Typically, that would be done on set, or very quickly after the shoot, and is done by either the art director or sometimes by the photographer, depending on the type of shoot.

Often, there's a group of photo editors that work for a larger company. They'll start, and one of them will be assigned to shoot. The photo editor will then dig in. They're all going to want to take part in the broad edit.

MM: At this point, Aaron, people are looking at raw files. Are they looking at them directly within a viewer, or a browser?

AH: Typically, where you run into the first challenge of the workflow is the fact that you have an on-location shoot, and you have to capture very large files. The cameras and software have gotten much better at producing preview jpegs. The whole purpose of those preview jpegs is to allow files to be moved around, because the raw files are just too big.

In our workflow — from the shoot, those preview jpegs are processed out and color-accurate, based on the color curve that was set by the photographer. That's very important, because the color that was captured on set is one of the most important components — it needs to be maintained. Those jpegs that are color-accurate are then ingested into GLOBALedit. That begins the collaboration process.

As soon as files are brought into the GLOBALedit environment where multiple people can access them — that then triggers often four or five simultaneous processes. That's really where the big win is. As soon as you get

photography or creative materials into an environment where people can collaborate, you remove the requirement that a process be linear and have stops. Now the photo edit can begin. Everything starts to happen much faster at this point.

MM: So you're saying that with respect to the raw files, you take the color-accurate jpegs? That's what gets put into the system, and that's what people start to collaborate around?

AH: A good example would be Showtime Networks. Let's say they've shot all the marketing materials for *The Tudors* — which is one of their big shows. Say they shot it at a castle in Ireland and it was a six-day shoot.

After the first day of shooting, they'll take all of the preview jpegs from the day's shoot — organized into the different setups. They'll send them back electronically and upload them to GLOBALedit.

From that moment, all the people at Showtime — the publicists, the art directors, the layout artists and often external people like talent management as well as talent — are given access, as appropriate, to do their different jobs.

MM: Then what happens in the workflow?

AH: Then you tend to break out into more industry-specific workflows. For a catalog production, they're going to want to make photo selects, and approve images for the different layouts that they need for their catalog. They'll get them retouched and then go to print.

If you've got a network or a film studio workflow, it gets more complicated. During the photo selection they're going to pick the ones they like. Talent has the rights, contractually, to kill images. That means that a certain image — if it's marked as a kill — cannot be seen or used, ever, in publicity.

So there are some rules that get put into place that allow systems like GLOBALedit to make workflows possible that weren't possible before.

That can cut out millions of dollars of expenditures from the workflow, simply because these things are managed within a system. If an image is killed, it's not available to other people that need to make selection. That's a *talent-approval workflow*.

Then you have advertising workflows similar to catalog. They're going to need to make some

photo selections and to start to drop photos into layouts and see how they look before they go into production. Increasingly, there's also a large web-publishing component. Print used to drive the production priorities. The people that were producing content for the website would often be left with whatever was there — whatever the print people had selected.

But now, people are very much viewing the print and the web audiences separately — so the web content publishers get to make their own decisions about what content they want to use. That's another scenario where having a collaborative system is very important, to be able to provide many people access to the photography very early on in the process.

MM: Would it be fair to say that there's a broader array of content that would be appropriate in a web context that simply wouldn't work in a print context for a variety of technical reasons? Is that fair to conclude?

AH: Yes. The website people are often looking for different things. They're often producing content that may be analogous to what's going to be in print, but not exactly the same. Their audiences are different. They measure their audiences differently. Their end product is different.

It's different, but it's still going to come from the same photo shoot — that's the interesting part. They're not going to create different shoots. However, you've got web-only organizations now that are doing 20–30 photo shoots a month. These are groups that are only producing web content.

MM: For the most part, that web content is not suitable for subsequent print output? Or is it just not part of their business?

AH: As an organization, these media companies may have web properties and they may have print properties. They don't want to be in a situation where they've limited their options down the road. So where the web property might only need a certain resolution or a certain format of the file for its output, they may decide later that they may want to go back and take one of those image assets and use it for a billboard.

Systems need to be smart about how they archive and manage content. They should not introduce limitations into the workflow. That's really where all the benefit is — in

understanding. The greater the understanding of the end goals that come out of these photo shoots, the better the system's design.

MM: You were talking about how photo selection tends to reflect a very specific functional workflow.

I would imagine that retail — merchandising and point of purchase — would be another typical workflow channel. Is that correct?

AH: Absolutely. Also, the retail operations are in line with an overall trend that includes a lot of visual content in either promotion or marketing. Most major retailers at this point — in their flagship stores, if not the majority of their stores — have very large visuals being swapped out every week or every two weeks, to drive promotions.

If you go down the main streets in Chicago or New York, you have multi-story visuals that are being swapped out every couple of weeks. It's not happening in one store — it's happening in many stores. This production process, from capture to layout to retouching to fulfillment to analytics, and then back into, "What should we shoot?" That information cycle is really important.

They're getting feedback from all of these systems, and then in two weeks they're going to put another set of visuals up in those stores that have to be retouched, and go through the same production process. The automation and the speed and the usability and the workflow components — all these things have to come together in order for these people to be able to create their competitive advantage.

MM: As you were taking us through that process, we printed to the catalog and printed to the web. It got pushed out to advertising or publicity or retail operations. Then you were talking about capturing some analytic data and doing things with that?

AH: Factors such as the color and the information and the intent of the shoot. How is that going to be carried through the product process in such a way that that value's not lost?

I can find out if I can use this again. Has this been used elsewhere? Am I being repetitive with this? When are my model rights up for this image? Can I use it in print? Can I use it in Australia? What are the conditions?

That's the layer of information where now you start to mine some really valuable stuff at the business level.

MM: A lot of companies have begun to develop rights information management systems, almost as stand-alone applications and services that plug into other workflows and processes. It sounds like you're then subsuming some of that rights information management capability into your platform. Is that correct? If so, can you expand on that?

AH: When someone comes to our system, we have to know whether or not they should be able to pull the asset. We have to know everything about what they're going to do, and whether they should have the rights to do it.

When they log in, we have to be able to know what they should be able to see, which versions they should be able to download. Should they be watermarked? Should they be able to get the high-res assets? Should they have to request access for the high-res asset?

As soon as you start to ask all those questions, you have to define rules. What are the rules that would govern whether or not this person is allowed to download that asset?

If you already have defined the rules, you need to inform the rules. The only way to inform the rules is to have enough of the business information there. So I have to know whether the model contract is expired. I have to know whether it's been used in a condition that precludes me from using it now. I have to be able to inform the system to the point where it can make decisions about whether or not a user can have an asset.

Driving those rules from the metadata perspective gives us a common platform to — at the very least — make sure that assets as they come on and off of the system can be normalized, and that I can make sure that no information is lost.

So we're operating from the perspective that we have clients with a lot of content in the system. We need to be able to control when and where they can pull the assets. We're finding that it's more important to do that at the moment that the asset is being requested than it is to abstract that.

MM: You mentioned that you were also then accessing business information out of other systems — be they accounting systems or

publishing systems, bringing that business information into your system as additional metadata.

Did I understand that right?

AH: Yes, to a degree. It's difficult, simply because metadata is early in its evolution. I think it's going to become a real foundation for how systems work, in general. We have situations where we need to pull XMP data from a legal system, and the legal system in its legacy format doesn't really have much awareness of XMP. But because it's an XML construct, we can do it.

But the real benefit envisioned for XMP that's genius is that it removes the need to tightly couple systems. It just means that if you have a system and if I have a system, and if we can agree upon a language, those systems can communicate. There's emerging pressure to adapt to the idea of an open communications framework. I think that most have settled on XMP.

So the short answer to your question is that it's still very challenging. We would love to get more data. We'd love to get more information that would inform the system, and give more information to our users. A lot of it is just hard to get.

I think that's going to be a huge trend, where we're going to get more and more information. That asset that I'm looking at — I'm not going to just be looking at the asset itself — I'm going to be looking at the history of the asset. I'm going to be looking at all of the information known about that asset. I'm going to have data feeds that relate to that asset. I'm going to have a lot of decision-making criteria at the moment that I need to deal with any individual digital asset.

MM: For the readers of this interview or listeners online, could you give us a good summary of XMP and how it works?

AH: Sure.

I think XMP is essentially a technology platform, a mechanism that allows you to define how you want to describe information in your world in such a way that it can be translated to how someone else describes information.

MM: First, we're talking about a technical standard — a publishing standard — originating with Adobe.

We're talking about embedding a certain amount of data into the file header. That is to

say, it's inside the file. But exposing it as XML allows it to be read externally, without necessarily having to open a file.

AH: XMP is when you put your translation earphones on at the United Nations. Someone's been talking for 20 min, and now you understand what they're saying.

The information has been there since the beginning, but the listener had no mechanism to translate the content for them.

The thing is, XMP on its own — you can have a million images on your server that all have XMP, and they don't mean anything to you. The other systems that you're using to address those assets don't speak XMP. So, XMP really has to be thought of as a language — because unless the system that you're working in is conversant in XMP, you will get no value from that XMP.

XMP only has value when systems are in place to handle assets that are XMP-enabled.

That's the only point at which the XMP has value. The good news is that XMP — at the industry level — has been commonly agreed upon as the metadata framework that has been adopted.

That means the guesswork is over. We don't have to worry. Is XMP the way to go? Am I going to waste my investment? Is this a bad decision that's been agreed upon? People work with it. IPTC has adopted it. All of those core questions have gone away.

Now we have a mature technology that's well known and can be implemented. It has enough of an industry momentum that it's very exciting.

What it allows you to do is to create very clean integrations, and move media around in such a way that over the lifecycle of an asset, you aggregate value points for everyone that's touched that asset.

At the end of it — when you put it into an archive, or when you retrieve it — you theoretically should know everything that's happened to it along the way.

MM: Maybe this is a good time to bring up the emergence of what you and I in the past have called digital supply chains. Largely comprised of what we refer to as “loosely coupled systems.” Could you just take us through that and, specifically, how you see that evolving in the market as an industry?

AH: I think that what seems apparent — at least from a philosophical standpoint — is that if you look at a workflow and you look at a digital workflow, throughout the production process, there are multiple steps. Each one of those steps represents a point on the digital supply chain. Now theoretically, in any system or workflow, if you map out your workflow at this moment in time, the simplest exercise you can do is look at everywhere there's a stop in the flow. You can pretty much assume that over time, those stops can go designed away.

Over time, an exercise can be to take everywhere that media stops, and remove that necessity. Instead, figure out where there should be choice. So you replace necessity with choice, and you say, "Well, if in the digital supply chain the assets never had to stop, where should they stop?"

Then you're designing a digital supply chain that's based on the inherent value of an interconnected and loosely coupled system. That means that there are going to be points in time at which decisions are made. Those may take a long time or a short time — but that's not really a systems constraint.

In a loosely coupled digital supply chain, wherever the files need to move, they're immediately integrated and available to the next system that needs them. That next system knows exactly what to do with them the moment it gets them. So you have a couple of things that you need to work on, in order to really make that a reality. That's what we're 100 per cent focused on.

Also very important, you need to be able to transfer files at high speeds.

MM: And by a "loosely coupled system" you mean multiple systems that may exist or live at different business entities, and certainly have different governance policies by which the systems are managed?

AH: To a degree. They could also live at the same business venture. You could have several systems in the same company that need to talk to each other to represent the best in breed for that system. In a digital supply chain, you might have a publishing system, a DAM system, a photo studio. These are all different systems that may be in the same company.

But the concept of decoupling the systems — from an integration standpoint — is basically

trying to reduce the interdependency of software versioning and scale.

MM: So if one system goes down, the entire supply chain doesn't necessarily go dark.

AH: Correct.

Also, versioning is difficult. You may have one system that's talking to the other system, and they're talking directly to each other. With one upgrade, that could fail.

MM: I did some work with IBM around their self-managing autonomic computing group. I was astonished with the research of IT systems delivery or IT systems management. The largest, most significant cause of system failure entailed provisioning and upgrade.

AH: Sure. Absolutely. That's the biggest change-management exercise.

MM: Yes, the idea was that versioning upgrades and so on is the principal cause for system failure today.

AH: We did an "integration" for the Showtime implementation with a DAM vendor. The president of the DAM vendor's company is a friend.

Our testing involved custom-creating schema on our software, putting some assets in and pulling them down. Sending them over to the client to test them. Then the client said, "Yes, it works." So they sent me their schema and we implemented the schema. We scanned some files. We sent them over. They populated the other system. I never spoke with the other vendor. The other vendor never spoke with us, and we have an integrated system. It took 10 minutes.

That's the goal.

MM: This gets to another trend that I'd like you to speak to. Sometimes it's referred to as a "mash-up." You have multiple systems that share some sort of metadata schema — either file-based like XMP, or another more traditional syndication framework. It links previously disparate systems together, in what might be minutes or hours as opposed to days or weeks.

AH: It's interesting, at the DAM level, just because of the heavy lifting. A data feed from your Facebook account with a news feed from Google — that's a pretty low-risk, low-overhead thing to do. You're going to take a couple of files and you're going to make a UI and present them.

Now if you want to move 100,000 files into a system that has one specific value, and then

make them available to another system that has a different type of value, it's a bit more complex. But from an architecture perspective, that's definitely where it's going.

We already do that internally with a couple of the different components that we utilize to provide the comprehensive solutions. So underneath the presentation of GLOBALedit, you have system-level services that are taking care of different components.

You have file transfer, conversion services, routing services, metadata. These are all service-level applications that provide functionality to the GLOBALedit suite. But on their own, they're also stand-alone applications.

One of the things that we're starting to do now is to make these available in different forms. We have clients that have specific projects where the most important element of what they need to do is to move large files around at high speeds. We can decouple the file transfer component and make that available to them. We can make the UI elements of the file transfer component available, but decouple them from the rest of the GLOBALedit workflow — because for that project it's not required.

It's the presentation of the application in different contexts that exposes the low-level system services.

MM: Aaron, you've described the service-delivery mode as both software-as-a-service, where it exists in your data center and — as you say — existing in the cloud. Alternatively, you have software that you license and subsequently deploy on-premises at a customer's organization.

Can you take us through that as an overall IT service management framework?

AH: Because we own and operate our data center, early on we started to view the application development process as a process of building software that runs on a data center. So we're not really building individual components that run on a desktop machine or that run on a server. GLOBALedit can't run on one server. One server is just not going to do it. It has many different moving parts, and it requires storage and processing and web services, as well as metadata services and various different engines.

Some of our clients want the deployed solution — GLOBALedit is available both as a

hosted product where clients utilize our infrastructure — but we also have clients that have adequate infrastructure and IT support, and people at their location where they want to utilize it. So there are enough clients out there at this point who have the need and have made the investment and have the infrastructure and the management in place to run their own data center, where GLOBALedit is designed specifically to make use of data center environments.

So, to bring together and enable software scalability by default, you're putting in place a software platform that's designed to scale, based on data center resources.

MM: Is there a particular type of business requirement or threshold at which point it makes sense to go from a rented provision-on-demand service to having the service installed on-premises?

AH: I don't know if I found any direct correlation regarding that point. I found that it has a lot to do with corporate culture and with whether or not the IT and creative departments are on the same page. It has a lot to do with corporate directives, and with content as well.

You have companies that see issues like piracy or privacy. Or their assets just have too much risk getting outside of their controlled environment. That's a simple one.

If it's too risky for the assets to ever transport outside of your environment, you build an environment where you can handle your workflow. Other than the value of the asset and the perceived value of the asset within the organization and whether it's seen as a core risk, if it travels outside, it really depends on the type of media that's being produced.

MM: As we begin to wrap up our conversation today, can you provide us with some more forward-looking insights in terms of how you see DAM evolve or specialize in increasingly exotic niches?

AH: Yes. I think that the user interfaces are just going to get so much better. There's a real awareness on the client side, that it's a necessity. If the creative user can't use the system fully, it is often a user-interface concern. In DAM, traditionally, user-interfaces have been awful. We're really trying to stress that they don't have to be.

MM: What about them have made them awful?

AH: You have a couple of key elements.

This is a cultural issue. It has to be engrained in the culture of the software development company. There has to be a real understanding at the DNA level of how the end-user wants to use the technology that's being put in front of them. There has to be a corporate, company-wide directive that this is crucial.

It has to do with the technology landscape and the evolution of technology away from a client-server model into more of an SOA model and more of a web-interface model. What you can do with a web interface is pretty much unlimited. You can expose a function in any way you want.

And if you can expose a function in any way you want, you're not limited by native desktop component development. You don't have to be limited by the buttons and widgets that are available to you when you're developing desktop software. So, when you're developing user interfaces that run in a web-based environment that give you unlimited flexibility — that can be a disaster, if you don't know what to do with it.

The user-interface side is about to undergo massive change, because it's now technologically possible. The other thing I see as crucial is file transfer developments. One of the things that we've spent a lot of time this last year developing is a partnership with our high-speed file-transport partner, Aspera. They're an excellent company with excellent technology.

We have clients that are moving hundreds and hundreds of gigabytes of content each month throughout their entire vendor chain. They're moving them to retouchers and printers and layout artists. These are multi-gigabyte files. Now everything is online and files can be distributed. They can be delivered electronically at high speed. It just changes everything. If you can deliver the files electronically, then you've truly got a digital supply chain. That's a huge change.

MM: What would be the secondary or unintended consequence that would likely emerge from that capability you just described?

AH: First is the awareness that a very wasteful and ecologically harmful process is not necessary. If you establish a workflow that transforms a shipping workflow and makes it all electronic —

if you show that that's possible, then the question becomes, "Why aren't you doing it that way?"

The other thing that comes out of this capacity is that people very quickly recognize the value of bandwidth. I think that over the next few years we could very likely see bandwidth replace the processor as a metric for innovation in software.

MM: Doesn't this in fact accelerate the overall cycle time throughout the entire value chain? And doesn't this hyper-acceleration of core workflows and business process produce interesting consequences at the point of consumption — be it a retail store operation, a website, a catalog or a CDR?

AH: You're right. There are two things happening.

There are unintended consequences when you start to speed things up. But I think it's important to make sure that you retain the elements of choice. At any point, you can analyze things and slow them down, if it's important. At the end of the day, that'll change — I think — how people do interact with consumers.

If you can give them more or you can give them less, faster or slower, ultimately businesses are going to ask what their consumers want. That again is driving choice.

MM: I recently had a conversation with your client over at Victoria's Secret — as well as a number of other companies that are in the retail marketing business — and I think there are several consequences of this acceleration and hyper-acceleration of innovation and supply chains.

One — the overall market space clutter has become overwhelming. Point of purchase has emerged as even more strategic and influential, in terms of impulse purchasing.

Retailers are beginning to take the traditions of direct mail — of targeted demographic and panels — and the whole notion that there are different kinds of buyers to whom they want to send different things. They're then taking that analytic mindset into the store, and executing what's now referred to as "shopper" marketing strategies.

In the retail operation, you're targeting very specific shopper. This then call for very specific, unique engagement strategies that often require

very specific visuals and point-of-purchase displays and merchandising techniques.

AH: Sure. Then at that point, the question is, if that's seen as being important, and it's seen as being a business model that needs to be developed and supported, can you do it — and can you do it quickly?

MM: I guess the point that I'm making is that as we get into the digital supply chain, some of the unintended consequences are a hyper-acceleration of innovation and marketing strategies, which only puts further emphasis on POP and in-store merchandising.

AH: Yes. I think that the trends we were just talking about are going to be huge, there, where you look at what the UI is at the point of purchase. Right? What is the speed of file

transfer to that point of purchase? What is the bi-directional metadata that's coming to and from that point of purchase? Those are the underlying elements, and they're consistent all the way through.

MM: Thank you, Aaron, for our conversation today. Hopefully we'll be able to pick up in six months or so with some additional client success stories and new developments.

AH: I hope so, too.

MM: Are there any concluding remarks that you'd like to share with us?

AH: I think this pace of change that we've seen accelerating is only going to continue. The organizations that master change will be phenomenally successful.

MM: Thank you.