Interview

Driving workflow innovations into large media operations – An interview with Vincent A. Di Paola of Moksa Incorporated

Vincent A. Di Paola

was one of the first graphic designers in the United States to recognize the power of Apple Macintosh technology for publishing. He is credited with being the first to produce a full-color catalog with digital technology from start to finish. Over the years, he has developed a strong but independent working relationship with Apple Computer. They often seek his advice, use him as a consultant to others, and send him software and hardware for evaluation. Today he is a much sought after speaker in the fields of digital publishing, prepress and image manipulation. He tours extensively, speaking at seminars and contributing to workshops, and has taught photography and graphic design at the university level.

ABSTRACT Michael Moon interviews Vince Di Paola, Principal at Moksa Incorporated, examining how the open plug-in architecture of Canto Cumulus transforms a DAM into a process innovation platform.

Journal of Digital Asset Management (2009) 5, 246–258. doi:10.1057/dam.2009.11

Keywords: Cumulus; DAM workflows; enterprise DAM; server virtualization; plug-in architecture; digital publishing

PERSONAL BACKGROUND

MM: We're here with Vincent DiPaola of Moksa. Vince, let's start off with a summary of your professional career, and then about Moksa, your firm.

VAD: Sure. Thank you, Michael.

I became involved with the technology we're now integrating and supporting, through the graphic arts industry. I have been involved in *graphic arts for over 30 years*, and still consider myself to be a graphic specialist.

When the technology helped my industry – the graphic industry – execute what we do more efficiently and gave us much *more control over the overall design process*, I moved immediately to it. I was fortunate enough to have the work I was doing recognized by companies like Adobe and Apple – who asked me to do a number of seminars and tours for them.

I wound up leaving the creative side of graphic arts, and more or less became a consultant to the industry on how to *implement technology into our daily workflows* and improve the process.

MOKSA INCORPORATED

MM: You started Moksa when?

VAD: In 1997, for the first time. Apple had tagged me while I was *VP of Technology* for large prepress operation called *Gamma One*. At that time, Gamma One was being acquired by what is now called Vertis. Apple asked me to become their international spokesperson. I used the opportunity to launch Moksa, Inc.

I traveled all through Europe – 18 countries, and all through the Far East and South America. I spent months on the road for them, all outside the US, as their spokesperson, for what at that time was called the *'Masters of Media' program* – which you and I actually launched together in the US if you recall.

'Masters of Media' means *the ability for graphic designers to master the communication capability* available through the Macintosh, creating documents for print or film – throughout all communication media.

Correspondence:

Vincent A. DiPaola
Moksa, Inc., 539 Main
Street, Studio 2, Branford,
CT 06405, USA
E-mail: dipaola@moksa.
com

That project was for Apple and Adobe. Canto Cumulus was one of the products I used in my presentations, as well.

PREPRESS FOR CATALOGS

MM: Over the years, what sort of client work have you performed as it relates to DAM, and DAM-enabled systems and processes?

VAD: The industry was creating an enormous number of digital graphic files. The greatest challenge was getting these digital files organized. You can well imagine, as the VP of Technology for Gamma One – a company that specialized in working with catalog companies like Victoria's Secret, J Crew and Bloomingdales – what it was like dealing with catalogs full of highly color-corrected digital images. *We would help the client get all of their digital assets under control* for catalog production. This prevented a lot of miscommunication – and a lot of image misprints.

You can imagine the level of color requirements involved in producing a catalog like a Victoria's Secret or a J Crew. If there was *any difference between the color of the product in the catalog photo*, when compared to what the product actually looked like, you'd have an *unacceptable number of expensive, money-losing product returns*.

So, we had to *apply a business understanding that would prevent product returns – resulting in our clients losing money* – this connected to all of the work we were doing on pre-press side to the clients production process. We had to ensure that color-correct images went into the client's printed catalog. I started building these DAM systems to address this issue, so that I could completely control every digital asset that was being used.

BROAD CLIENTELE, VARIED NEEDS

MM: Some of the other client work that's tangential to DAM included what?

VAD: I had installed many different DAM systems while I was with Gamma One. When I started Moksa in 1997, the principles behind their application were second nature to me. I expanded the practice to the point where I was able to apply these same principles to many different industries. Moksa expanded to take on clients like Madison Square Garden, Scientific

American, Military and Government installations; we are presently working with the United Nations and the Metropolitan Water Authority of southern California.

In each and every instance, the lowest-level requirement is that all of the assets come under your control. DAM was the perfect technology for addressing this requirement.

In the beginning, we struggled with many different DAM product offerings. Each vendor's product had a great many shortcomings. I know that Moksa, in particular, had to use a lot of ingenuity in order to make a product meet the business requirements. It never works if you try to wrap the client around the product. *The product must always be wrapped around the client*.

Moksa then expanded into what we do today – that is, go into very large organizations and wrap the product around the business needs by expanding the core functionality of the base products we deploy. Most of the clients we're working with are larger organizations, with diverse requirements.

The first challenge is, 'How do I get all of my assets under control?' Then, 'How do I work with assets in a manner that *allows my business to perform the way it needs to perform*, and to accomplish a high level of control over how those assets are being used?'

NEWSDAY CALLED

MM: From the point of view of a DAM practice when you first engaged Newsday, what did you find – in terms of an organizational profile and its systems, processes and accountabilities?

VAD: Well, at Newsday we had a unique situation. I was brought into the company by James Kober. He had been working with another Canto partner.

That partner understood Canto Cumulus as an application, but lacked an understanding of what Newsday was trying to do with it. When James brought me in, I started where I would always start. That is to dial into what the real business applications are for the client. It's always a process that needs to be handled delicately. Unless the system meets business requirements – and this is the whole point of the discovery process – the project is doomed from the outset.

A client often doesn't really know what the technology can do for them. It falls to the consultant to ask the kinds of questions that will extract from the client exactly what they're really trying to accomplish. Not simply the short-term goals, but to understand where their business is headed. So, long term, you can configure the system properly. In short, only if you know the business problems, can you deploy a true business solution.

PARADOXES IN PIONEERING DAM

MM: That gets to an underlying paradox, Vince, of many early DAM pioneers, and the evolution of their work: clients almost never know their business in a structural sense. That takes looking-from-the-outside-in. The client is immersed in their day-to-day operations and the meetings – the personnel and so on. In many respects, to really understand their business needs, clients have to get outside of their day-to-day and structure what's going on. In many cases, clients well intended and full of information – a vast inventory of experiences – really can't tell you what their needs are, because they're too close to it.

VAD: Exactly.

MM: Then whatever they do give you – which they believe they've given you with candor and honesty – is completely off base. However, as soon as you start introducing DAM-enabled capabilities, they learn to think through a workflow or process in a more systemic way. So the mindset and their understanding of their organization changes as they learn to get deep into the DAM mindset.

VAD: Exactly.

MM: This also makes it really difficult and sets up the paradox: If you go in with a fairly complete end-to-end solution, not even a quarter way into their adoption of the solution, they want to start making wholesale architectural changes. That leads into a long, hard, expensive journey that quite frankly many DAMsters don't complete. They give up. They run out of funding.

VAD: That's exactly right! We find a lot of white elephants in that graveyard that you're talking about here. The project gets completely abandoned.

Newsday, with James Kober, is a perfect example of this. James inherited the Cumulus

system. In the beginning, his predecessor had one concept of what was to be done. But when James arrived he realized the potentiality of the system, he asked Moksa to create functionality extensions to Cumulus that would really alter the way that Newsday was doing business. *James was fortunate that he did not start with a top-down framework*, with millions of dollars sunk into professional services. We were able to make significant changes right away.

THREE WORKFLOW AUTOMATION PHILOSOPHIES

MM: Great. Can you explain more about workflow and process automation and, specifically the three basic approaches of automating activities, workflows and processes?

Let's start with the top-down strategic process management platform. You see this commonly in the enterprise content-management space – with business process management capabilities – all within a robust DAM framework. It takes considerable effort and skill to wire the platform to how a company does business. It takes often many, many man years and man decades of effort.

VAD: Right.

MM: Then there's the tactical, bottom-up automation through scripts and scripting – and actions that are part of a suite, for example, Adobe Creative Suite. The problem with many of the scripting technologies – while very powerful – and very easy to make incremental innovations to – is that they tend to remain specific to the platform and the user. There's usually a considerable amount of maintenance associated with keeping all of these scripts up and running.

Then there's a third – the middle road: 'We'll innovate as a function of having a technical framework supported by an open plug-in architecture – plug-ins for the client or for the server, extending or changing the functionality of the core software product.'

Could you speak to your personal experience of dealing with those three modes of process and workflow automation, speaking to their upsides and downsides?

VAD: I'd love to. *I'll start with the 'bottom-up' approach.* That's the one with which most companies started. These systems are especially prevalent in prepress houses and graphic

agencies. You'll find hundreds of *home-built, scriptable little systems, playing task-oriented roles in the day-to-day requirements of a department.*

In a fairly short period of time – as requirements expand – controlling these systems is like trying to herd chickens. This is the problem; as I said earlier, these systems are really task-oriented, and *when it comes to getting these systems to exchange information, the challenges become overwhelming.* While they may have been a time-saver in the beginning, getting yourself out of this situation is harder than anyone imagined.

I find the users to be the biggest obstacle to breaking free from these systems, and the biggest stumbling block to a replacement solution's success. Even if a new system is properly configured and installed, people tend to dig their heels in when they are presented with change. They want to return to old familiar ways. Companies must focus on the long-term goals, build for the present, but plan for the future, so you don't paint yourself into a corner.

Now, let's discuss the 'top-down' approach that you mentioned, it's one that I would really caution against. While they believed they were purchasing a product, the client finds that they've *inadvertently created a 'one-off' version of a very expensive system.* It often becomes an isolated – what I called earlier a 'white elephant,' – and ends up being a 'silo,' one that no one can operate and no open system can communicate with.

With these 'top-down' systems you often have a central core of technologies, but with each and every update installation, the implementation of that core technology moves further and further away from its origins – so much so that *upgrades become impossible.* Due to extensive customizations, any changes made to the core system will break the product, and the software won't run. The client ends up not being able to move from one version to the next, rendering that entire implementation useless.

OPEN PLUG-IN ARCHITECTURE

MM: As I understand the Canto suite – both server and client – its plug-in architecture for end-use client and servers really impressed me – Embedded Java Plug-ins (EJPs) for the end-use client or a Java-based Embedded Server Plug-in (ESP) – all inherently cross-platform.

And because of the unique permissioning of Cumulus, I can create a new plug-in, either on the client's side or on the server side – such as a new item in a pull-down menu – and only provision that change to me, the developer. Or to one or two power-users that would be able to deal with whatever kind of hiccup might be discovered in the Q/A process.

VAD: Exactly.

MM: So I can create incremental innovations – vis-à-vis very discrete enhancements through plug-ins – either on the client side or the server side and provision those changes only to very specific users without having to take down the system. Did I understand that right?

VAD: You stated that perfectly.

MM: Could you expand on that notion in terms of a process-innovation platform?

VAD: Of the three, this is the one that I've really centered on, because this is *the only one that I believe works, the one that offers the greatest chance of success,* and has the least chance of becoming obsolete.

Here, unlike the 'top-down' architecture that we discussed earlier, you really have a core product. With every version release, and with every expansion we make to the core functionality, we ensure its *applicability to every installation that we've put in.* Every one of our clients gets to take advantage of these new features.

With the ESPs – the Embedded Server Plug-ins – and the EJPs – those client side plug-ins – working through the APIs, we ensure that the system keeps moving forward. They're really *enhancements to the over-all system.* We take advantage of the core functionality through these plug-ins – both on the server and client sides – utilizing the full system.

The ESP/EJP architecture allows the development team at Canto – based in Berlin – to focus on future development of the product, expanding it with each new version. They're now focusing on version 8.0 of Cumulus. All of the development that we did at Newsday or Scientific American will be able to take advantage of each new version that they release.

That simply enhances the validity of the extensive programming that we've done for our clients, making it easier for companies like Newsday to meet some of their challenges going forward.

DEVELOPER ECOSYSTEM

MM: This calls attention to the developer ecosystem. Around the Cumulus platform, hundreds of independent developers/consultants/integrators have created lots of nifty little plug-ins and utilities – if not applications – using the software developer kit in the C++ libraries, published APIs and open plug-in architecture. Canto also enabled a marketplace for these third-party plug-ins to thrive.

This marketplace extends the functionality of the basic platform in a way that a small software company could never hope to achieve – because of the inherent diversity of the applications and the special niche of focuses that it takes to create those specialized applications.

It creates a marketplace for end-user customers to source additional capability much, much more cheaply and almost immediately. So the time-to-value is almost immediate as a function of using a plug-in – as opposed to paying somebody their US\$100 an hour to write these specialized things.

It gives you, the developer, a place to produce incremental revenue and make money while you sleep, as a function of being able to take these tools that you develop to get a particular job done – but then put it into the marketplace and make some additional money.

VAD: Exactly.

MM: Would you explain how you've taken advantage of this ecosystem? And specifically, how Newsday benefited from sourcing a solution from this ecosystem?

VAD: Yes! Exactly, kudos to what you just said about the functionality. That is the reason that this particular space is working so well for everyone, for all the Canto Partners – and it's been so beneficial to the industry. It really has made a big change in the way we all interact. *There isn't any reason to duplicate efforts; we leverage each other's products*, affording us the opportunity to build additional functionality that meets business requirements. As an example, Moksa works very closely with Modula4 – in Berlin, and XLent Software – in Belgium, to offer a wide array of solutions.

First, regarding the ecosystem. Well, when you create these *top-down 'one-off' systems* we spoke about earlier, you can see the obvious problem: *'Who is going to support it?'* Especially

when internal staff moves on. Support is usually the biggest obstacle to the longevity of these systems.

Whereas, if you have the *type of ecosystem the Canto partners have formed*, each developing a wide array of extended solution, each understanding how to maintain the Cumulus core, there is less likely a chance that your system will become isolated. Bottom line, the clients are able to easily get the help and expertise they require, and become less likely to create that 'white elephant'.

Moksa has been able to leverage and benefit from what has been developed by other Canto Partners. I have a unique background in graphic design and graphic arts, spanning over 30 years. I can take the knowledge I've gained and leverage these plug-ins, services and client side development to really enhance my client's workflow.

Cumulus is widely used in education, museums and art galleries. People who understand those industries, and what those industries require, can leverage their knowledge by expanding Cumulus to meet the requirements.

Because we can *take an application built for one business and reuse it in another*, the developer can leverage their time and selling potential. For example, Moksa, as a Canto Partner, can take what was built in Newsday, and implement it in many other places.

Because we don't need to charge the same as we would for a 'one-time' solution, when we build new functionality for the originating client – my client benefits. Because I know I can utilize it in other installations – all my clients benefit. And, when I offer it to other partners, the benefits are wide reaching.

Let's also take the case of Newsday, where we've created a PDF advertising management system. For Newsday, as I improve that system and sell it to other people, I am extending the functionality. Because they initially financed Moksa's building of the PDF advertising workflow, *I can offer them the upgrades for free.*

RAPID DEVELOPMENT AT NEWSDAY

MM: So would you take us through some of the developments and successes at Newsday, specifically calling attention to how you used

Cumulus as a framework for rapid development of purpose-built applications using DAM?

VAD: Okay, let me see if I can provide you with an example.

Initially, one of the biggest problems in any DAM system is getting the proper metadata into the proper fields. Out-of-the-box, *Canto* offers several ways of managing the ingestion of metadata. But there must be a method to ensure that the people who are at their workstations ingesting the assets contribute as much as they possibly can to the process.

To that end, Moksa has put together a program we call the *Metadata Manager*. When the end-user puts assets into the system, the Metadata Manager *actually brings up a form that requires them to populate mandatory fields* – fields that are required by the department manager, in this case, James Kober. We have expanded the functionality to include the ability for an end-user to assign an XML document or Excel spreadsheet to assets during the ingestion process. If the end-user decides to skip a file, the Metadata Manager will omit it from the ingestion process and place it in a designated ‘holding’ directory. All this happens in single or batch ingestion processes.

MM: So that would be similar, Vince, to a job ticket? Except this would be a FedEx-like packing slip used for cataloging a newly created asset?

VAD: That’s exactly right. So when the asset comes into the system, it has immediate usage. You can search and retrieve it, and you know what you’re looking at once you’ve found it. And most importantly, you know what it relates to.

The end-user gets a chance to actually become one of the contributing members in how these assets get into the DAM properly. Without a structure like this, assets would just be loosely gathered in that database. You’d never really get them all together at one time. Your database would fracture or become compromised.

MM: You did that as a function of an EjaP embedded Java plug-in. Right?

VAD: That’s correct.

ASSET INGESTION, MASTERED

MM: Another neat aspect already in the Cumulus product is the ability to create a whole

bunch of category structures. As you’d drag an asset into a category, you can drag it into several – three, four or five different categories at once. Each time you drag it into a category, it gets assigned to that category and can be found based on that association.

VAD: Yes, that’s correct. So, let’s talk about a new product Moksa is just launching, called TypeTrax. It takes that to a whole *new level of automated ingestion and tagging*.

The TypeTrax product really helps you manage an entire creative project. What *Canto* offers out-of-the-box is great, but it puts the responsibility on the end-user to manually drop assets into the logical structure that Cumulus allows the end-users to construct.

TypeTrax leverages *the logical category structure in Canto* in the following way. With TypeTrax, when you start a project – graphic designers work pretty much the same – we create a Cumulus category, and that becomes a digital job bag. Inside the digital job bag are your InDesign documents and images and fonts, etc. There’s line art and logos as well. Creative briefs, documentation such as model releases, etc. all of that gets put into a single category. It’s typical, and I think unjust, to require your graphics people to try to find the categories that each of those elements should go into.

TypeTrax allows you to *import the entire project through a single-menu option*. TypeTrax then goes through and analyzes the contents of the hard drive folder. It analyzes every component that’s in the folder. It acknowledges the images. It actually parses the InDesign document, sees all the relationships and then goes back to the folder, to try to understand that these are referenced assets. It captures the fonts that have traditionally been poorly handled – or not handled at all, as a foreign file type – by DAM systems.

FONTS MISSING NO MORE

MM: I want to explain your solution to a pernicious and largely unsolved problem – and as someone who has a publishing operation, let me summarize the very annoying problem. My graphic artist and editor, Iris, will create an Illustrator diagram for me, containing our special embedded fonts. When she sends it to me, I open it up and get ‘font missing,’ even though I have the font on my machine. For some reason,

they don't automatically sync. Really, really annoying if I'm traveling and don't have the particular font that only lives on her machine. So I have to go into the Illustrator document and reinstitute the type that was already there. In other cases, I have to get back to Iris and have her use another font. Damn annoying.

VAD: Thank you! You stated that perfectly.

MM: So let me go on further. Then I put this Adobe Illustrator file into an InDesign document for a white paper, and in instances where we send it out to a printer, often times the printer comes back and says, 'Well, we don't have this particular font. It didn't show up in the PDF.' Or, 'It's an illegal use of the font.' Or something like that! So we have to go through and re-distill the PDF and scrub it of that illegal font.

VAD: That is precisely the issue *TypeTrax* addresses. Working together with *FontAgent Pro*, *TypeTrax* actually parses the documents, locating the fonts that are inside the document. Let's take an InDesign document, for instance. It tracks what the font is doing through something called a 'font descriptor.' The font descriptor 'describes' what a designer did with the font, that is, did he/she make kerning adjustments?

Of course, at the time the page is committed to print (or digitally to PDF), if the proper font is not installed, it causes all kinds of problems, from rendering to the reflowing of items on the page. Instead, *TypeTrax*, working together with *FontAgent Pro*, ensures that the font being supplied matches the 'font descriptor'.

FONTS AS A DIGITAL ASSET

MM: So it treats each font as a digital asset – creating a profile of the font with metadata and linkage to its instances in files and other users.

VAD: That's correct. *TypeTrax* takes the 'font descriptor' and builds a DAM record around it inside *Cumulus*. Working together with *FontAgent Pro*, it pulls all of the font information out, and puts it into a special view called a 'font view.' We place all that information that you just described, Michael, in that font view descriptor.

The end-user doesn't really need to, but could, use this information. *TypeTrax* brings that 'font descriptor' into *Cumulus* with the projects. This allows the user to match the font that he wants to pair with that font descriptor. Or we'll do it dynamically.

TypeTrax will, working together with *FontAgent Pro*, search for all possible matches to the 'font descriptor' it can find. It then rates the results and based on this rating, *TypeTrax* can confirm 100 per cent for sure that this is the font being used on the page. But, if the designer prefers to use a font in the search result that's not the one *TypeTrax* chose, they have the capability of doing so.

MM: So if I were to quickly summarize here. *TypeTrax* and *Cumulus* document each instance of a font used in a particular document or an illustration, profiling it as an asset. Does this mean that I can go into my *Cumulus* catalog and say, 'Where and in what files, documents and PDFs have I used this font?'

VAD: Right. Exactly.

MM: Does it let me sort on the instance of a very specific font, in all of its documents. That's really important. Because in many cases, I discover that we have an illegal font family.

VAD: Yes it does.

But even more importantly, *TypeTrax* will now go out and get that font, and bring it into a protected area in *Cumulus*. So, after several years have passed – even if you wiped out your local system – if you search in *Cumulus* for a project or document, and you open it with *TypeTrax*, you're going to get the fonts that you actually used during the creative process.

In fact, we work so closely with *Font Agent Pro*, if you try to launch the InDesign document, *TypeTrax* will query *Font Agent Pro* to ensure that the proper font is loaded. All this takes place while the end-user is seeing the progress bar that states, 'Loading Fonts.' So *TypeTrax*, working together with *FontAgent Pro*, is actually loading and activating the fonts, and then releasing the InDesign document.

The designer never again sees that annoying dialog box that reads, 'Missing Fonts'.

CALLING ALL FONTS: WHERE HAVE YOU BEEN?

MM: It also then would allow you to say, 'Okay. This one particular user in our Boston office is using the wrong fonts. Therefore, we need to contact him and say, 'Hey, we've moved off of that Monotype over to an Adobe family'.

VAD: Exactly. *TypeTrax* can also execute an 'Export project' command, taking all of the

project elements, Michael, that we're talking about, including the model releases, etc., and put them all into one package – one folder – for the designer. *So the end-user can deliver it anywhere.* Because we have everything in Cumulus, we can take it right out, and hand it right back to them at any point in the process, and all version of every element in the project is also maintained.

MM: That's a great way of then archiving a project, and then encapsulating it so that the project now has the ability to move across a digital network and have work done on it without any missing pieces.

VAD: Yes. It addresses one more really important thing. *Every company* that I go into – *has archived years of work, often off-line on CDs or tape* – hundreds of projects, in the manner we've discussed earlier.

This means that the *firm has the same fonts and images stored in each project folder*, hundreds or thousands of times. So when after 10 or 15 years, and they want to put their projects into the new DAM, they discover just how many *hundreds or thousands of duplicate fonts and images* that they just ingested and clogged up everything.

Because all of these folders come in from different locations, *the typical asset management system processes all the fonts and the same images each and every time a project folder is ingested.* TypeTrax controls the process, it only lets it in once – whether it be a font or an image – and link each page where it's used to the single asset stored in the database. This is a very, very important functionality.

PDF AD APPROVAL SYSTEM

MM: Let's get back then to the capability that you deployed and continue to enhance at Newsday. You talked about the PDF ad workflow. What are some of the other innovations that you brought in?

VAD: We talked about the pre-filler – the Metadata Manager – that we developed there. So let's talk about the *PDF Ad Approval* system that we developed for Newsday.

The creative team at Newsday builds thousands of ads each week. As they build these ads, the creative and production teams ingest the ads and media components into Cumulus. James wanted Moksa to develop an ad-tracking

plug-in to Cumulus that provides information on the time spent by a given designer on a particular ad. Each ad needs to pass through an approval cycle. Moksa further developed the 'PDF Ad Workflow'. When an ad is ready to begin the approval cycle, all the *designer has to do is select the 'Publish Ad' menu option* and the 'PDF Ad Workflow' system will send it to the InDesign Server (Moksa is an Adobe InDesign Server developer). Working directly through the APIs of the InDesign Server, *a PDF file is created.* The PDF file is then published to a Web-enabled Cumulus catalog. All required metadata is also transferred to the PDF's file asset record. The system then *sends an email to the appropriate sales individual*, with a link to the catalog record, indicating the ad is ready for review and approval.

The salesperson (or client) simply *clicks the link embedded in the email.* That brings them right into the published catalog where they see only the PDF under review. They can then *annotate the ad, submit it*, and, based on whether or not changes are required, the 'PDF Ad Workflow' system sends it back to the graphic designer who originated it, or marks it as approved.

So you have a loop that continues until the ad is approved. This process eliminates the need to notify sales of the ad statuses, as every step in the process is captured in the ad record. Cumulus keeps each version of the ad; if the client decides to return to an earlier version, the system can do that instantly. This automated process really expedites the ad workflow at Newsday, and really cuts down the time it takes to get an approval.

MANAGEMENT REPORTING ON ASSETS, ACTIVITIES AND USERS

MM: What other sorts of things, in terms of automation? Either in terms of image transforms, video encoding, file distribution?

VAD: Another thing that Newsday, like *every production department, needs to do is generate reports.* It's extremely important to track the time spent on each ad. We developed another extension that has helped Newsday track who's working on what files. We do this by dynamically tracking the time the file has been 'checked out' of the system. James Kober can generate this report, *showing how long an ad has been worked on, who worked on it, when they did it* (through

comment fields) and *what shift the ad was worked on*. So James can then begin to orchestrate resources within his department.

AD-BUILDING SERVICE

MM: Would you explain an ad-builder, where a design team puts together a template – whether it's in an InDesign server or not – can enable channel partners, dealers, resellers and field sales people to customize a pre-built brand-consistent ad-flyer or direct mailer?

VAD: That was something not done for Newsday, but we did help Trader Corp. in Canada accomplish. *Trader Corp. is a division of Yellow Pages*, where we worked out of the Montreal office.

You can imagine the tremendous number of ads that are submitted over the Internet. People are selling boats, cars, tractors, equipment or whatever. In this instance, Moksa brought in both Canto partners Modula4 and XLent Software to work with us. (Another example of leveraging that eco-system we spoke of earlier.) Together, we really expanded the capability of Cumulus. We monitored the directories on the Trader Web-based ad submission system, at a designated interval (every 10 min). *We parsed the website directories and picked up all of the images and ad metadata received*. We'd bring that all into Cumulus as ad components – a single ad might have several photos involved. Think of an auto ad where the seller would want the front, back, sides and interior of the car displayed. We store all these photos in designated fields in a single Cumulus record. We'd then *send it on to a template-based ad-creation system*.

The strength of the system came when there were correction cycles. In these instances, where a simple template-based system would simply return to the untouched original photos and start again, using a DAM solution like Cumulus, the retouched photos could be supplied to the system. It made for a much better process.

PAGE-BUILDER FOR CATALOGS

MM: Would you explain things like automatic page-building for catalogs and other sorts of things that tend to have a lot of labor associated with it as well as many design elements that someone has to manage and track?

VAD: Sure. That's the major reason you want to use a DAM system. It's the major reason

why Yellow Pages/Trader Corp. made that decision to go with a Cumulus DAM system.

Often, advertised items don't sell immediately. *Sometimes it will take several weeks of running the ad*. So, the way some photos originally come in, say an automobile, you may need it touched up – to remove distracting background features. If the ad repeats, you may want to display the ad on a different page altogether. If the graphic designer has done *a lot of retouching on the images*, you *certainly don't want the design to go back to the core image* that was downloaded from the website (still with the distracting background feature or whatever).

With a DAM, you can have your original assets along-side retouched versions, especially the most recent or the 'current' version, as we'd call it. The DAM can *'hand' that corrected version of the file to the automated page-builder* or a manual page-building team using InDesign or Quark. Thus, the DAM manages the workflow and the most current version of the asset. *That prevents a lot of rework*.

ASSET VENDING MACHINE

MM: Would you explain an automated submission process using Cumulus, and the kind of value you'd anticipate a company like Newsday would get from using DAM to manage the inflow of a lot of new photos from the field?

VAD: That is a very important and special feature in Cumulus 8.0. Let's imagine the following scenario. In the past, *people have had to go to the DAM system and really be pretty much on their own*. Do searches. Try to find items themselves. Often, make multiple searches in order to get all the components that they'd need for the production of a page or a slideshow or anything that they're looking to do.

With an automated submission process, you simply fill out a form for the items that you'd like to have delivered to you. You then simply submit that to Cumulus 8.0, and the entire order that you have – just like you ordered it online – arrives to you through email or through a Web front-end. So you can have it, manipulate it and do whatever you'd like with it. It really is going to revolutionize the way we look at a DAM. Finally, the *DAM system will become more of what it was meant to be: an appliance or vending machine*.

A major problem in almost all of the DAM systems – that I’ve worked with up until this point – is the placing of *unnecessary burdens on the end-user to perform even the simplest of tasks*. As an example, assigning user rights, or simply ingesting an asset with metadata: users are *forced to navigate multiple screens and populate multiple fields*. Cumulus 8.0 really simplifies and automates a lot of this.

For instance, you could bring an asset in and assign metadata. You could do transformations. You could *do a number of what they’re calling triggers-and-actions within Cumulus 8.0*. So you could get it fired off at one time. So instead of having to have your people work through each one of these tasks manually, scripts or macros help to automate ingestion and tagging of assets.

ENTERPRISE DAM SOLUTIONS, EXAMINED

MM: As we begin to move toward the end of our session, I’d like you to compare and contrast Cumulus with other DAMs in the marketplace. Specifically how do you compare Cumulus with the high-end DAM systems from North Plains, OpenText or EMC Documentum?

VAD: I’m very familiar with Artesia *TEAMS*, now owned by OpenText. I was on the product management team for almost 2 years. I have also worked extensively with Documentum, as well – if you’re talking about the top-down solution we covered earlier, these are probably No. 1 and No. 2 in the industry today.

Let’s start with Documentum. I used them in providing a solution for a large, well-known medical association in the Boston area. We really looked at the solution and saw that it had a lot of functionality that they required in a DAM and workflow system. But when it actually came time to build it, *we had to build-out most of the real media services and workflows we thought were out-of-the-box*. The same holds true for Artesia as it does for Documentum.

So, in both cases you’re stuck using a *robust but unconfigured framework to build a solution from the ground up*. You have to build each and every component, and make it fit your client’s needs. Both Artesia and Documentum developers believe that they will put together a system that the client wants and needs. It’s just so difficult and involved that it takes forever to patch it together and it’s EXPENSIVE to this!

As outside consultants, we brought these products to our clients. I knew what I was looking for. We saw the product demo, or thought we did. Problem is, the team that comes to implement the solution is not the team that presented the system. *The new team knew almost nothing about the requirements* the first team took months to study and understand – just to build the demo.

So from the get-go, I had to get the second team, the deployment team, to understand how to utilize the tools in Documentum that the first team, with deep knowledge of the account, demonstrated and promised: *how to accomplish the kind of functionality that the first team promised the client*. In the case of Artesia, the same thing happened with a large media company. In both cases, the end result was a disaster. It just wasted months of time and lots of money.

Cumulus is completely different. You really do have a core set of functionality that simply needs to be configured – not customized. What do I mean by ‘configured?’ It means I open windows within Cumulus and I click to control how the system behaves.

I really can work with the core functionality of the system. And then, where there is functionality that’s not in Cumulus – and we all know this right up front – then I need to develop either the Java server plug-ins or the EJPs required to extend the system to *accomplish the workflows that I want to accomplish*. It’s a completely different approach.

MAC-BASED ENVIRONMENTS

MM: Because the creative workflow largely remains a Macintosh-based environment, would you speak to how these larger enterprise DAMs handle Mac clients in general?

VAD: Well, you really hit the nail on the head, there – this is a significant point. In the case of the Documentum installation I spoke of earlier, the client *emphasized Mac support from day one*. In reality, the system had little real Macintosh support. The Documentum system is really *optimized to work with the Windows platform*. Test results showed a 25-to-1-speed advantage when using the Window systems over a Macintosh system – and this was purely because of their lack of understanding and support for the Macintosh platform. They used a number of

excuses to explain the problem, each making absolutely no sense to any of us.

The problem really is this – these systems have never had support for working at the desktop application level. Here is where solutions like Cumulus are far and away the best solution. Interacting with products like the Adobe Creative Suite is what products like Cumulus were built to do. While Cumulus offers a simple thin client, it has a robust thick client as well. This client can take full advantage of the capabilities of either the Macintosh or Window's desktop work.

So, these larger systems really have a *tough time with very large media files*, especially in creative workflows with lots of works in process. You're really much better off with a solution like Canto Cumulus, which really works well on the creative desktop – which is a Macintosh application on a Macintosh platform. It integrates completely with all your other desktop applications, at a drag-drop level. Canto Cumulus just becomes a menu extension in many of the applications you'd want. It's completely integrated with your desktop.

MEDIA WORKBENCHES

MM: That gets to Cumulus as native, mature desktop applications and with a great plug-in architecture, where you can start with basic DAM functions and capabilities. Later when you add plug-ins, you can quickly customize workflows and interfaces to support very specific work cells. That's neat because innovation really happens from the bottom-up. So, when DAM becomes tightly integrated with the other tools and utilities that the designer has come to depend upon on a minute-to-minute basis, bottom up innovation lands as a media workbench.

VAD: As always, Michael, you stated it perfectly. That's exactly what I was trying to say.

MM: Would you compare and contrast Cumulus with mid-market DAM systems such as Xinet Web Native, MediaBank from WaveCorp, and Media Beacon from BrightTech?

VAD: Yes. One of the most advantageous aspects Canto in the mid-realm relates to is a product called the *Cumulus Workgroup edition*.

That means that it's built for a smaller group or a smaller company. The Enterprise version of Cumulus remains very feature-rich and has a lot of complexity. It benefits from having a true administrator. *You can start small and grow*. At some point, you grow into an enterprise DAM but you don't really change the user experience – unless you want to!

PROCESS-INNOVATION PLATFORM

MM: Great, that leads into how Cumulus works well as a process-innovation platform. Would you explain how plug-ins, a developer ecosystem, and terms of license for server support ongoing innovation of new processes?

VAD: It really starts with the *licensing scheme of Cumulus*. You can start small in a Workgroup configuration that fits the budgetary constraints of some of the smaller organizations and most pilot projects of larger firms. Canto also provides an Enterprise version: more functionality, more users, larger servers, more catalog administrators etc.

MM: How does this compare with Extensis Portfolio?

VAD: I really don't think that Portfolio should be considered as being in the same category as Cumulus. Why? Portfolio doesn't really have an enterprise database at its core. They'll say that they offer Microsoft SQL Server or Oracle databases. But really, Portfolio just parks metadata. Cumulus has matured into a specialized and dedicated DAM metadata database.

METADATA DATABASES NOT ALL THE SAME

MM: Can you explain how DAM metadata databases differ from relational databases used by other DAM applications?

VAD: Sure, *other DAM applications process metadata differently*. As the assets come into Cumulus, everything is fully indexed. Everything around that asset is put into the database in such a way that you get an instantaneous response. Even with hundreds of thousands of assets.

You get a *return to a search instantaneously* because the database was optimized to do that. With other DAM applications using an external relational database that processes metadata in a different way, you often wind up having extended

loops. Remember, a secure search takes an extremely long time, especially a Boolean search or a complex search.

I've done comparisons to the database, and have found complex searches on an Oracle or a SQL Server to take 10 or 20 times longer than with Cumulus. A really big difference that ultimately means that most users give up trying complex searches and waste lots of time doing simple searches.

OPTIMIZED FOR METADATA

MM: We find the same thing in comparing various business intelligence data warehouse and data visualization tools. It really sets up the very interesting technical contrast of a relational database, where you're loading thousands or millions of rows of data with hundreds of data fields, and a columnar or analytic database that load just 10 columns of data with a few bytes of data from the same thousands or millions of records – many orders of magnitude in difference.

VAD: Because *it's so completely indexed, the Cumulus database lets you work with the collections you're working with. It just doesn't go running off and bogging down.*

In particular, the way Cumulus ingests new assets becomes very important. In most DAM systems, you start by digitizing assets (creating thumbnails) and then apply metadata to the thumbnail and underlying database record. In Cumulus, you just drop an asset or an entire nested folder of assets into a collection bucket, and you've tagged it. A simple drag and drop into other buckets adds more metadata. Thus, you can get creative folks, normally who hate tagging the assets by typing in metadata, to tag away using drag-and-drop tagging of Cumulus.

PRE-CACHED COLLECTION, INSTANTLY

MM: So an asset catalog administrator could look at high instances of duplicate searches and go, 'Wow. For whatever reason, people are searching on this particular kind of collection. Why don't I create a shared search query?'

VAD: Yes. You can have an administrator intelligently package information about how their system is being used in a *set of collections, pre-built and ready to download.*

MM: And that the collection of search queries simply appear to users as a menu item or an icon within a palette of quick searches.

VAD: Correct. Exactly.

MULTI-THREAD OPERATIONS SUPPORTED

MM: Could you then speak to the notion of multi-threaded processes? What does that mean, and why is it a big deal?

VAD: The biggest drag on a database is opening 'write' threads to the repository. A thread constitutes a single user trying to write data back to the database, sending information back and forth in the data fields. Each thread sends information to the repository.

In short, the application can direct a process to a specific processor, or group of processors.

Let's say you have a multiprocessor box. Eight processors. I'm going to allocate these three processes to a specific group of processors. These two processes are less process-intensive, so I'm going to put them through that single processor. This capability allows me to fully optimize the hardware available.

It allows you to determine, based on the way your users function, *how to customize the processing of your assets in order to achieve optimal performance from your hardware.*

MM: Thus unlocking additional speed and value of these two, four or eight microprocessor-based servers as well as offloading some operations to special graphic processor or vector-processing chips right next to the general-purpose microprocessor.

VAD: Precisely, Michael, and that's why reporting is so important. Because the *more information you get on how your people are working, the better you can allocate physical computing resources* (servers, storage, routers and bandwidth). Again, Cumulus will take even greater advantage of multi-processing capabilities when version 8.0 ships in summer, 2009.

IT SERVICE MANAGEMENT WITH LINUX VERSUS WINDOWS

MM: I'd like you to speak to the IT infrastructure – and more specifically, the service delivery and service management of DAM-enabled workflows and applications. First, speak to the importance of having Linux and Solaris as part of your overall IT infrastructure.

VAD: That's really a very important consideration. In fact, my firm will help Newsday move from Windows to Linux servers, for the specific requirements that you just outlined.

In particular, *firms like Newsday that really, really need to have system failover happen dynamically*. This requires that you use logical structures in a network. So when you have failures at different areas of the network (such as a storage array or a server going offline), you can still maintain continual uptime. To achieve that, you need to have other resources pick up where resources failed or went offline.

For instance, the asset storage system might have gone down. Maybe one of the asset catalogs went offline – where you have a database structure for the asset repository. You need *the ability to have a second database come online in an instant* as the backup database. So without skipping a beat, your users who are so dependent upon a system don't notice a thing.

You know, graphic designers in a production department really become dependent on the uptime of workflow and DAM systems. To have it so there's absolutely no down time. Backup, failover, mirroring and all of these things are critically important. They come to you when you start to use Linux or Unix.

MM: So as a function of being into Linux, can you then explain why having a Linux and Solaris enhances overall IT service management?

VAD: It's really very practical in a business sense. For years, Microsoft has held large organizations hostage. You know, at a large medical center that we talked about a little

while ago, *Microsoft extracts huge annual license fees* – the renewal fees to support hundreds of essential servers. Once you get yourself locked into a supplier like Microsoft, *costs become just horrendous*. I think that as consultants, we need to make sure that corporate clients understand the Microsoft excise tax. We need to remind corporate clients that *OpenSource software such as Linux, MySQL, Apache and PHP really becomes a get out jail free card*. Open Source gives you flexibility and payment options; it keeps vendors honest in the game.

BEYOND SERVER VIRTUALIZATION: BACK TO BASICS

MM: What other final thoughts do you have about service management of DAM operations, especially as it might relate to other IT executives who may not know a lot about digital asset management?

VAD: I've found that most of the companies struggle with a *tremendous number of security patches*. As a result, they have to maintain huge number of IT specialists to support complex Windows environments – especially if they use *virtualization* – with hundreds and thousands of servers. *A single Linux or Sun server could replace many, many of these servers*. When you use Linux or Sun, it really simplifies your life and reduces cost. Why? Because *Unix stays up for so long*, it maintains itself so well. Linux or Sun reduces a lot of maintenance requirements that you experience on the Windows side of the fence.

MM: Sounds like a great place to conclude. Thank you.