

About the Author:



- ≡ Ralph Windsor is Project Director for Digital Asset Management Consultants, Daydream [☑](#) and also editor of DAM News.
- ≡ He has worked in the IT industry since 1994 as a software developer, project manager and consultant.

Building and managing scalable DAM solutions

Digital Asset Management [☑](#) strategies are never set in stone. High scalability [☑](#) and flexibility are crucial, enabling users to adapt the system for future requirements. In this white paper, the DAM expert Ralph Windsor explains, what needs to be considered when it comes to implement a sustainable DAM system.

Introduction and context to DAM scalability

By any measure, the rate at which new digital assets are being created is increasing at an exponential rate. This mirrors the growth in the volume of data being stored in general and is a function of processes being transferred from older analogue methods of production to more flexible digital alternatives.

The increased use of digital media delivery devices like tablets, mobile phones as well as distribution via websites and social media are also key drivers.

Although the reasons behind the growth in demand for DAM solutions are easy to understand, the consequences are anything but. Digital Asset Management systems are becoming central hubs for a diverse range of digital media which may be used across organizations for a multitude of purposes.

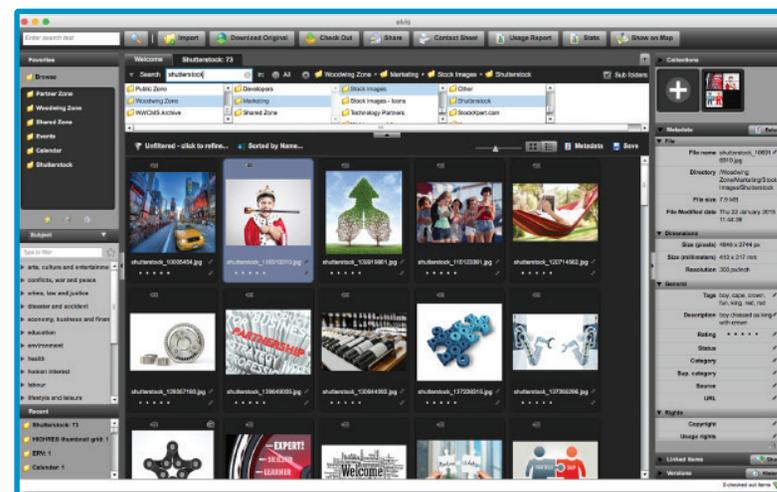
They are the media equivalent of airports, railway stations or shipping ports and rapidly acquiring all the logistical challenges that the efficient operation of those facilities demands.

Despite the rapid increase in asset volume and the critical role that DAM solutions will hold, the issues about scaling them up to meet high levels of demand are not always as widely considered as they should be – at least until it becomes too late to deal with

them without causing disruption. As with a transport hub, it is not enough to build one that just has a lot of capacity, they need to be carefully planned, skillfully operated and continuously monitored so that incremental improvements can be introduced.

1. Scalable architecture and risk analysis

One factor with scaling any system (not just Digital Asset Management) is that in nearly all cases, the software architecture



Given the ever-increasing stock of digital assets, scalability in every aspect is an important trait of digital asset management systems.

“

Risk and scalability are closely correlated to each other in DAM systems architecture. There should be transparency, openness and collaboration between vendor and client for this to be successful.

”

has to be made progressively more modular. This makes it more feasible to exchange one component with another that is more flexible, has greater capacity or possibly both those attributes.

Many scaling issues can arise due to a lack of granularity where one process is composed of a number of elements that are tightly connected together. This requires the replacement and/or re-engineering of all of them and that is usually the root cause of unexpected costs and disruption for users.

To avoid these problems, it is essential to have a detailed model of how assets will arrive, get processed and then be made available to asset users. This needs to be broken down in some depth, even if they are aggregated as descriptions to senior stakeholders who are less interested in the technical details.

While these are mostly considerations for the implementation team (on the customer and vendor sides) project management personnel from the customer should have some awareness of what is happening behind the scenes. This knowledge enables them to assess risk and work with the vendor to develop mitigation plans for them or highlight potential concerns.

For example, if assets are likely to be required by a lot of Asia Pacific users, but the servers are all located in North America then there is a performance risk which will increase as the volume of assets and users increases.

A comprehensive log of scalability risks must be maintained; not all of them may be considered worth mitigating yet, but it is essential they all are understood and whether the probability of them is rising or falling.

Risk and scalability are closely correlated to each other in DAM systems architecture. There should be transparency, openness and collaboration between vendor and client for this to be successful.

If the product and/or implementation team have used hacks and shortcuts to get some key processes working (or have been given a budget to work with which is insufficient for the requirements) then scaling up a system will frequently find those sins out.

This does not mean it is incumbent on the implementation team to aim for perfection (as that would increase costs for potentially limited benefit) – but it does mean that if there are trade-offs, they should be up-front about what they are and pro actively identify any new risks that might be introduced as a result. These should be recorded and continuously assessed for probability and severity of impact.

2. Major scalability pressure points – asset handling

The second part of this series will give a brief introduction to the major pressure points likely be to encountered when scaling DAM solutions. This is by no means an exhaustive list, but it should prompt some ideas and consideration of areas that might not have been fully examined previously.

Uploading of assets (accession)

This is the point where digital assets get introduced into the system: uploading (or 'accession' to use the more formal term). With web-based applications, many DAM solution developers who lack experience of scalability issues provide a file upload control and a single operation which transfers the asset to its storage location.

This creates problems when either the range of methods to receive files expands (e.g. FTP, e-mail, direct imports from a file server etc) or the number of simultaneous uploading users becomes too great.

A more scalable method is to use one or more interim holding areas that all files are received into before they are dispatched to a long-term asset file storage facility. Most scalability techniques

depend on highly modular design principles being rigorously applied and importing assets is no exception.

Taxonomy & metadata scalability issues

An entire series of articles could be written about taxonomy [↗](#) and metadata [↗](#) scalability issues, but here are a few example questions:

- Is the taxonomy still usable after numerous nodes are added, especially if these are nested many levels deep?
- Will multiple taxonomies need to get integrated into the one for the DAM, or perhaps even will a separate taxonomy be introduced for a given asset type (or some other arbitrary criteria)?
- When controlled vocabulary [↗](#) entries are added, what effect does this have on the usability [↗](#) of the system (for example, if the number of entries in a given field increases from 10 to 200)?
- Is there a governance or management process to determine how terms are added and who needs to manage that? Will requiring these to be checked each time create bottlenecks? What is the risk of not checking them? The trade-off between both risks is either slowing asset approval throughput or allowing unsuitable assets to be released (with potential business impacts as a result).

Systems architects and stakeholders need to liaise closely with each other to assess which risk has the most negative consequences. In some cases, it might be necessary to sub-divide governance based on other criteria (e.g. the intended usage for assets).

To deal with this scalability pressure point, the delivery team needs to envisage what will happen when the range of metadata grows exponentially.

The impact on DAM system metadata of having one million as opposed to 1000 assets might be very significant and render other

parts of the system (e.g. search or user interface) to become unfit for purpose. These impacts need to be considered in advance, including their consequences and any plans to mitigate them.

Asset growth & findability Issues

This is closely related to the metadata and taxonomy topic described previously. The nature of how users interact with DAM systems changes according to the number of assets they contain.

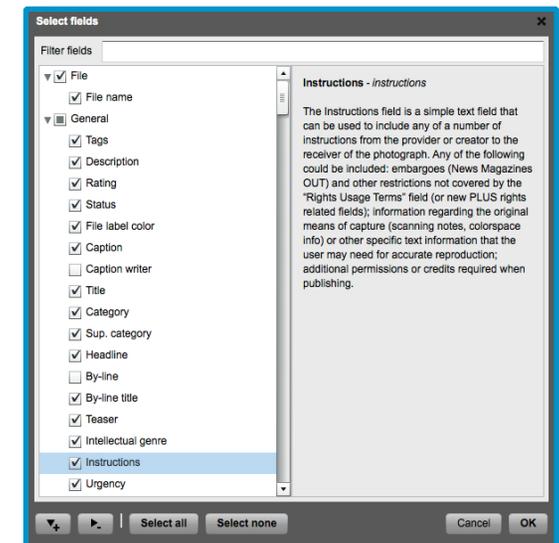
Most DAM systems narrow searches when users supply more terms to queries, for example a search for "dog cat horse" will produce less results than just "cat" because all three terms must be present in the former example.

Shortly after a system is launched (when it usually contains fewer assets) users will search with single word queries because they quickly realize that using longer queries frequently produces no results.

“

The nature of how users interact with DAMs systems changes according to the number of assets they contain. As more get added, single word searches begin to generate an excess instead and users gradually switch to multi-word queries.

”



The metadata capabilities of a DAM application are also relevant for the scalability of the system.

As more get added, the single word searches begin to generate an excess instead and they gradually switch to multi-word queries. Some may also make use of advanced filtering tools that they previously had no need for.

When most DAM systems get initially tested as part of an acceptance process, asset growth tends to not be part of the test criteria, however, it can have a dramatic impact on the user's perceptions of the system.

Ideally a good number of assets are either migrated or ingested from the start, but if that is not possible, tests with dummy assets or samples should be carried out using a mix of real and random data. It is essential to get a good idea of what the system will be like for users when it contains both very few and a lot of assets to maintain usability irrespective of the quantities involved.

Asset search

Until recently, many DAM solutions employed relational databases for both asset metadata storage and searching. When users initiated a search query, the database only was checked for matching terms.

Relational databases [are](#) good for structured data such as categories but they tend to be less effective at larger volumes for text search and other natural language queries.

Recently, database concepts such as NoSQL [have](#) been embraced by some DAM software developers. These include search technologies have been specifically developed to handle a much larger repository of assets via products such as Solr [, elastic-search](#) [and MongoDB](#) [. These technologies can cope better with search scalability challenges.](#)

DAM systems that need to be able to process millions of assets and allow users to perform searches across all of them can benefit from integration with them, however, they can increase com-

plexity, for example, merging results from multiple sources and this is a risk factor that needs to be acknowledged and mitigated.

Generating proxies & renditions

Nearly all DAM systems will now generate at least one representation of an asset, for example a thumbnail so that users can see what they are about to download and decide if it is suitable.

These get referred to by a variety of names such as 'surrogates', 'derivatives', or 'proxies' (which is the one I tend to use). In addition to thumbnails, a range of other proxies might be needed such as larger previews at different sizes so users can zoom in.

As well as proxies, many DAM systems support various forms of asset manipulation or 'renditions', such as resizing, cropping, rotation etc for images. An even more sophisticated range of options can come into play if the underlying asset is a video or a complex document such as a PDF or InDesign file.

Typically, the core DAM system itself will not provide these functions, instead it will provide instructions to some associated third party component that processes the file. But, that does not mean



Search engine powered by

elasticsearch.

Logos of companies using Elasticsearch: METFLUX, Path, Jira, github, foursquare, W, SONY, yelp, SOUND CLOUD, b: XING, subforn desk, IGN, mozilla, Bloomberg, theguardian, stackoverflow, KLOUT, infochimps, SCOUT 24, The New York Times, NETWORKED INSIGHTS, Fog Creek, Revinate, picturesafe.

5 ways to improve managing your assets in today's digital marketing world

elasticsearch is one of the technologies that can cope better with search scalability challenges than relational databases.

the system is automatically scalable. These operations are usually resource-intensive and will utilize a higher proportion of the server's memory and processing capabilities.

If they are running on the same device as the DAM system application then they might cause the system to slow down for all other users. To avoid this, a scalable solution should allow the permit other servers to be dedicated to this task and to signal when they have completed operations. To make this possible a separate renditions generating service is needed which can be hosted on different hardware and potentially at a completely remote location away from wherever the core system is held.

3. Major scalability pressure points – IT-aspects

From storage to security to integration and more – the third part of our series cover the more IT-related aspects of scalability.

Asset storage

Asset storage is probably one of the better understood DAM scalability issues, but can still catch delivery teams out, especially where the volume of assets stored increases to a level not anticipated in advance.

This can be due to a variety of factors, for example, a new larger group of users being introduced to the system. More recently, the switch from SD to HD video caused storage requirements to increase for video-oriented repositories.

One architectural technique used with scaling up asset storage is that each asset record needs to be able to point to an independent storage location, which may be shared with others, but could be totally separate if desired.

This allows the location to be changed later and the data storage logistics to be carried out behind the scenes without interrupting users (or organizing outages to allow time for assets to be

copied). Further, it allows new storage repositories to be incrementally added as needed rather than everything being held in one large container.

If the DAM system is accessible over public networks, another development more prevalent now is to use Cloud based storage including Amazon S3 [↗](#), Google Cloud Storage [↗](#), Microsoft Azure [↗](#) etc. In this case, the assets are held on a remote server that is hosted separately from the DAM system and can take advantage of the very high capacities that the provider of the Cloud services can offer.

Those benefits are partially offset by additional complexity in managing security and ensuring access is controlled. Most of the storage providers offer pre-rolled methods to address these concerns, but these need to be integrated with whatever asset file delivery mechanism the DAM uses (and options to switch between methods and/or providers might also be required).

Authentication & security

Many DAM applications are now required to integrate with some enterprise-wide security and authentication [↗](#) system, for example, Active Directory [↗](#). If the organization using the DAM operates globally in many different regions or they have recently acquired another firm, they may have several of different instances of these coexisting with each other.

This can make the integration process more complex and the DAM system might need to switch methods based on some characteristics of the user requesting access.

This is another area where modularity is essential for scalability. If the system has been designed in a component-oriented manner then it might just be a case of adding different modules for each one. If everything is self-contained inside a 'black box' then someone will have to open that up and alter it (with the potential for bugs, disruption etc for existing users).

“

Authentication and security issues can make the integration process more complex. This is another area where modularity is essential for scalability.

”

“

It is essential to get some idea of metrics (i.e. at what levels will bad things happen and how close to them are you now, or might we be within a given period). These need to be more than 'finger in the air' estimates.

”

Scaling integration

It is rare now to find modern enterprise DAM solutions that are not integrated with at least one other third party system. Integration-related scalability issues can arise from either end of the process and the impact of poorly planned implementations can be similarly multi-dimensional.

If the DAM is receiving data from a third party system, scaling issues can unexpectedly present themselves if the volume of data introduced increases rapidly without warning. To give an example, a DAM application might receive records of sales leads from a CRM application which assets get associated with during cataloguing (so users can find all assets used in pitching materials for a given sales lead).

A simple method is to receive a daily file and then re-load everything to ensure it is up to date, but if the number of records increases then the update might never complete or have other knock-on effects.

A more scalable method is to issue these on-demand via a dedicated service or alternatively a 'delta' file (i.e. only the modified records). Both latter options are more complex to implement, but are also more scalable.

The other side of the coin is the DAM supplying assets to another system. If assets are pushed to something like a Web CMS for authors to use and there is an unexpected rise in volumes, this could have effects on the corresponding system and cause it to become unstable and might even cause a public website to stop working.

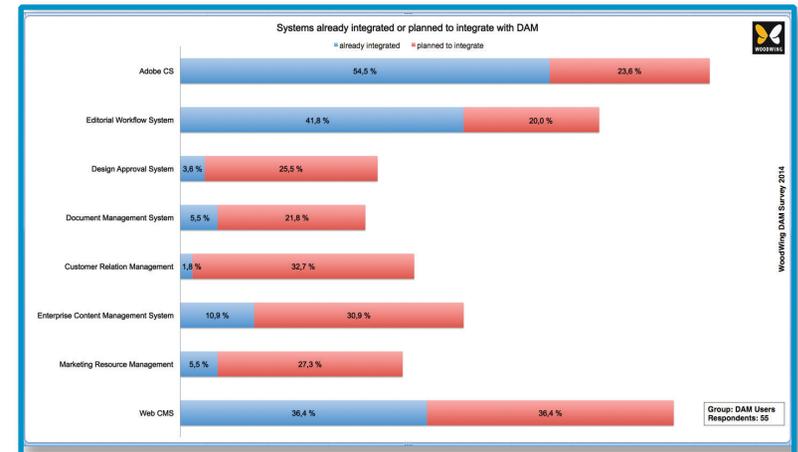
As should be apparent, a risk analysis has to be carried out when DAM systems get scaled up – but this is particularly crucial with integration since the consequences can occur in at least two different applications (and sometimes even more if those solutions in-turn issue data to further endpoints).

Analytics to gain wider insights

It is essential to get some idea of metrics (i.e. at what levels will bad things happen and how close to them are you now, or might we be within a given period). These need to be more than 'finger in the air' estimates that the "dev team" invent on the spot in a project meeting, they have to be tested so there is some scientific basis for them.

In my experience, the majority of scalability metrics evaluations are never checked and are often best-guesses, because there is not the time nor resources to carry out the task in any great depth. But if the stability of your implementation (including any integration with other systems) depends on data volumes remaining within a given threshold and the business impact is significant if they are exceeded, you must carry this out.

It is almost certain you will need to keep reviewing the numbers you collect and continuously check that any process which depends on remaining within those same thresholds will not silently fail and/or create a weak link in the digital supply chain as you scale up.



Most of DAM users have integrated their DAM system with at least one other system or plan to do so (Source: WoodWing DAM survey 2014).

“

Scaling up DAM services offers an opportunity to generate significantly higher ROI and often is more cost-effective than wholesale replacing solutions –with all the upheaval for users that this implies.

”

Service Oriented Architecture [☞](#) and asynchronous applications One of the other core principles of scalable systems that they are 'asynchronous'. In simple terms, this means you can commence a process (e.g. batch resizing thousands of images) and they then allow you to do something else on the same system (with a notification either on-screen or by email etc when your request is complete).

By contrast, synchronous systems require the user to wait until they have completed before returning control back to the user. As well as proxies or renditions, this can apply to other intensive operations such as zipping files or generating large reports.

Designing and developing asynchronous applications is more complicated than synchronous ones because some other process has to reside in the background and monitor what is going on (amongst other essential tasks). Systems designers usually refer to this as a 'Service Oriented Architecture' (SOA [☞](#)). It means not only is each core component modular and replaceable, but they can be run independently of the core DAM as quasi-independent services.

One other major benefit of SOA is that there is also the possibility that a given service can be used by some other application also, so the range of functions to which the DAM can be put to use expands considerably.

This can come into play with the integration topic described earlier. If your DAM system is required to manage large-scale automated on-demand re-purposing of assets then this is a significant benefit and takes the range of uses for your DAM to a level above cataloguing and search.

Conclusion

As should be apparent, the topic of DAM scalability is a highly diverse one with many different interrelated subjects to consider. The act of modifying one component of your DAM architecture creates consequences that may necessitate changes elsewhere also. It is rare that you can just upgrade one element and be done with it since the issue limiting capacity will just move further down the digital asset supply chain.

It is instructive to view digital assets in larger systems as you might flows with liquid characteristics (e.g. money, people or water). Having an overview of asset data movements which is appropriate for the level of detail under consideration is helpful for everyone involved since it makes it easier to understand the context of any problem.

It is preferable to use a continuous improvement approach to scalability, which should be informed by an intimate understanding of what happens to assets across their entire lifecycle (especially as that relates to the DAM).

This knowledge is what will provide the insight required to be able to propose resolutions to scalability challenges with some confidence that they will be effective and enhance long-term ROI rather simply incurring unnecessary expense.

Scaling up DAM services offers an opportunity to generate significantly higher ROI and often is more cost-effective than wholesale replacing solutions (with all the upheaval for users that this implies). With that said, those opportunities are not without manageable risks, so you need to be aware and in control of them as soon as a DAM initiative is first commenced.

About WoodWing

WoodWing Software develops and markets a premier, cost-efficient multi-channel publishing system, Enterprise, and the next generation digital asset management system, Elvis DAM. WoodWing's solutions are aimed at magazine, newspaper and book publishers, corporate publishers, agencies and marketing departments to reach their goals for quality, economy and time-to-market.

WoodWing's publishing system Enterprise – including the editorial management application Content Station – coordinates and streamlines the process of creating, managing and publishing static, dynamic and interactive content for all media channels – print, Web, social, smart phones and tablets. Elvis DAM enables users to securely store and efficiently manage the increasing collection of rich-media files.

WoodWing Software, founded in the year 2000, has its headquarters in Zaandam, The Netherlands, and has regional sales offices in Europe, the Americas and Asia Pacific. Customers are served locally by over 80 selected partners in more than 100 countries. WoodWing's long-standing relationship with Adobe as a Technology Partner and its close cooperation with a large number of other technology vendors worldwide, confirm WoodWing's position as one of the leading suppliers of publishing software. WoodWing is a privately owned company, with all founders actively engaged. More information can be found at www.woodwing.com.

