


5.x CMS (Content Management System) - Adobe AEM

- [Why Use AEM](#)
- [Architecture Components](#)
 - [Adobe Experience Manager Shortlist of Project Documentation](#)
- [Technical Notes](#)
 - [General](#)
 - [Which Data Store \(MicroKernel\) to Use](#)
 - [Adobe Acronyms](#)
- [Version Information](#)
- [Download](#)
- [Next Steps](#)
- [References](#)

Why Use AEM

Leader in Gartner Quadrant in 2017, now 7 years in a row.

Architecture Components



Component Front-End to Back	Purpose	Summary	Notes	Key R Mater
Content Delivery Network				
Akamai	CDN (Content Distribution Network)			
Content Management System				
Apache Web Server with Adobe Dispatcher Module	Web Server, Caching and Load Balancing	<p>Caches into static content pages generated by Publisher to improve performance and provides a layer of security protection.</p> <p>Requests from User stops at Dispatcher which initiates a cleaner request to the Publisher.</p> <p>The Adobe Dispatcher Module is proprietary Adobe software.</p>	<p>The Adobe Dispatcher module loads like any other Apache module but is closed source.</p> <p>Note that there are two versions (standard and SSL) with the same extracted name.</p> <p> Is the Adobe Dispatcher module what also receives cache expiry information from the Publisher?</p>	<ul style="list-style-type: none">• T• C• a• In• C• N
AEM Publisher	CMS Publisher	Assembles and serves the published content as static html.		<ul style="list-style-type: none">• In• A• A• C• C• C

Content Repository	CMS Database Cache	See Content. Repository further below AEM Instance for Author for details.		
AEM Instance for Author	CMS Author	<p>Create, upload and edit content and to administer the website.</p> <p>Powered as follows,</p> <ul style="list-style-type: none"> • Uses Apache Sling for Web framework to expose the repository RESTfully. • Bundles CRXDE Lite web interface to manage the repository. • Powered by Apache Felix (implementation of OSGi) running inside of Jetty. <p>Has two user interfaces, old and new (add more to this later).</p>	If possible, ensure only available within the Intranet.	
Content Repository	CMS Database	<p>NoSQL File based data storage for content-centric applications, implementing Content Repository API for Java Technology (JCR 2.0).</p> <p>Uses Apache Jackrabbit as JCR implementation and Apache Jackrabbit Oak as module manager (searching, versioning, transactions, etc.) of repository with additional customization on top of Oak called CRX3 (Code Repository Extreme 3) to add features (synchronous updates, key-value persistence, clustering, static extension, individual workspace index).</p> <p>This finally communicates to a MK (MicroKernel).</p> <p>Default and recommended storage format is to use tar files TARMK (Tar Micro Kernel) which also offers Cold Standby to sync repositories.</p> <p>Alternative is MongoDB but not Adobe does not recommended except for specific use cases. More details on selection further down in Technical Notes.</p> <p>Note Azure blob storage is not supported. AWS S3 is supported with Adobe's S3 Datastore Connector.</p>	<p>There are always limits to any type of file store, for example maximum number of files in a directory. Adobe is not clear on how the product deals with this.</p> <p>Planning should also be done around categories of assets, ie media, images and normal content. It especially becomes a problem when restoring.</p>	<ul style="list-style-type: none"> • T • A • C • 6 • T • C • A
DAM (Digital Asset Management System)	Management of Production ready Digital Assets	As best practice, the DAM connected to AEM should be used for <u>production ready</u> , not high fidelity digital assets.		
Deep Content Management Technology Components				
TarMK				
Mongo		Sharding no longer supported. Only added as of 6.1.		
Solr powered by Lucene	Native System Search Engine	Used by the native system to perform lookups and performs indexing.		
Solr powered by Lucene	Site Search Engine	Bonsai Framework evaluation is to not use embedded and keep the CMS function purely to Authoring. Instead use an external Solr based search engine.		

JCR

Core

MK

	CMS Author Database			
AD or Single Sign-On	Centralized Authentication Directory	<p>In companies with effective identity onboarding it is strongly recommended to leverage.</p> <p>Out of the box, AEM's authentication does not meet minimal standards (enforcing password complexity, password recovery, etc...).</p>	<p>Optional and should fail back to local authentication repository if AD fails.</p> <p>Currently AEM does not support one-time email users. As such decision is to not use AD in a large org to quickly provision Business users in and out of the system.</p> <p>Confirm users deactivated do not take license seats  . Check if deleted users can be recreated with same GUID to more easily pass audit and compliance .</p>	
CRXDE Lite		Adobe's web user interface to interface with Apache Sling. CRXDE Lite looks to be Adobe's own extension sitting on top of Apache Sling Web Console and Apache Felix Web Console .		
Apache Felix	Implementation of OSGi Specification	<p>Can redeploy code pieces without requiring a restart.</p> <p>Each bundle requires a unique name space.</p> <p>Manages service discovery and registry.</p>		
Apache Sling	Control to expose and access resources.	REST Architecture.		
JCR	Object Oriented Database			
JackRabbit		Handles multi version control with refresh function to handle real-time updates.		
Microkernel		Provides an abstraction layer for the storage of content in JSON.		
Vault		Extract out of TarMK for normal file system.		
Jetty	Container	Ahhh... not super clear but in Error 404 starting up noticed it's powered by Jetty.		
Granite				
Roy Fielding	Senior Principal Scientist at Adobe Systems	Significant achievements include, HTTP Specifications (Principle Author), originator of the Representational State Transfer (REST), co-founder of the Apache Software Foundation.		
David N		JCR		

		Note to incorporate more clarity from this article in https://www.terrabeta.com/2015/03/aem-interview-questions-what-is-the-technology-stack-of-adobe-experience-manager/ Also more good stuff in this technical deep dive http://dev.day.com/content/ddc/en/gems/technical-deep-dive-of-the-aem-6-platform.html		
Site Search				
Solr powered by Lucene	Website Search	Used internally for DAM but also available as module for Customer facing website.	Keep distinct from CMS system for Customer facing website. Some products offer as a plugin, but only implement if outside of the CMS system for load and security purposes.	
Development - Author				
Kits		Samples, Tools and Commons		

Adobe Experience Manager Shortlist of Project Documentation

- [Typical Deployments](#)
- [Project Checklist](#)
- [Documentation Page](#)
- [License](#)

Technical Notes

General

AEM 6.x with Apache Sling uses TarMK with Apache Jackrabbit Oak. Key aspect of speed is the combination deserializes the repository into memory. It is not caching the repository in memory. The repository exists in memory and uses the tar files to serialize and persist the data. It will deserialize the entire repository into memory if enough memory available. So give AEM lots of memory.

Though AEM and Apache Sling runs in a JVM, the memory they use for repository is not the JVM heap, but the "off-heap" memory. You may limit it via Java parameter `-XX:MaxDirectMemorySize`. One consideration to gain memory is to set the Pool Size parameter of Apache Sling Job Thread Pool to 1. This will limit the memory used by jobs at the cost of speed of job completion when jobs are queued.

JCR implements a tree of nodes with associated properties with all sorts of pointers (think object oriented database). Persistence manager by default is Tar, transaction aware, file system agnostic and uses Lucene to support SQL queries.

Which Data Store (MicroKernel) to Use

Since most Enterprise use a CDN, uses TARMK. From [Adobe's microkernel recommendations](#),

- TarMK is designed for performance
- MongoMK is used for scalability of Authoring servers

"Adobe highly recommends TarMK to be the default persistence technology used by customers in all deployment scenarios, for both the AEM Author and Publish instances".

MongoDB is an alternative data storage [introduced](#) in AEM 6.0 which provides active site redundancy. However, Adobe does not yet endorse full fledged as of 6.0 for active site redundancy (ie sharding support was problematic and dropped). Adobe's only [reason to recommend MongoDB](#) is if author instance has load requiring one instance. Even then, a larger Author machine likely the better solution.

SDD or HDD? - Ask Adobe how often the microkernel writes?

Adobe Acryonyms

AEM was previously called CQ.

CRX - references the bundle of technologies put together by Adobe.

Version Information

The BonsaFramework instructions were made for,

- AEM 6.2
- Apache 2.4.x
- Ubuntu Server LTS 64-bit (exact version should not matter)

Download

June 2017 - Downloads do not seem to work with Chrome. Have had success with Safari.

No public downloads as far as I can tell. Log into [Licensing Website](#). Click LICENSES and Downloads. Click Adobe Experience Mgr.

You will see two downloads,

Download	Size	Comment
Adobe Experience Manager 6.2 Quickstart Setup	476MB	The version you want for Bonsai Framework instructions.
Adobe Experience Manager 6.2 Quickstart Deployment for Application Server		It downloads in some browsers as .zip but should be renamed to .war. It is for deploying to Application Servers like Tomcat or WebSphere. * I'll speak to Adobe about this because the docs are not clear at all.

Click LICENSES. Do not select your product. Just click the **Search** button.

You should see "Adobe Experience Mgr" but it says "No Serial Number Required". This is not true. Instead look for "AEM Standard" with Production Version "5.6" and you will see a serial number under the "Serial" column. Copy and paste into a standard text file for later.

AEM is built using many of the same concepts endorsed in the Bonsai Framework, hence the entire system in initial download is a single file in two flavours,

- **cq-quickstart-6.2.0.jar** - executable to run as a Java process which includes Jetty as an embedded container
- **cq-quickstart-6.2.0.war** - install in a container like Tomcat

Next Steps

In a large organization you will also need to consider the following,

- [Brochure and Onboarding](#) - Explain what to expect and limitation enforced by your organization to be enterprise.
- Security - Restrictions to enforce
- Operations Runbook - DR procedures
- Author Development - ...
- Author Workflow - ...
- Author - ...
- Contributor - ...
- Subscriber - ...
- Get ready for Adobe 6.3 https://solutionpartners.adobe.com/home/news/2017/03/get_ready_for_adobe_.html

References

[AEM Good Bad and Ugly](#) - Pretty good technology feature with description dump.

[Amazon Offering of AEM as a Service](#) - (to Read) Possible ideas of how we setup explain our service. p15, interesting points on dynamic content ESI, page 15 mod_rewrite in Apache prevent request passing to Dispatcher module.