

Hardware and Network Considerations



Production Environments

- 32Gbs minimum of RAM recommended
- Minimum free on hard drive depending of the requirements of the project

Development and Testing Environments

- 8Gbs minimum of RAM recommended
- 10Gbs minimum free on hard drive recommended

Hardware

Basic Aspire has a fairly small footprint. For evaluation purposes, it can run on your laptop. For most production implementations, we recommend:

- A 4-core server running 64-bit Linux or Windows
- Minimum of 8Gbs of RAM (**32Gbs recommended for production**) assigned to the Java Virtual Machine.
 - The memory can be modified at the startup files or check [Run as a Windows Service](#) or [Run as a Linux Service](#) if Aspire runs as a service.
 - 8GB is the minimum recommended for testing and lite Aspire instances. This may change depending of the content sources and the data to process.
- Minimal disk: Aspire itself currently runs in less than 1GB, we recommend a minimum of **10Gbs for development and testing environments**
 - We don't recommend the minimal disk space for production environments
 - *This value does not include logs or data generated due crawls or other component activity*
- Network connection speed of 100Mbps minimum, and preferred 1Gbps or 10Gbps

Factors Affecting Recommendations

What you need to run Aspire optimally in a complex, enterprise production environment depends on multiple factors, including:

- The total number of data sources to be crawled, and for each data source:
 - What type of repository holds the data (i.e. SharePoint, Documentum, RDBMS, etc.)
 - Total number of documents
 - Total number of bytes
 - How long an initial indexing crawl should take
 - How long a periodic incremental indexing crawl should take
 - Are real-time updates required?
 - Network speed between repository and Aspire server
- Data metrics for query client connections:
 - Number of client connections expected
 - Number of queries per second expected
 - Query latency times required / expected
 - Is authentication, authorization, and group expansion required for client search requests?
- Will distributed processing be required?

If you need assistance in sizing an Aspire Enterprise system, please contact your Search Technologies Account representative (we'd be happy to help).

Network

Aspire makes substantial use of [Maven](#), in both the build and execution phase. During the build process, Maven will download bundles (jar files) and plugins from several Maven repositories located on the internet (including, but not limited to <https://repository.searchtechnologies.com/artifactory/public> and <http://repo1.maven.org/maven2>). During execution, Aspire will download applications and components (OSGI bundles or jar files) from <https://repository.searchtechnologies.com/artifactory/public>.

If allowing Aspire access to the internet is not in line with your network policy, see [Setting Aspire with No Internet Access \(Offline\)](#) or [Using your own Maven repository with Aspire](#).

Another network consideration is the speed of the network between the content repository and the Aspire server. When an initial crawl or incremental crawl is performed it is typically fetching the entire document and associated metadata. Although the text of each document may be rather small, the full binary can be quite large and if the actual network speed, particularly for WANs, is slow, the overall performance of the system will be slow. There are solution architectures that can be employed to address this situation. Please contact Search Technologies and we would be happy to discuss these issues and potential solutions.