

Synapse Sizing Guidance

Synapse uses a microservice architecture. Inasmuch, the deployment is typically divided among multiple hosts, with the individual components communicating over TLS. Vertex has recommendations for provisioning infrastructure based on the number of expected Cortex nodes, Synapse UI users, JSONStor metadata storage, and expected Axon storage. We break these into Small, Medium, and Large deployments. These deployment recommendations are meant as a starting point to estimate infrastructure costs, but are not comprehensive. We would be happy to discuss your use case in greater detail and help to tailor a deployment plan to your specific needs.

Size	Number of Nodes	Number of Users	Metadata Storage (JSONStor)	Number of Blobs (Axon)
<i>Small</i>	10 Million	10 Users	16 GB	100,000 blobs
<i>Medium</i>	100 Million	25 Users	64 GB	250,000 blobs
<i>Large</i>	1 Billion	100 Users	256 GB	500,000 blobs

Cortex

The Cortex is where the knowledge graph and query engine lives. Sizing a host for a Cortex is similar to sizing a database server. A Cortex performs best when it has enough RAM to hold much of the graph information in memory, but this is not a strict requirement. We recommend that the backing storage for the Cortex is done on Solid State Drives (SSDs) as opposed to spinning disks (HDDs). Using an NFS-based network attached storage (e.g. AWS EFS or Azure Files) should **not** be used to run a Cortex. Deploying Cortex mirrors typically requires the same level of hardware support for each mirror being deployed.

Axon

The Axon is used for object storage, e.g. for file samples. This can either store objects directly or via an S3 compatible storage interface (e.g. AWS S3, Digital Ocean Spaces, MinIO). The later examples shown assume an average object size of 1MB each using AWS S3 storage.

Synapse UI

The Synapse UI is the host that serves the web-based user interface. The Optic host requires its own storage for user data and server-side caching.

JSONStor

The Synapse JSONStor is a centralized metadata store, which is used to store runtime metadata for the Cortex, the User Interface, and other Synapse services. This service can be co-hosted (i.e. run on the same host) with the User Interface or with Advanced Power-Ups.

Advanced Power-Ups

Some of the Synapse Advanced Power-Ups may be deployed separately from the Cortex and the Synapse UI and will require their own resources in order to operate. Many of these can be co-hosted and some will work best when they have their own dedicated storage available.

The following Power-Ups are appropriate to be co-hosted without any issues:

- Synapse-Fileparser
- Synapse-Nettools
- Synapse-Maxmind
- Synapse-Playwright
- Synapse-Sidepocket

Two other Power-Ups, Synapse-Metrics and Synapse-Search, perform specialized indexing, hence have storage and CPU requirements similar to the Cortex.

Example Infrastructure and Estimated Costs (AWS)

The following estimated yearly costs for deploying Synapse in AWS include provisioning the following example infrastructure:

Size	Cortex	UI	Axon S3 (Storage)	Misc Services	Total Yearly	Metrics & Search	Total Yearly with Metrics & Search
<i>Dev</i>	\$348	\$0	\$60	\$0	\$408	\$0	\$408
<i>Small</i>	\$696	\$0	\$60	\$847	\$1,603	\$1,392	\$2,995
<i>Medium</i>	\$2,580	\$564	\$240	\$1,637	\$5,021	\$5,160	\$10,181
<i>Large</i>	\$5,928	\$1,128	\$720	\$1,750	\$9,526	\$11,856	\$21,382

Dev

- Developers may want their own hosted infrastructure to test on in a safe environment.
- A single t3.large instance (8GB Ram, 2 vCPU) with 30GB of Gp3 storage, running a Cortex, co-hosting the Synapse UI and any services as needed.

Small

- Cortex / Metrics / Search can be hosted on separate r6i.large (16GB; 2 vCPU) instances with 200GB gp3 storage.
- Co-hosted services, including Optic, can be hosted on a m6i.xlarge (16GB, 4 vCPU) with 64GB storage.
- The JSONStor should be allocated at least 16GB of storage.
- This assumes ~200GB of AWS S3 storage.

Medium

- Cortex / Metrics / Search can be hosted on separate r6i.2xlarge (64GB, 8 vCPU) instances with 800GB of gp3 storage.¹
- Optic can be hosted on a r6i.large (16GB, 2 vCPU) with 64GB of gp3 storage.
- Misc Services m6i.2xlarge (32GB, 8 vCPU) 32GB of gp3 storage.
- The JSONStor should be allocated at least 64GB of storage.
- This assumes ~800GB of AWS S3 storage.

¹ The RAM footprint for the services is more important than the vCPU count.

Large

- Cortex / Metrics / Search can be hosted on separate r6i.4xlarge (128GB, 8 vCPU) instances with 2TB of gp3 storage.
- Optic can be hosted on a r6i.xlarge (32GB, 4 vCPU) instance with 128GB of gp3 storage.
- Misc Services m6i.2xlarge (32GB, 8 vCPU) with 32GB of gp3 storage.
- The JSONStor should be allocated at least 128GB of storage.
- This assumes ~2.4TB of AWS S3 storage.

Example Infrastructure and Estimated Costs (Azure)

The following estimated yearly costs for deploying Synapse in Microsoft Azure include provisioning the following example infrastructure:

Size	Cortex	UI	Axon Azure (Storage)	Misc Services	Total Yearly	Metrics & Search	Total Yearly with Metrics & Search
Dev	\$330	\$0	\$55	\$0	\$385	\$0	\$385
Small	\$863	\$0	\$55	\$908	\$1,826	\$1,726	\$3,552
Medium	\$3,453	\$602	\$220	\$1,672	\$5,947	\$6,906	\$12,853
Large	\$7,674	\$2,161	\$665	\$1,787	\$12,287	\$11,509	\$23,796

Dev

- Developers may want their own hosted infrastructure to test on in a safe environment.
- A single B2ms instance (8GB Ram, 2 vCPU) with 30GB of Premium storage, running a Cortex, co-hosting the Synapse UI and any services as needed.

Small

- Cortex / Metrics / Search can be hosted on separate E2ds v5 (16GB; 2 vCPU) instances with 200GB Premium storage.
- Co-hosted services, including Optic, can be hosted on a D4ds v4 (16GB, 4 vCPU) with 64GB Premium storage.
- The JSONStor should be allocated at least 16GB of Premium storage.
- This assumes ~200GB of Azure Blob storage in a Hot availability tier with Zone Redundant Storage.

Medium

- Cortex / Metrics / Search can be hosted on separate E8ds v5 (64GB, 8 vCPU) instances with 800GB of Premium storage.²
- Optic can be hosted on a E2ds v5 (16GB, 2 vCPU) with 64GB of Premium storage.
- Misc Services can be hosted on a D8ds v4 (32GB, 8 vCPU) 32GB of Premium storage.
- The JSONStor should be allocated at least 64GB of Premium storage.
- This assumes ~800GB of Azure Blob storage.

² The RAM footprint for the services is more important than the vCPU count.

Large

- Cortex / Metrics / Search can be hosted on separate E16ds v5 (128GB, 16 vCPU) instances with 2TB of Premium storage.
- Optic can be hosted on a E4ds v5 (32GB, 4 vCPU) instance with 128GB of Premium storage.
- Misc Services can be hosted on a D8ds v4 (32GB, 8 vCPU) with 32GB of Premium storage.
- The JSONStor should be allocated at least 128GB of Premium storage.
- This assumes ~2.4TB of Azure Blob storage.