

This FAQ provides answers to common questions about **Reveal Private Deployment (RPD)**, including architecture options, infrastructure and storage requirements, security considerations, integrations, and operational best practices to help customers plan, deploy, and manage Reveal in their own environment.

Deployment Architecture & Platform Requirements

Q: Are these the only deployment options available?

No, not all. The architectures described in this document represents minimum guidelines for likely deployment scenarios. Discussions with Reveal technicians for what-if scenarios are strongly encouraged as every organization's eDiscovery needs will vary greatly based on document complexity, number of graphics versus text in various documents, and the desired workflow within your eDiscovery operations.

Q: Can Reveal be deployed on virtual or physical systems?

Yes. However, as eDiscovery needs rarely involve static usage, deployment onto Virtual Infrastructure is strongly recommended both for short-term spikes, as well as ongoing growth of usage.

Q: The documented Reveal solution appears to be heavily reliant upon Microsoft products; is this a requirement?

Yes. While some components (such as Elasticsearch and the SMB Network Share) can be deployed on other Operating Systems, Reveal relies on core Microsoft technologies including IIS, MSSQL and .NET/C#.

Q: The documented Brainspace solution appears to be heavily reliant upon Linux operating systems; is this a requirement?

Yes. The Brainspace solution is deployed solely on Linux systems.

Q: Does Reveal require Microsoft SQL Server Enterprise for any special features?

No, Reveal does not take advantage of anything specific in Enterprise except for the additional RAM support that Enterprise offers for scaling.

Storage and Data Management

Q: What is the recommended storage setup?

- S3 (S3 Compatible storage)
 - revealdata-s3stage- Load file staging and temporary storage.
 - revealdata-s3store- Permanent project data storage.
- NFS: One share; subdirectories managed by the application/MSA.
- SMB: One share; application creates subdirectories as needed.
- Local storage/temp is also required on the node for temporary operations.

Best practice is to utilize **Full Recovery model** for MSSQL. Storage allocations may be trimmed slightly if using the **Simple Recovery model**.

Q: Can SAN or NAS storage technologies be utilized in the solution?

Yes, these are very common. Just keep in mind that during high-load scenarios (Importing, Indexing, Production, etc.), the speed of storage will be a critical factor in the duration of those activities.

Q: Both the Processing Agent and the Elasticsearch components stress the use of high-speed storage. Is this required?

No, but it is strongly recommended. Both of these components are very I/O intensive, so having local, high-speed storage can greatly reduce the duration of many activities.

Processing, Agents, and Client Components

Q: Why are so many Reveal Processing Agents recommended?

Currently, the Reveal Processing Agent software can perform various tasks, but not at the same time. Multiple Processing Agents help ensure that no activities will be queued waiting for another activity to complete.

Q: Is Microsoft Outlook required on each Reveal Processing Agent?

Yes, as of now, this is a technical requirement.

Q: Is Microsoft Office required on each Reveal Processing Discovery Manager system?

No, this is not a technical requirement, but many users find that having MSOffice installed on these systems can assist with QC and troubleshooting activities.

Q: Do the Discovery Manager and Review Manager thick clients have to reside on their own systems?

No. Typically, customers will deploy Discovery Manager on its own System, with Review Manager either being deployed on the same system as Discovery Manager or one of the Review Application Servers.

Q: Do all components involved with the Reveal solution have to remain online at all times?

No, in many scenarios, idle systems can be taken offline when not actively in use. An example of this would include Processing Agents when no Processing is actively being performed.

Searching, Indexing, and Integrations

Q: Does Reveal use DT Search or Elasticsearch?

Reveal uses Elasticsearch.

Note

Discovery Manager supports DT Search indexing for keyword searching during processing. This index is only used within Discovery Manager and is not used in Review. All documents loaded into Review are reindexed into Elasticsearch.

Q: Is Elasticsearch only used for Review searching?

No, the Reveal Review product stores the actual text from documents in Elasticsearch. Elasticsearch needs to be online for almost all Reviewing activities.

Q: What integrations exist between the Reveal products?

Reveal Processing can directly export to Reveal Review although the "Uploader" feature in the Reveal Review web front-end is only available currently in the Reveal SaaS offering.

Backup, Recovery, and Operations

Q: How should backups be configured for Reveal Solutions?

Best Practices vary by component.

- SQL: Full Recovery mode with native SQL backups
- Elasticsearch: Built-in backup restore using Curator
- Brainspace, file servers, VMs: Virtualization-aware backup solutions

AI, GPU, and AWS Requirements

Q: What GPU requirements exist for AJI and ASK features?

To enable AJI and ASK, the platform needs to connect to a Consilio instance of AWS Bedrock. The GPU hardware can be comparable to an NVIDIA T4 GPU.

Q: What AWS Services will access be needed to from the AI node?

The following services with S3 compatible storage will be needed compatible

- AWS Bedrock for AJI/ASK
- AWS Transcribe
- AWS Translate
- AWS Rekognition

Q: How many AWS S3 buckets will be needed, and what will be their purpose?

Only one bucket is required, with multiple prefixes used in a single bucket, such as for sending to Transcribe & Translate.

Q: How many IAM users will be needed?

There will only be one IAM user required. In Reveal's SaaS environment this is one per MSA (Customer) with different roles created that AWS will assume as needed for specific work (Translate and transcribe).

Q: Will all data stored in AWS S3 be encrypted at rest?

All data stored in AWS S3 buckets are encrypted using AWS Key Management Service (KMS).

Scaling, Performance, and Throughput

Q: Some of the resource requirements are significant; can we start smaller and scale up later?

Yes, but keep in mind that adjustments in allocated resources may not be deemed necessary until in-progress activities are underway, often with a pending deadline, and adjustments to some resources (typically CPU and Memory) may require an outage which can be difficult to schedule if large batch jobs are executing. The

minimum resource allocations described in this document target eventual usage scenarios, but feel free to start smaller, and adjust upward, when necessary, always keeping a close eye on resource consumption and live throughput.

Q: Some scenarios list throughput estimates, are those guaranteed results?

Unfortunately, no. While those estimates are derived based on real-world usage of the Reveal-Brainspace solutions, numerous factors exist (document complexity, storage speed, etc.) which will impact overall throughput. Also, timeframe expectations vary between organizations. In most scenarios, throughput can be increased with either more robust systems or greater quantities of some components. Please feel free to discuss possible areas to improve throughput for certain activities with Reveal-Brainspace technicians during initial planning or deployment as well as over the long-run.

Networking, Security, and Access

Q: How many external IP addresses are necessary if we need external reviewers?

If using a proxy or load balancer one IP would be necessary. However, Additional IP addresses may be necessary for multiple web servers or publishing multiple services such as Review, Brainspace, and Keycloak authentication services.

Q: How will internal certificates be handled?

Internal certificates are provisioned through Cert-Manager, which interfaces with the Active Directory Certificate Services (ADCS) Issuer. The ADCS Issuer, in turn, communicates with the enterprise ADCS Certificate Authority to request and obtain the certificates.

Q: Does Reveal support Managed Service Accounts (MSA/gMSA)?

Yes, for Kubernetes you can use MSA as per Kubernetes documentation. Internal to the cluster, traditional accounts are required.