

### Physical and logical components

### Scale-out decision points

#### Overview

Search in Microsoft® SharePoint® Server 2010 is re-architected with new components to create greater redundancy within a single farm and to allow scalability in multiple directions.

Each of the components that make up the query architecture and the crawling architecture can be scaled out separately based on the needs of an organization.

**Query architecture**  
The query architecture includes query components, index partitions, and property databases.

- About index partitions:
  - An index partition is a logical portion of the entire index. The index is the aggregation of all index partitions.
  - Index partitions are associated with query components. You deploy a query component that is associated with a particular index partition to a specific server. In this way, index partitions are spread across query servers. For example, in a farm with three index partitions and one query component per partition, each query component contains one-third of the total index.
  - Deploying query components that are associated with index partitions across different servers creates a faster query architecture because the processing power of multiple query servers is used to respond to queries.
  - Index partitions can be associated with one or more query components. Multiple query components (mirrors) for a given index partition can be deployed across query servers to achieve redundancy. Typically, two query components are configured for each index partition, and these query components reside on different query servers to achieve redundancy of the index partition.

- Crawl architecture**  
The crawl architecture includes several components that can be scaled out based on crawl volume and performance requirements:
  - Crawl component** — Multiple crawl components can be deployed to crawl simultaneously. Each crawl component is associated with a crawl database. Crawl components reside on application servers. Crawl components produce portions of the index (per index partition) and propagate them to the servers that are running the query components associated with the given index partition.
  - Crawl database** — Manages crawl operations and stores crawl history. You can assign multiple crawl components to each crawl database for redundancy. In this case, each crawl component will crawl different content during a crawl.
  - Property database** — Also considered part of the query architecture, stores properties for crawled data. The number of required property databases depends on the volume of content that is crawled and the amount of metadata that is associated with the content.

- Architecting search for enterprise environments**  
SharePoint Server 2010 includes an updated services architecture that enables farms to connect to multiple farms to consume cross-farm services. In large environments, consider deploying one or both of the following options:
  - Enterprise services farm** — A farm that hosts the most commonly used cross-farm services, including search.
  - Dedicated search farm** — A farm that is optimized for search.

#### Server roles

- Web server**
  - Hosts Search Web Parts and Web Part pages for answering search queries.
  - In dedicated search service farms, this role is not necessary because Web servers at remote farms contact query servers directly.
  - This role is necessary for farms that include other SharePoint Server 2010 capabilities.
  - In small farms, this role can be shared on a server with the application server role.

#### Application server with query components

- Serves query results to Web servers.
- Holds the entire index if one index partition is configured. Otherwise, it holds the portions of the index that are associated with the index partitions as configured by the administrator.
- Runs the query processor (Search Query and Site Settings Service). The query processor forwards query requests to one query component for each index partition and then merges results to display to users. The query processor is also responsible for security trimming, duplicate results detection, and gathering and displaying properties. Multiple query processors can be configured (one per application server) by starting the Search Query and Site Settings Service on the Services on Server page in Central Administration.
- At least one server in a farm must host the query components.
- Two or more query servers provide redundancy based on the configuration of index partitions and query components.
- Adding query servers may increase performance.
- These components can be shared on a server with the Web server role or with other application server roles.

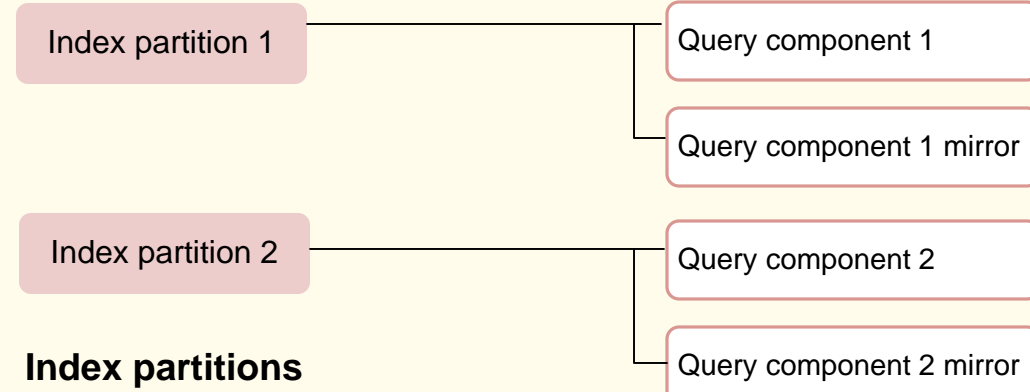
#### Application server with crawl components

- Runs one or more crawler components.
- Typically hosts the Search Administration component.
- At least one server in a farm must host the crawl server role.
- Two or more crawl servers provide redundancy based on how crawlers are associated to the crawl databases.
- Add more crawl servers to increase performance and to scale for capacity.
- This role can be shared on a server with the Web server role or with other application server roles.

#### Database server

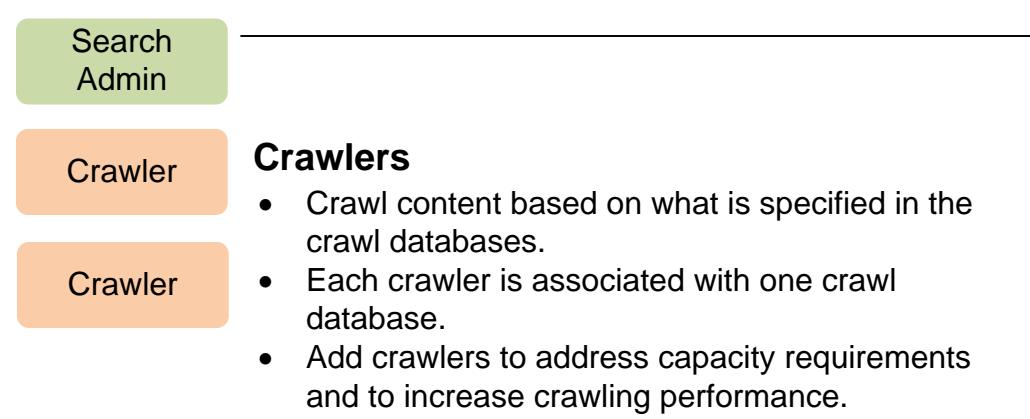
- Hosts search-related databases.
- Can host other SharePoint Server 2010 databases.
- Can be mirrored or clustered.
- To increase performance and capacity, consider adding disks to the database server or adding database servers (depending on the bottleneck).

#### Logical components



#### Index partitions

- The index can be subdivided into discrete partitions that can then be distributed across multiple query servers.
- The index is the aggregation of all index partitions.



#### Search Administration component

- Monitors incoming user actions and updates the search administration database.
- Only one Search Administration component per Search service application.
- Can run on any server, preferably either a crawl server or a query server.

- Property database**
  - Stores the metadata and security descriptors for the items in the index. It is involved in property-based queries and returns standard document attributes for all query results.

- Crawl database**
  - Stores the crawl history.
  - Manages crawl operations.
  - Each crawl database can have one or more crawlers associated with it.

#### Volume of items

How the number of items in your corpus affects farm sizing and the decision to scale out.

Number of items	Action
0-3 million	All Search roles can coexist on one or two servers, as shown in the Limited Deployment architecture example.
1-10 million	Move crawl components to another server, while the query components remain on the Web servers, as shown in the Small Farm architecture example.
10-20 million	Add a crawl server. Each crawl server has one crawler. Create another index partition with query components and distribute these across query servers, as shown in the Medium Shared Farm architecture example.
20-40 million	Add index partitions with distributed query components. Add another crawl database, and then add a new associated crawler to each crawl server.
40-100 million	Isolate each topology layer into server groups in which each role is deployed to its own servers. Each server group can then be scaled out to meet specific requirements for the components in that role. Add Crawl and Query databases as shown in the architecture diagrams below.

#### Performance metrics

Key performance characteristics and how they relate to scaling of specific search components.

To improve this metric...	Take these actions
<b>Full crawl time and result freshness</b>	Add crawl servers, crawlers, and crawl databases. Each crawl database contains content from independent sources. Each crawl database can have several crawl components associated with it, and those crawl components can be distributed among many crawl servers. If you have several content sources, multiple crawl components and associated crawl databases allow you to crawl the content concurrently.
<b>Time required for results to be returned</b>	If query latency is caused by high peak query load, add query servers and index partitions. Each index partition can contain up to ~10 million items. You can also add a mirror for each query component for a given index partition. Place the mirror copy on a different server. Query throughput increases when you add index partition instances. If query latency is caused by database load, isolate the property database from crawl databases by moving it to a separate database server.
<b>Availability of query functionality</b>	Deploy redundant (failover) query components for the same index partition on different servers and use clustered or mirrored database servers to host crawl and property databases.
<b>Availability of content crawling and indexing functionality</b>	Use multiple crawl components on redundant crawl servers, and add crawl databases. Crawlers associated with a given crawl database can be distributed across crawl servers for availability and load distribution.

#### Redundancy and availability

Which search components are redundant and how the redundancy or failover mechanism works.

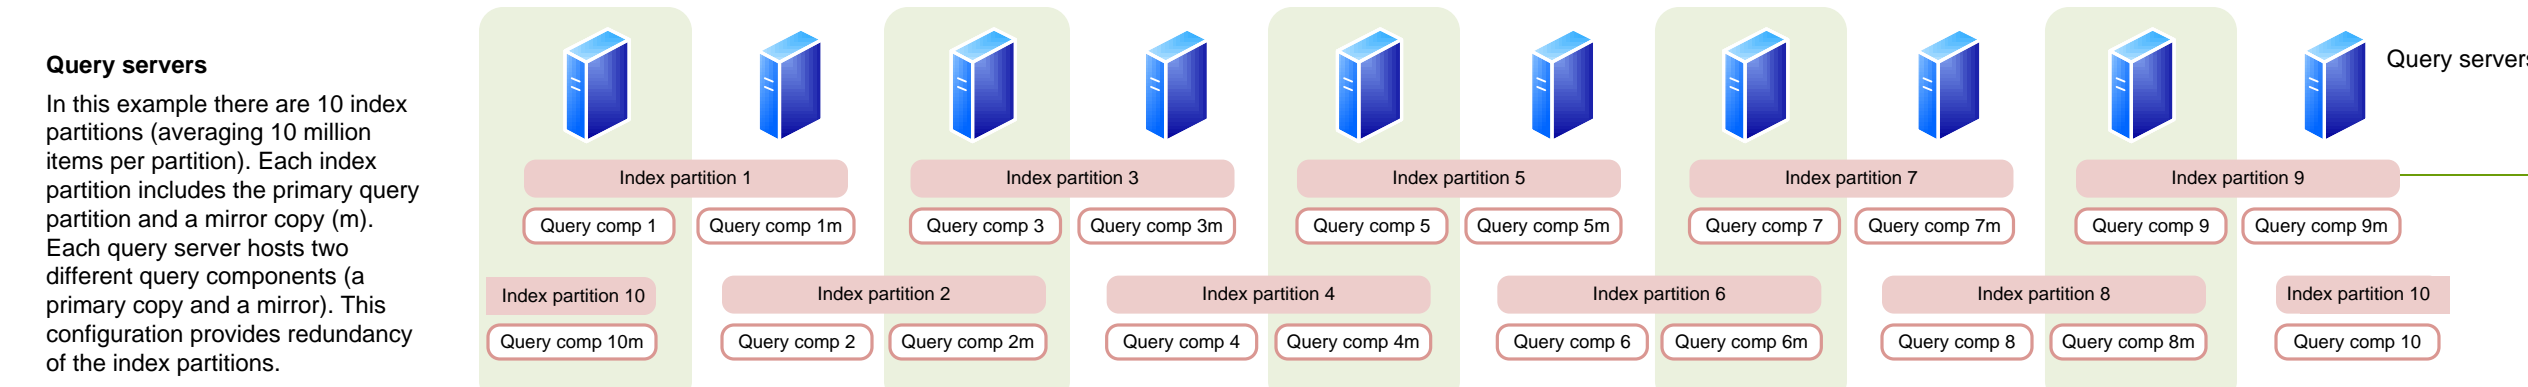
Component	Can be deployed to multiple servers?	How redundancy or failover works
Index partition	Yes	(See Query Components.)
Query components	Yes	You can create a failover copy of a query component (associated with the same index partition) on a different server for redundancy. Generally, one failover copy of a query component that is deployed to a different server is sufficient to fulfill redundancy requirements. You can also create an active copy of a query component associated with the same index partition on a different server for improving query throughput. Ensure that the different server also has a query processor (Search Query and Site Settings Service) enabled to avoid a bottleneck during query processing.
Search Administration component	No	There is only one Search Administration component required in a farm. A Search Administration component can be moved to a different server by using a Windows PowerShell™ cmdlet (Set-SPEnterpriseSearchAdministrationComponent).
Crawl component	Yes	You can have as many crawl components on a given crawl server as resources permit. You should plan for reserving four processor cores for each crawl component. Each crawl component on a given crawl server should be associated with a separate crawl database. For example, if you have two crawl databases due to the amount of content in your index, you should have at least two crawl servers with one crawl component from each database on each server. See the example diagram below for details.
Property database	Yes	The property database contains metadata for all crawled content. You may need more than one property database per 25 million items if there is a large amount of metadata associated with crawled content. Databases are redundant if they are mirrored or deployed to a clustered Microsoft SQL Server® computer.
Crawl database	Yes	The crawl database contains crawled content, and should be maintained on a separate storage array from the property database as a best practice to prevent I/O contention. If the crawl window overlaps with times when users are querying, or several crawl components are connected to a crawl database, consider deploying the crawl database to a separate database server. You can also have multiple crawl databases with different crawl components connected to them. Databases are redundant if they are mirrored or deployed to a clustered SQL Server computer.
Search Administration database	No	There is only one Search Administration database per Search service application. The Search Administration database is only redundant if it is mirrored or deployed to a clustered SQL Server computer.

### Example scaled-out architecture

#### Large dedicated search farm

Large dedicated search service topology (~100M items)

Web servers: Not necessary for a dedicated search service farm.



#### Crawl servers

- Administration component on one crawl server. Each crawl server has two crawlers:
  - Each crawler is associated with a crawl database.
  - Each crawl database has two different crawlers associated with it that are located on different crawl servers.

#### Database servers (clusters or mirrors)

- Two database server sets containing:
  - Two crawl databases, one on each database server (with x spindles).
- Two database server sets containing:
  - Search Administration database on one server.
  - Two property databases, one on each database server (with x spindles each).
- All other SharePoint databases.

#### Index partitions

**Load balancing**  
The index is distributed across query servers to load balance query processing and to increase query performance.

SharePoint administrators decide how many index partitions to create and how to distribute the query components across query servers.


**Redundancy**  
For redundancy, create a mirror copy of each query component and place the mirror copy on a different server.

The illustration shows how index partitions and query components (including mirrors) can be spread across servers. Each server hosts a primary copy (query comp n) and a mirror copy (query comp nm).

### Additional example architectures

#### Limited deployments

##### Limited deployments

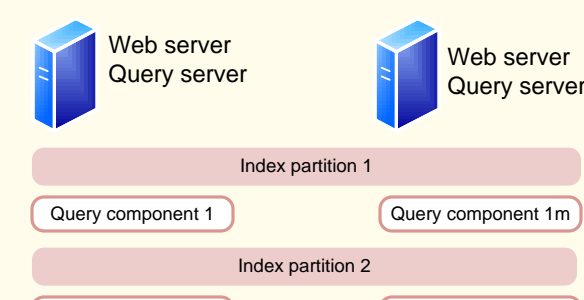


#### Medium shared architecture

Medium shared search topology (~20M items. At this scale point, consider using a dedicated search service farm.)


##### Combined Web and query servers

Place the Web server role on the same servers with the query components. There are two index partitions. For redundancy, each index partition includes a mirror copy (m) that is placed on the other server.



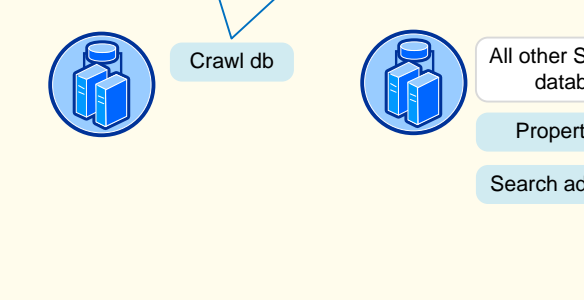
##### Crawl servers

Administration component is on one server. Each server has one crawler. Both crawlers are associated with the crawl database.



##### Database servers (clusters or mirrors)

- One database server set for:
  - Search Administration database
  - One property database
  - All other SharePoint databases
- One database server set containing the crawl database, which is isolated for performance reasons.



#### Medium dedicated search architecture

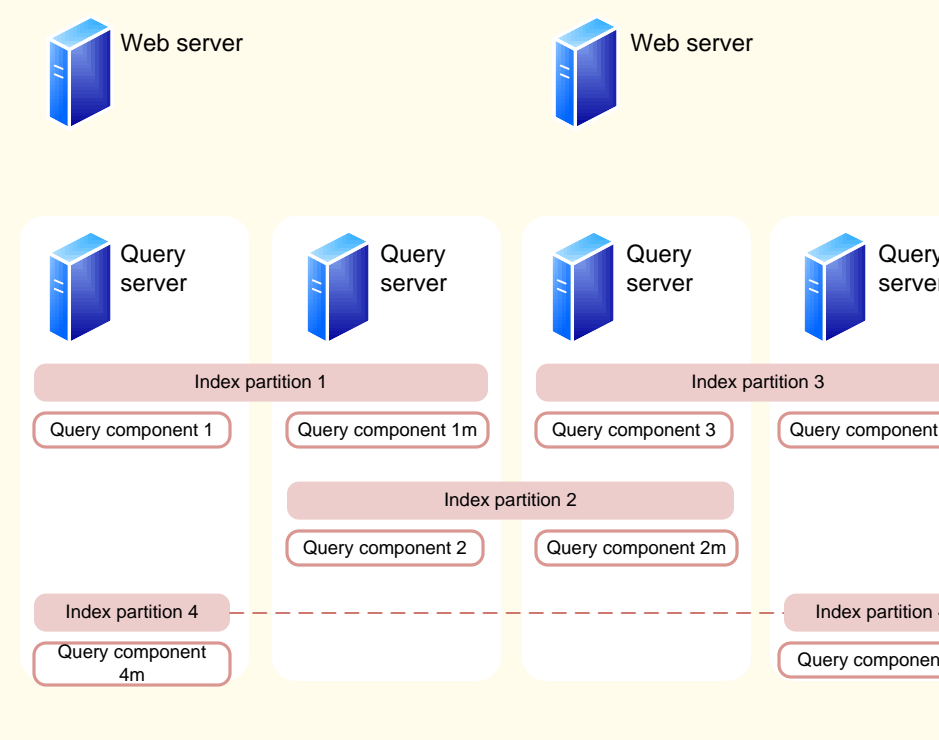
Medium dedicated search topology (~40M items)

##### Web servers

Web servers do not need to be hosted on dedicated search farms. However, farms that use the dedicated search farm must have Web servers.

##### Query servers


In this example, there are four index partitions. Each index partition includes two query components, a primary copy and a mirror (failover) copy (m). Each query server hosts two different query components. This configuration provides redundancy of the index partitions.



##### Crawl servers

Administration component on one crawl server. Each server has two crawlers:

- Each crawler is associated with a crawl database.
- Each crawl database has two different crawlers associated with it that are located on different crawl servers.



##### Database servers (clusters or mirrors)

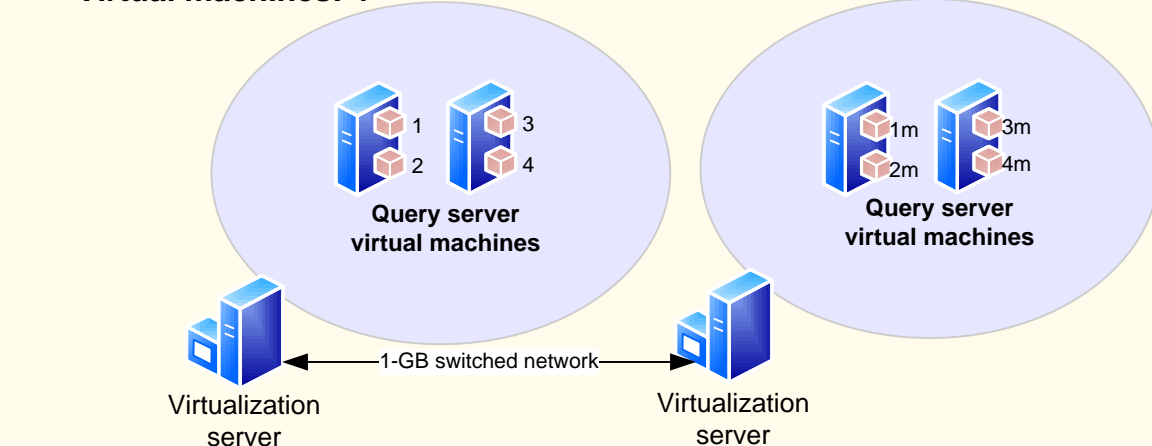
- Two database server sets containing:
  - Two crawl databases on one server.
  - Property database, Search Administration database, and all other SharePoint databases on one server.

#### Virtualized topologies

Virtualization is an effective way to reduce total cost of ownership by consolidating several servers on a single physical host. Using technologies such as Hyper-V™ for Windows Server 2008 and System Center Virtual Machine Manager, you can virtualize all or part of the SharePoint Server 2010 search environment. Virtualization provides benefits in the areas of hardware and server maintenance costs, data center space, manageability, ease of deployment, and rapid scalability. The virtualized topologies below show examples of how servers fulfilling different application roles can be grouped in a virtual environment.

##### Virtualization of query servers

Virtualization servers: 2  
Virtual machines: 4



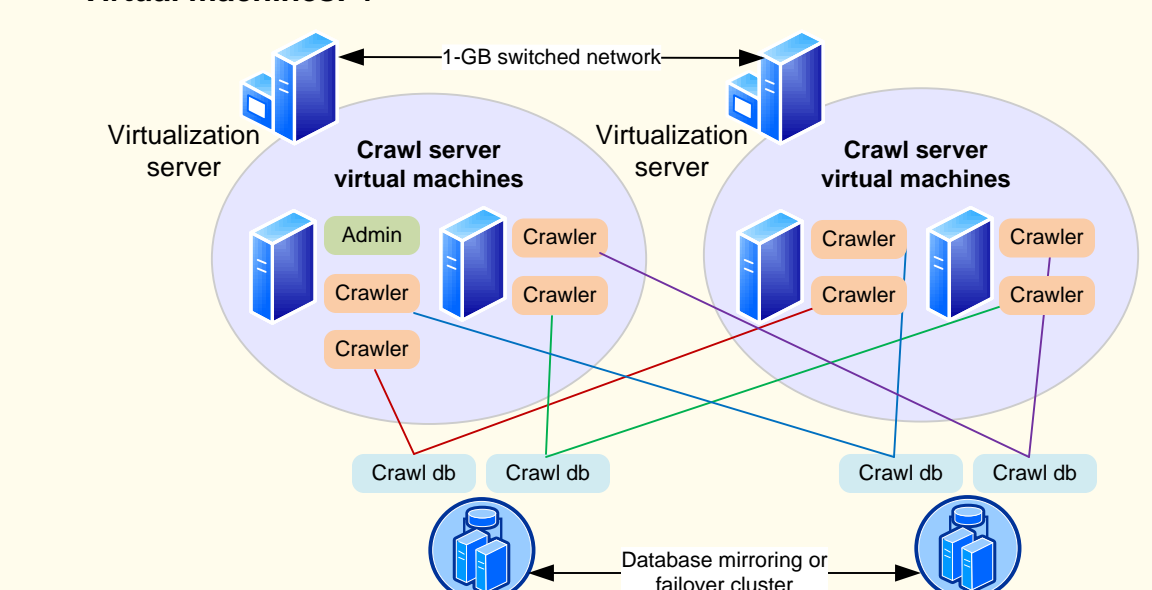
In this example, virtual query servers are running on two virtualization servers that are running Hyper-V. Redundancy is achieved by using two physical server computers that are running Hyper-V and by deploying redundant roles on virtual machines also.

In this example, there are four index partitions in total. Each virtual query server contains copies of two different query components or query component mirrors. The primary and mirror copies of each query component are located on different servers.

In this way, either virtualization server or any individual query server virtual machine could fail without loss of availability.

##### Virtualization of crawl servers

Virtualization servers: 2  
Virtual machines: 4



Crawl servers can be run in virtual machines to provide redundancy and scalability. In this example, four virtual crawl servers are running on two virtualization servers. Note that the number of virtual machines on a virtualization server is driven by resource availability on the virtualization server, so it is possible to host many more virtual machines per virtualization server.

The Search Administration component is located on one crawl server. Each crawl server has two crawlers, each crawler associated with a different crawl database to maximize load distribution.

Each of the four crawl databases, which are running on physical database servers, is associated with two different crawlers located on different crawl server virtual machines on different virtualization servers.