

PLANNING AND USING THE DISTRIBUTED CACHE SERVICE

ABOUT THE DISTRIBUTED CACHE SERVICE

The Distributed Cache service provides in-memory caching services to several features in SharePoint Server 2013. Some of the features that use the Distributed Cache service include:

- Newsfeeds
- Authentication
- OneNote client access
- Security Trimming
- Page load performance

These features use the Distributed Cache for quick data retrieval. The Distributed Cache service stores data in memory and does not have a dependency on databases in SharePoint Server 2013. However, some SharePoint features may store data in both the Distributed Cache and databases.

This model describes the planning process, installation steps, and configuration tasks to perform and to implement the Distributed Cache service in your SharePoint Server 2013 farm. You must do this to ensure that the Distributed Cache service is healthy and supports dependent services as required. Use this model together with other TechNet content to plan and implement your deployment of the Distributed Cache service.

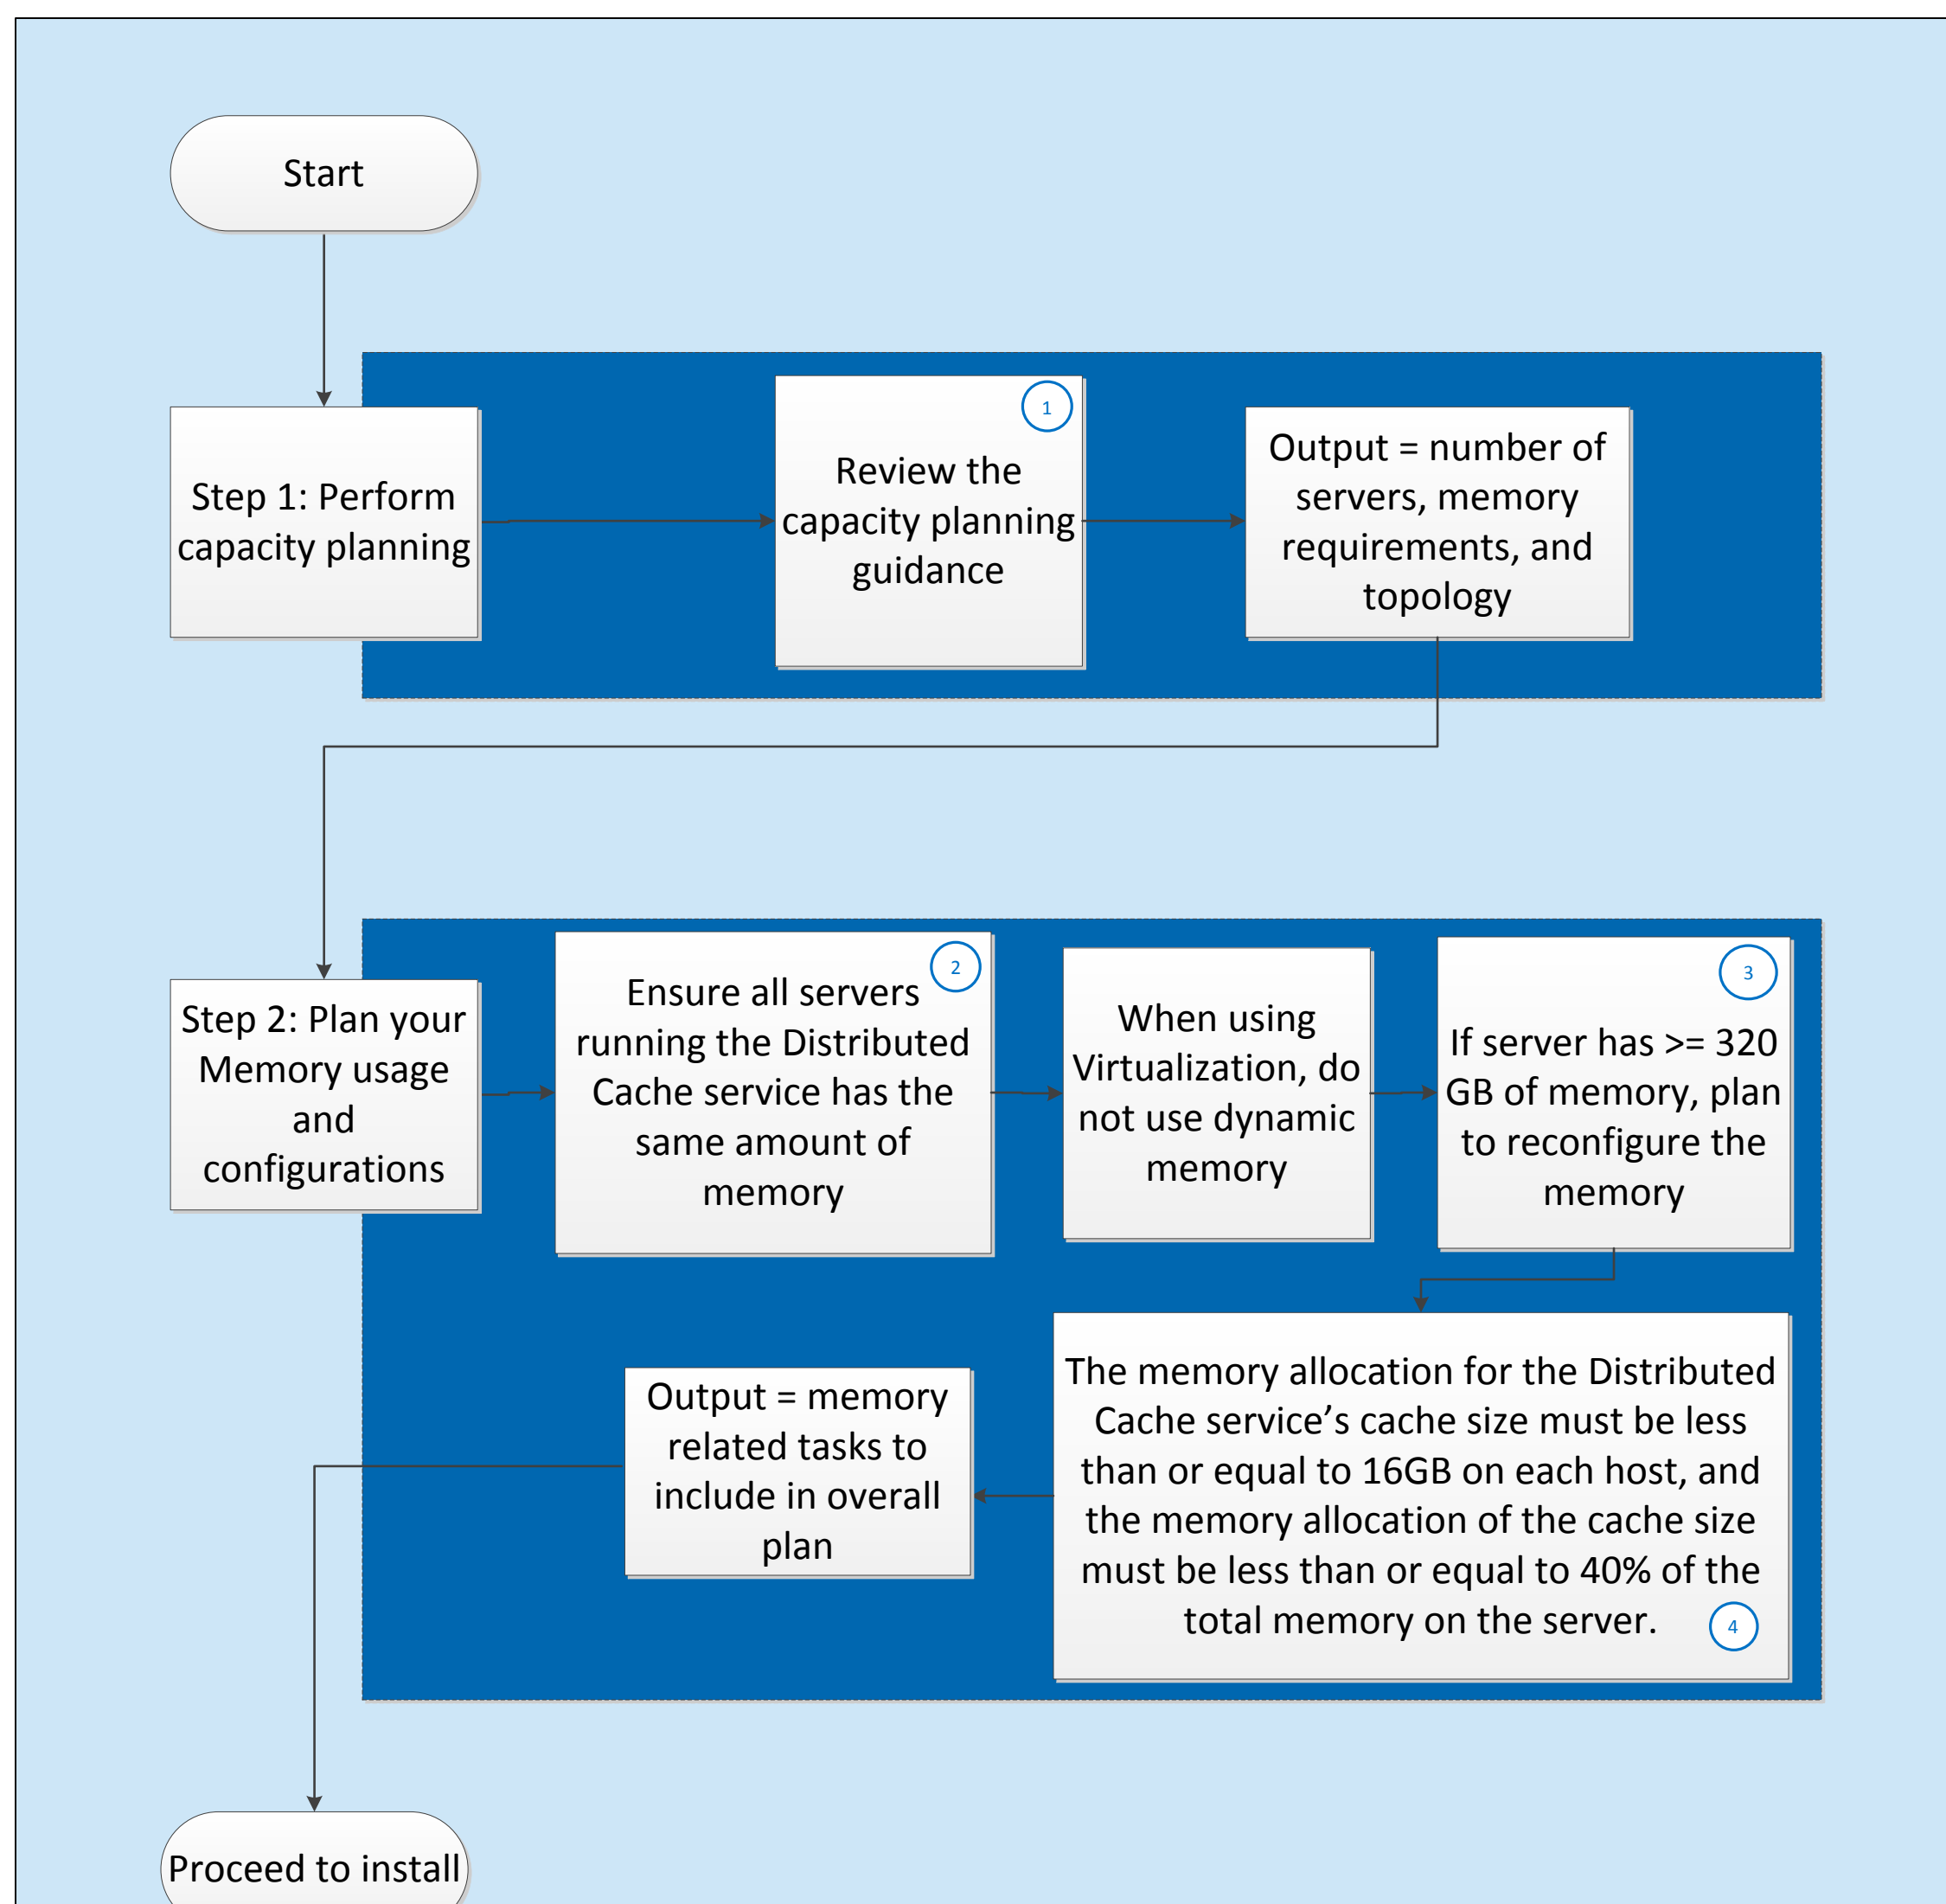
Other Caches which do not use the Distributed Cache service

- **Blob Cache.** The Blob Cache is a disk-based cache that caches files such as images, JavaScript files, and so on. This cache exists on each Web Front End server in the farm.
- **Output cache.** The Output Cache is an ASP.NET feature that provides caching functionality. The Output Cache is used to store frequently-accessed pages as a means of increasing the throughput of the system. On every Web Front End server in the farm, the Output Cache stores the rendered output of an .aspx web page.

PLANNING FOR THE DISTRIBUTED CACHE SERVICE

Planning for the Distributed Cache service helps ensure that the Distributed Cache service remains healthy, and reduces potential issues on dependent services. When planning for the Distributed Cache service, you should:

- Perform capacity planning. When performing capacity planning, you will consider your workload and amount of usage to develop the capacity recommendations in terms of number of servers, memory requirements, and topology.
- Use the capacity planning results to plan which memory configuration works best for your server farm.




Legend

- 1 For more information on Capacity Planning Guidance, see <http://go.microsoft.com/fwlink/?LinkID=269567>
- 2 All servers running the Distributed Cache service must match exactly in terms of total memory specifications on all servers in the cache cluster.
- 3 During installation, 5% of server memory is automatically allocated to the Distributed Cache service's cache size. The maximum amount of memory that can be allocated to the Distributed Cache service's cache size is 16GB per cache host in the server farm. When installing the Distributed Cache service on a server with 320GB of memory, the default memory assignment for the Distributed Cache service's cache size is automatically set to 16GB. For servers with > 320GB of memory, you must plan to reconfigure the memory allocation of the Distributed Cache service's cache size.
- 4 If an administrator allocates more than 16GB of memory to the Distributed Cache service's cache size, the server may unexpectedly stop responding for more than 10 seconds. The recommended minimum total physical memory on the server that runs the Distributed Cache service is 8GB (this is different from the cache size). If you are using a cache cluster with more than 1 cache host, ensure the memory allocation for the Distributed Cache service's cache size on each cache host is set to the same value.

ABOUT CACHE HOSTS AND CLUSTERS

A *cache host* is any server in the server farm that runs the Distributed Cache service. A *cache cluster* exists when one or more servers in the farm run the Distributed Cache service.

The total cache size for the server farm is the sum of each cache host's memory allocation for the cache size. The cache cluster's cache spans all cache hosts and saves data on each cache host. Data is not duplicated or copied on other cache hosts in the cache cluster. **A cache cluster cannot be configured for High Availability.**



Cache Host Management

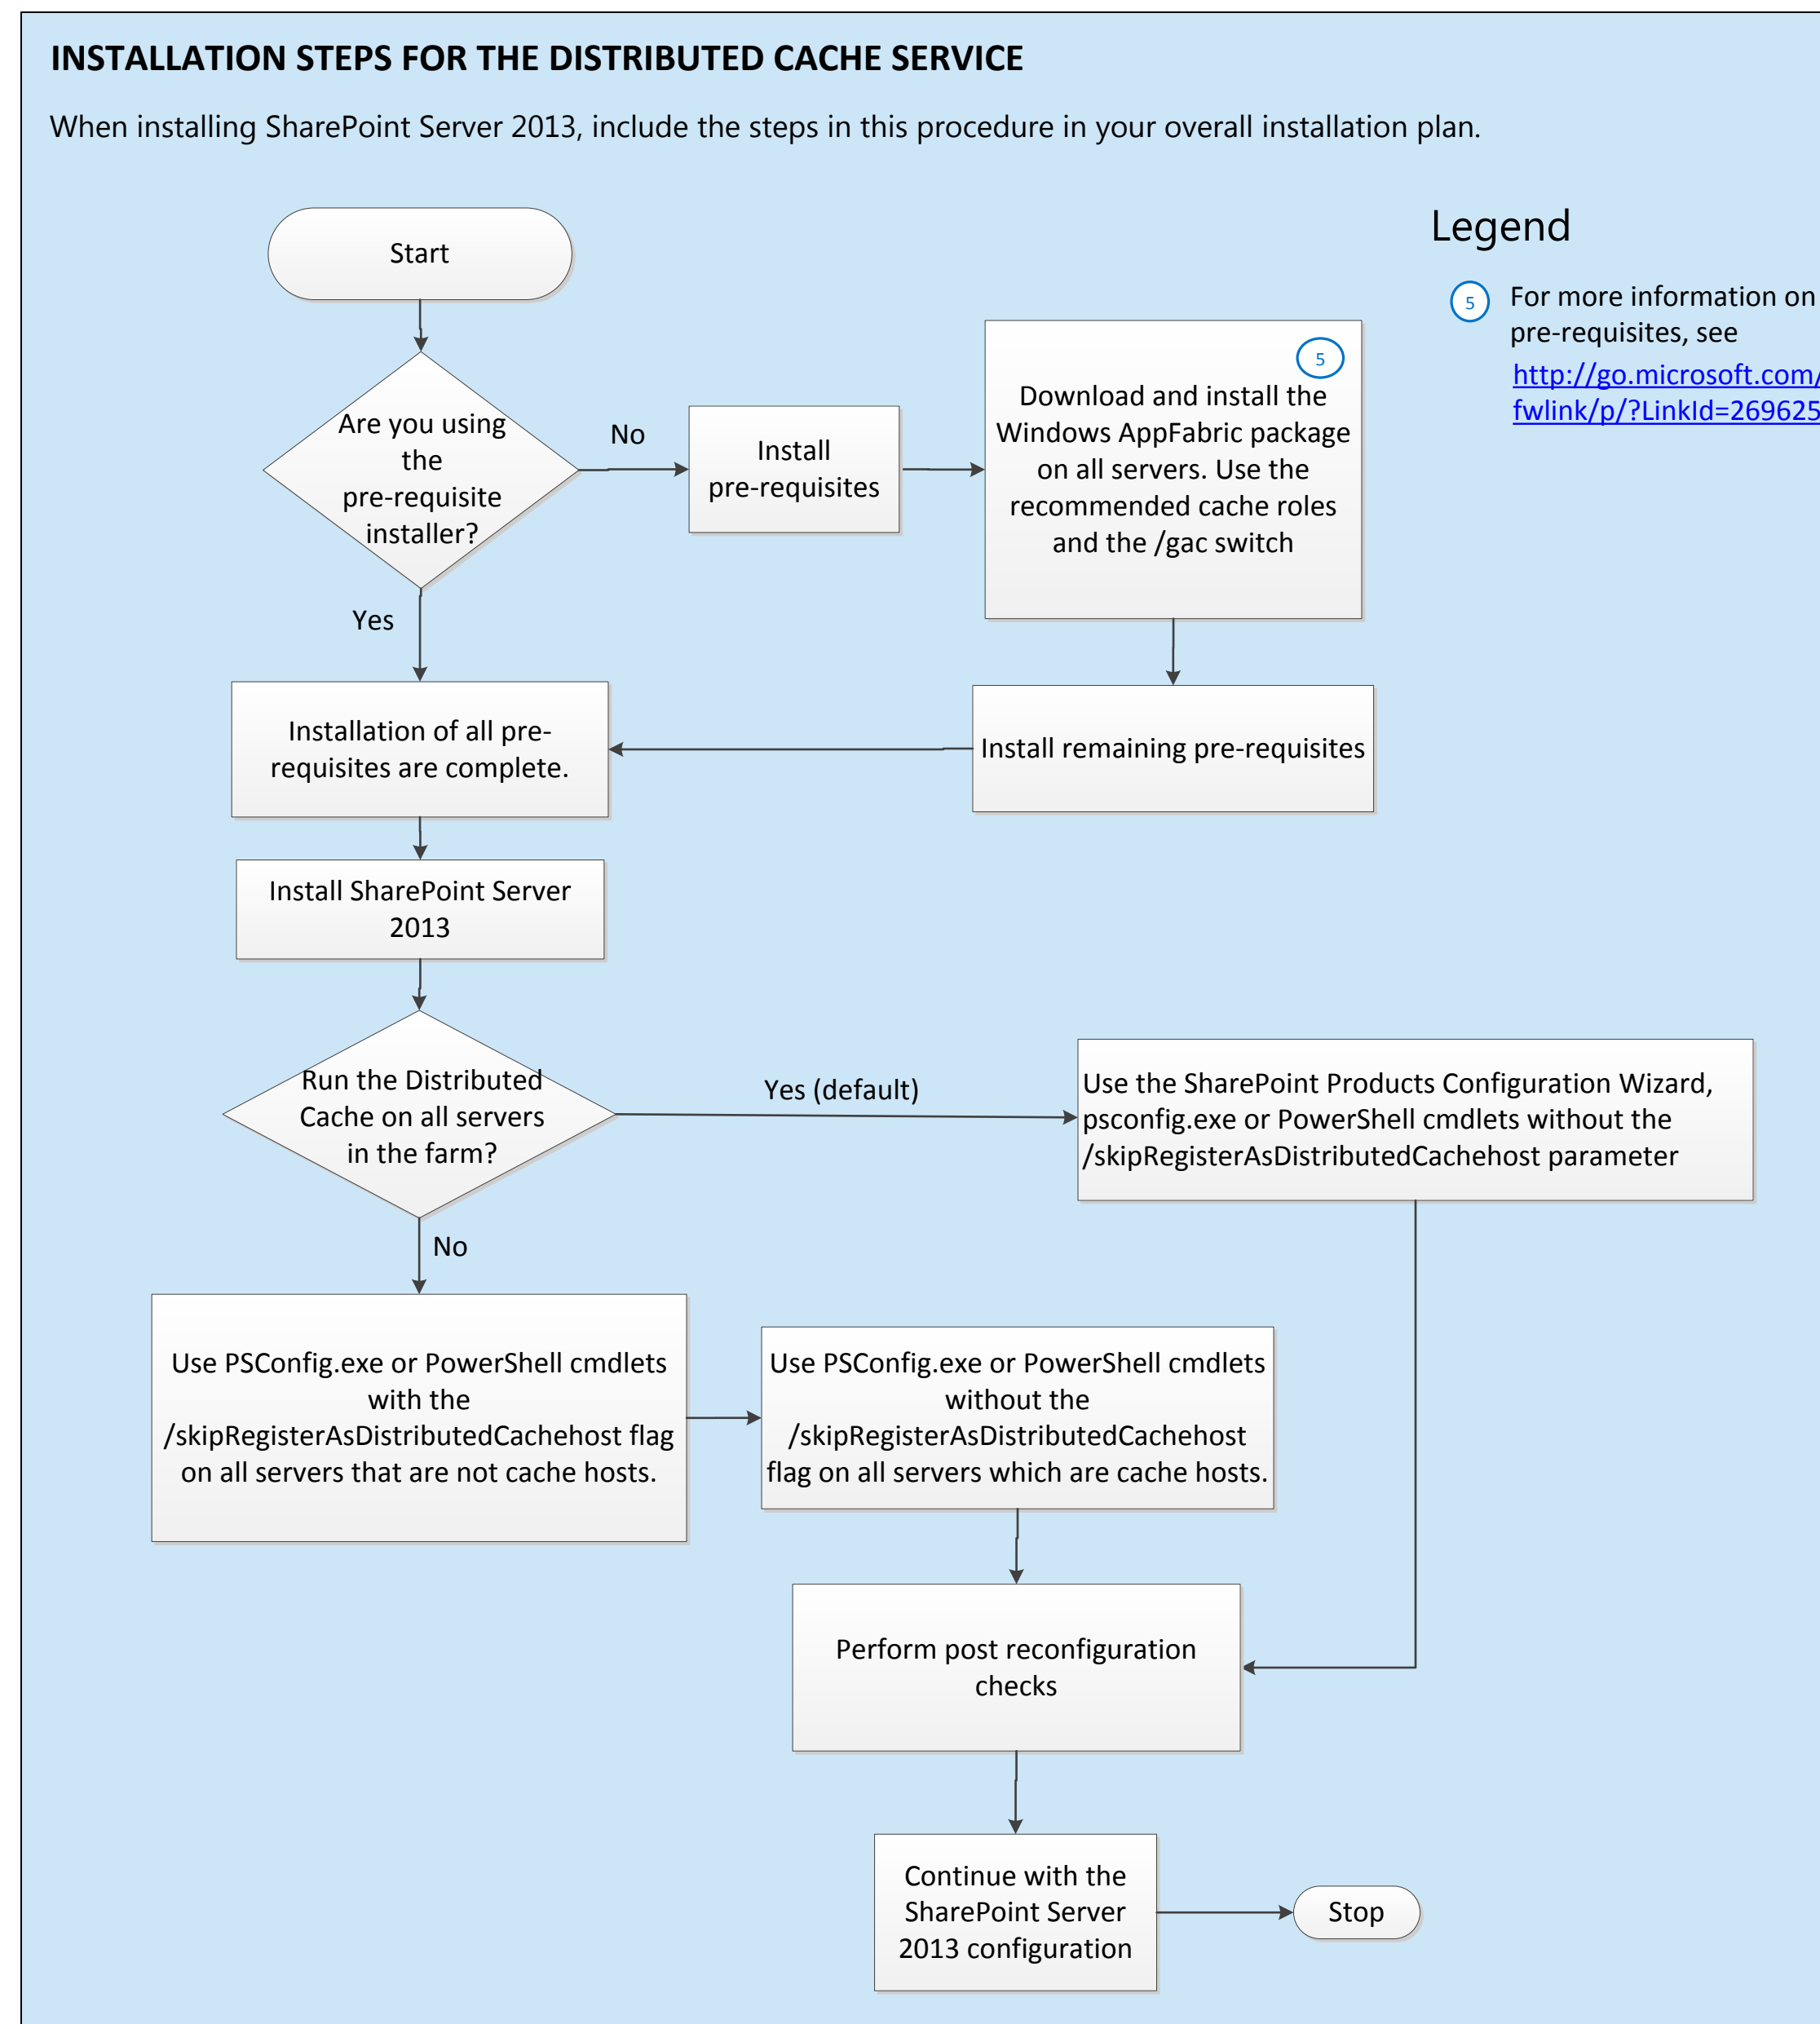
When the Distributed Cache service runs on a server together with other services and the server's memory resources near 95% utilization, the Distributed Cache will start throttling requests. This means that the Distributed Cache service will no longer accept read or write requests until the server utilization reduces to approximately 70% utilization.

Consider the following best practices to manage the Distributed Cache service:

- Do not run the following services on the same server running the Distributed Cache service: SQL Server, Project Server, Excel Services and Search services.
- Review SharePoint Server 2013's Health Rules for Distributed Cache service issues
- Change a non-cache host to a cache host, or a cache host to a non-cache host
- Add a new server to the cache cluster

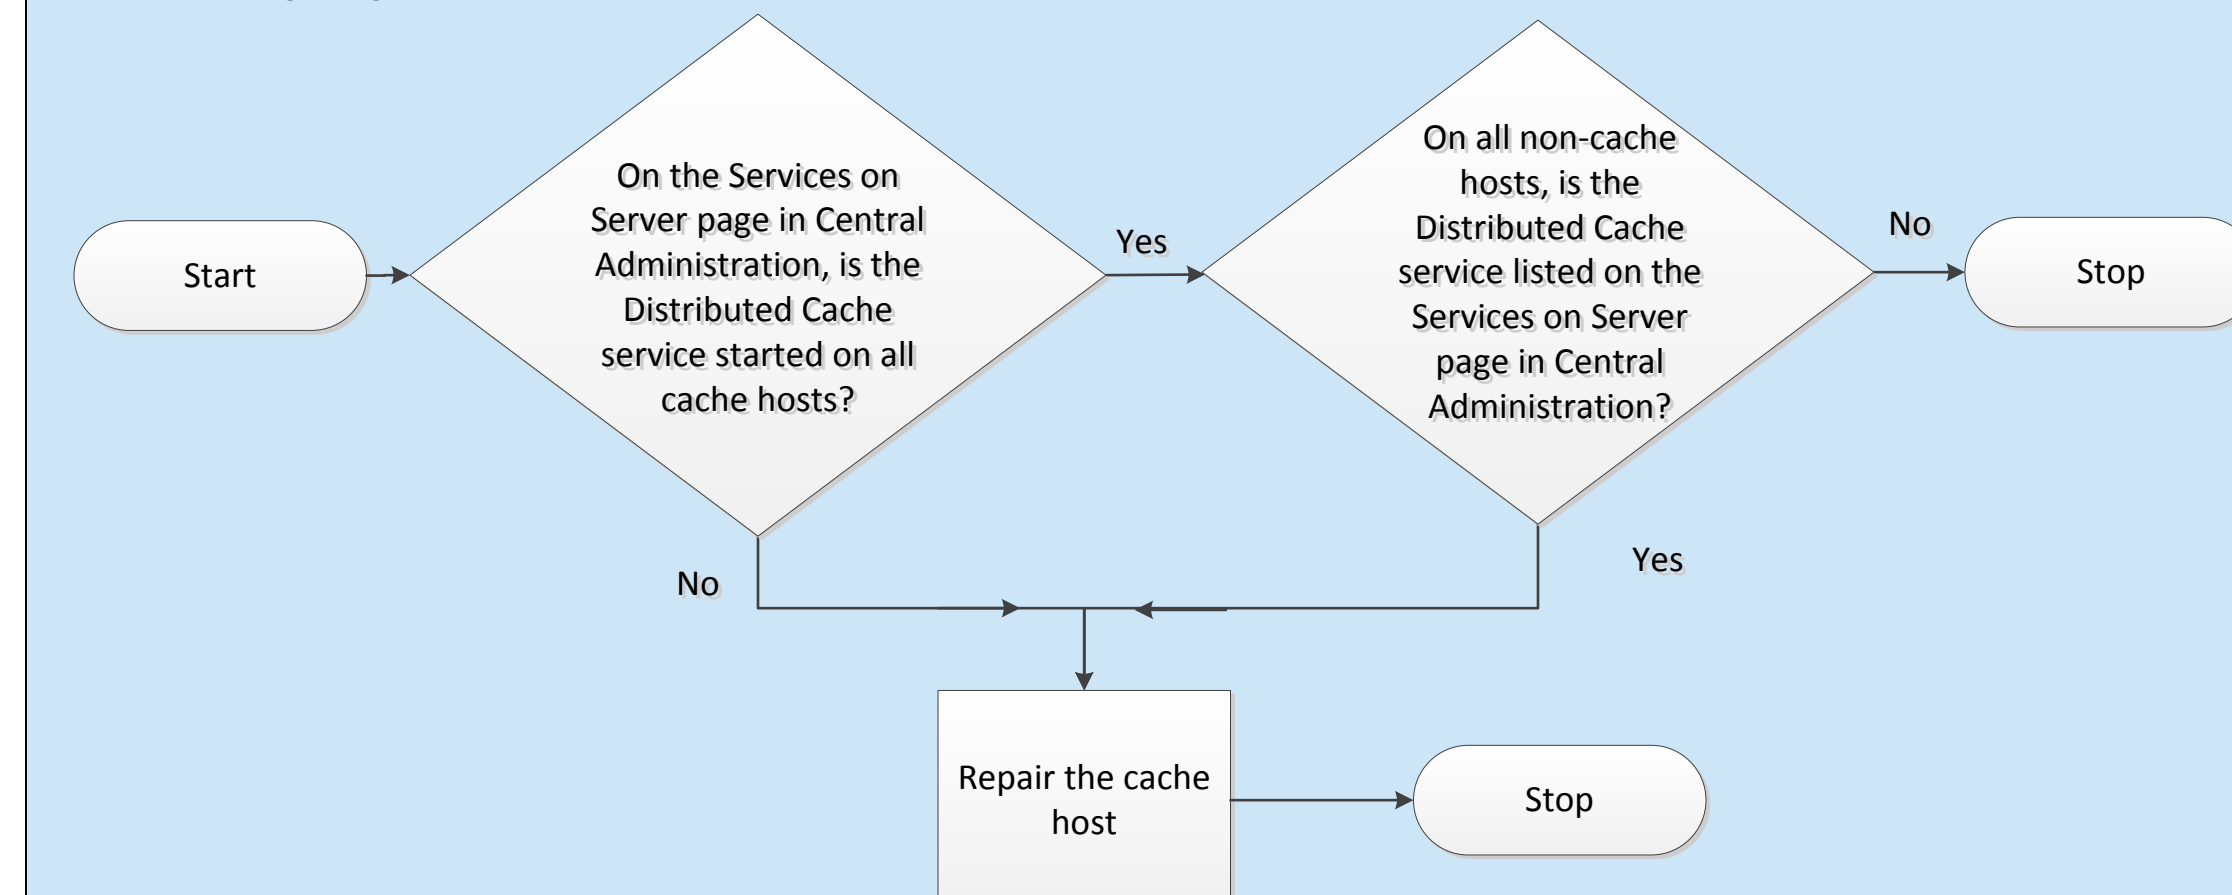
INSTALLING AND CONFIGURING THE DISTRIBUTED CACHE SERVICE

When installing and configuring SharePoint Server 2013, an administrator must perform certain steps to ensure the Distributed Cache service is installed and configured correctly.



POST RECONFIGURATION CHECKS

When reconfiguring the Distributed Cache service, perform this procedure to ensure the Distributed Cache service is available.



MAINTAINING THE DISTRIBUTED CACHE

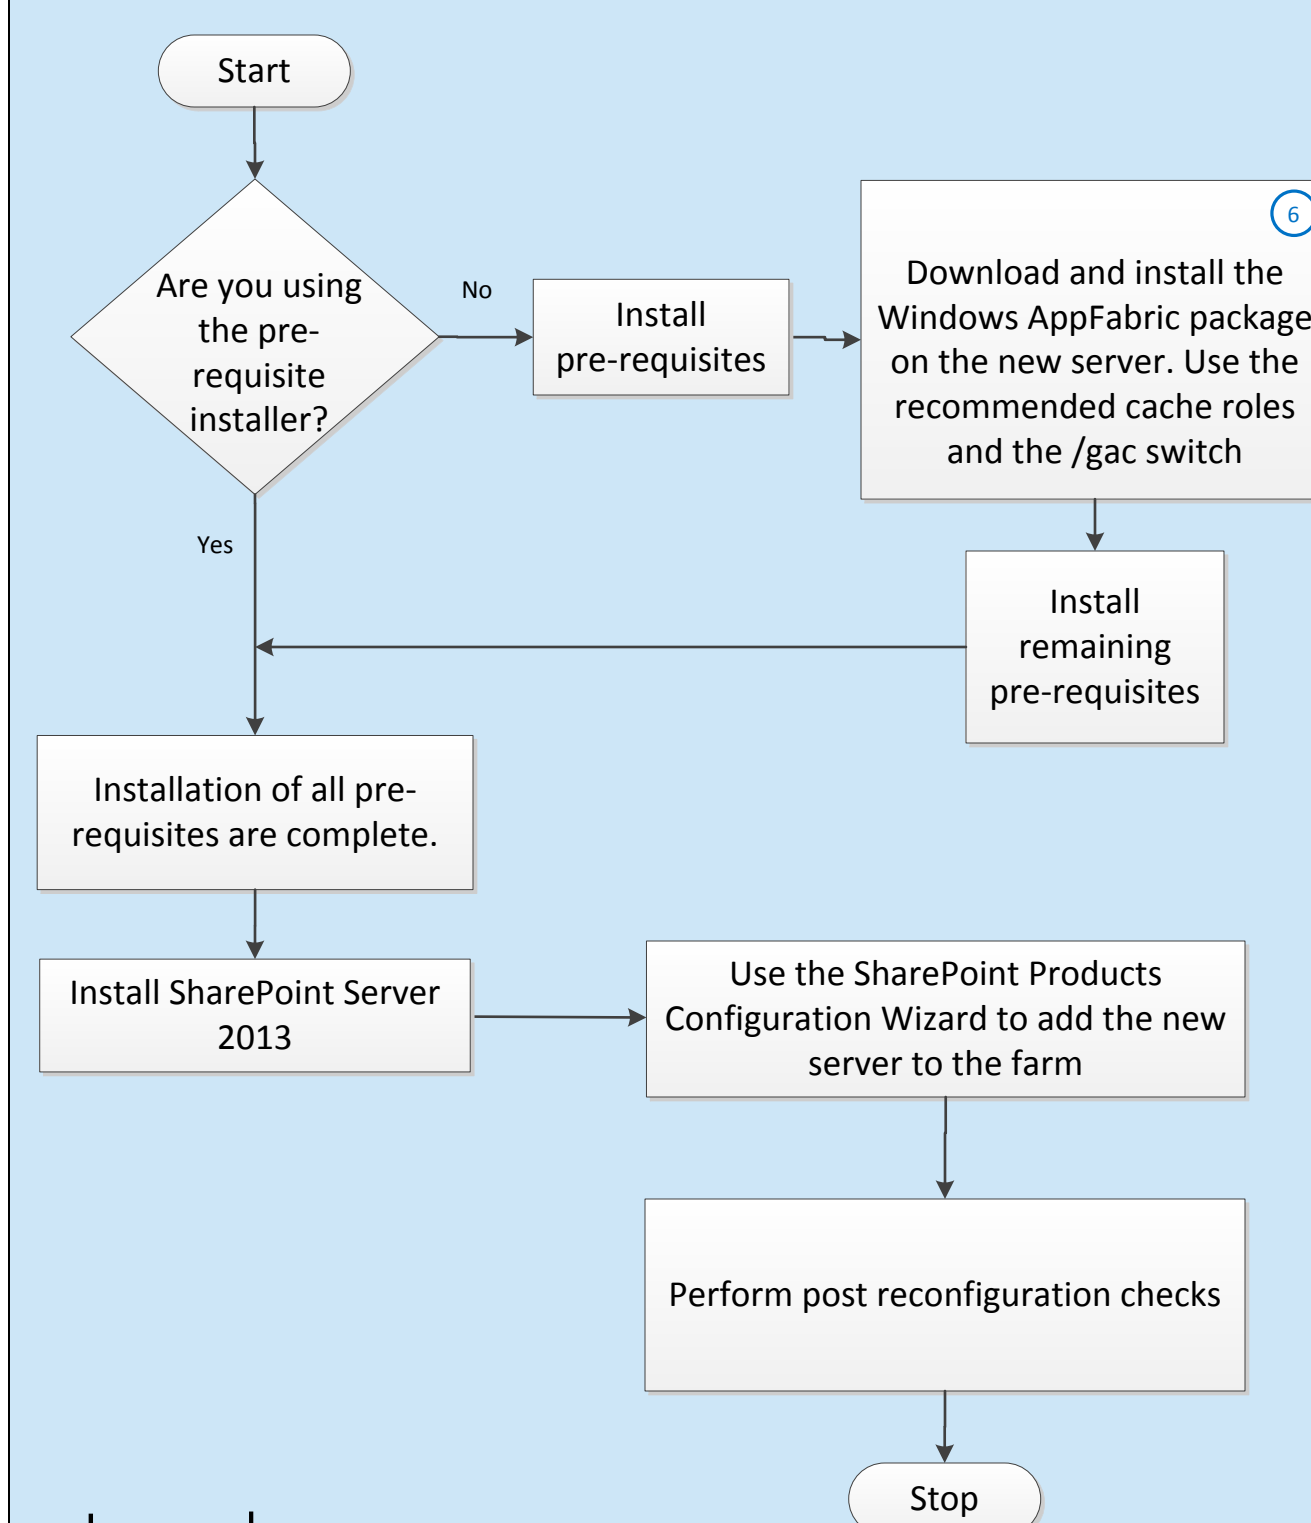
Administrators might need to perform required maintenance and operational tasks to manage their SharePoint Server 2013 deployment. The following flowcharts describe how to perform some of these maintenance activities as it relates to maintaining the Distributed Cache service.

INCREASING AVAILABLE RESOURCES FOR THE DISTRIBUTED CACHE SERVICE

As your SharePoint Server 2013 farm grows, you might need to increase the resources that are available to the Distributed Cache service by adding another cache host to the cache cluster. You can do this by either adding a new server to the farm (and the cache cluster), or changing a non-cache host to a cache host. After the new cache host is added, SharePoint Server will start using and populating the Distributed Cache service with data on the newly added cache host automatically. This may be a gradual process.

ADDING A NEW SERVER TO THE FARM AND THE CACHE CLUSTER

Use this procedure if you are adding a new server to the farm, and using it as a cache host

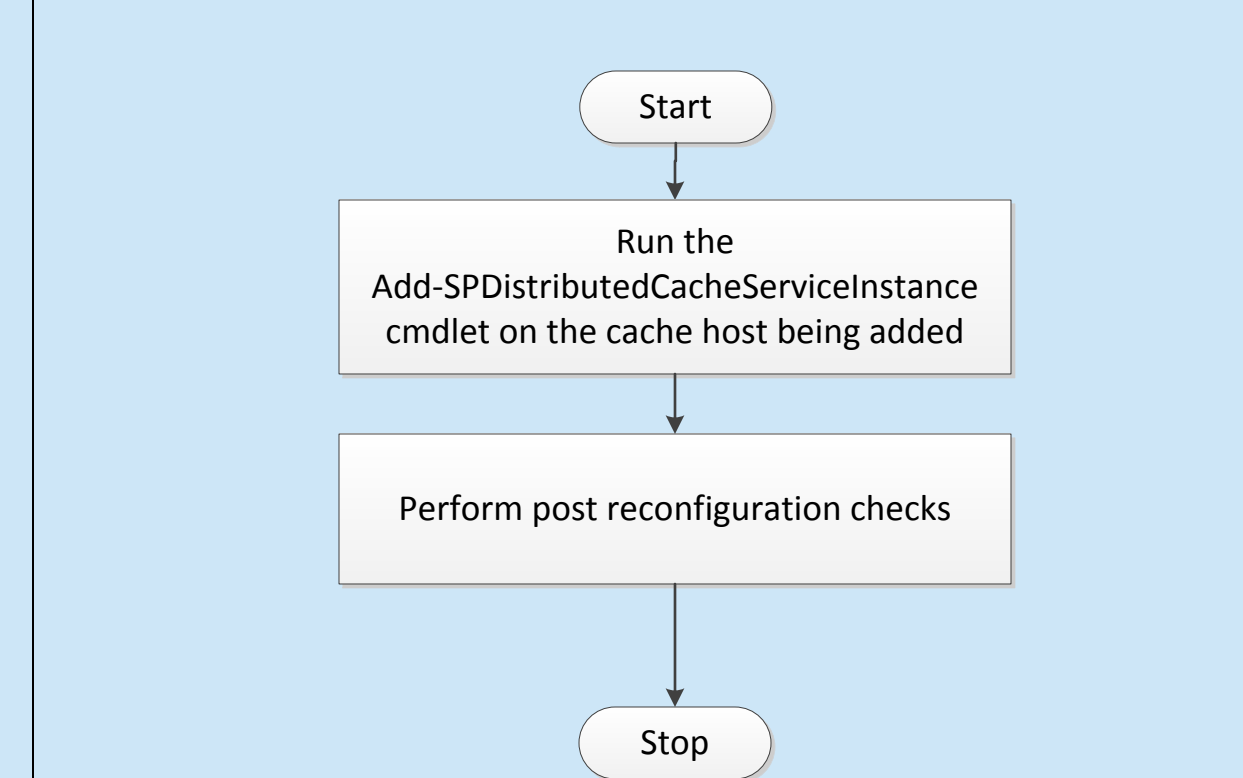


Legend

- 1 For more information on pre-requisites, see <http://go.microsoft.com/fwlink/?LinkID=269625>

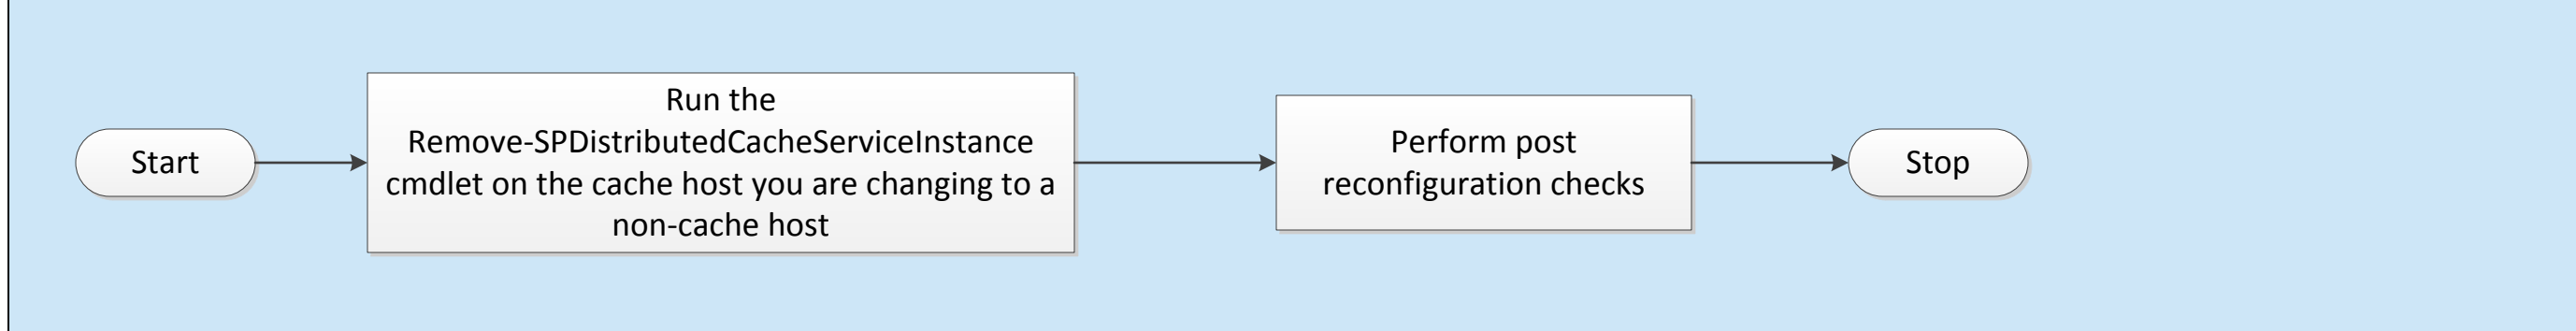
CHANGING A NON-CACHE HOST TO A CACHE HOST

Use this procedure if the server you want to use as a cache host is already joined to your server farm.



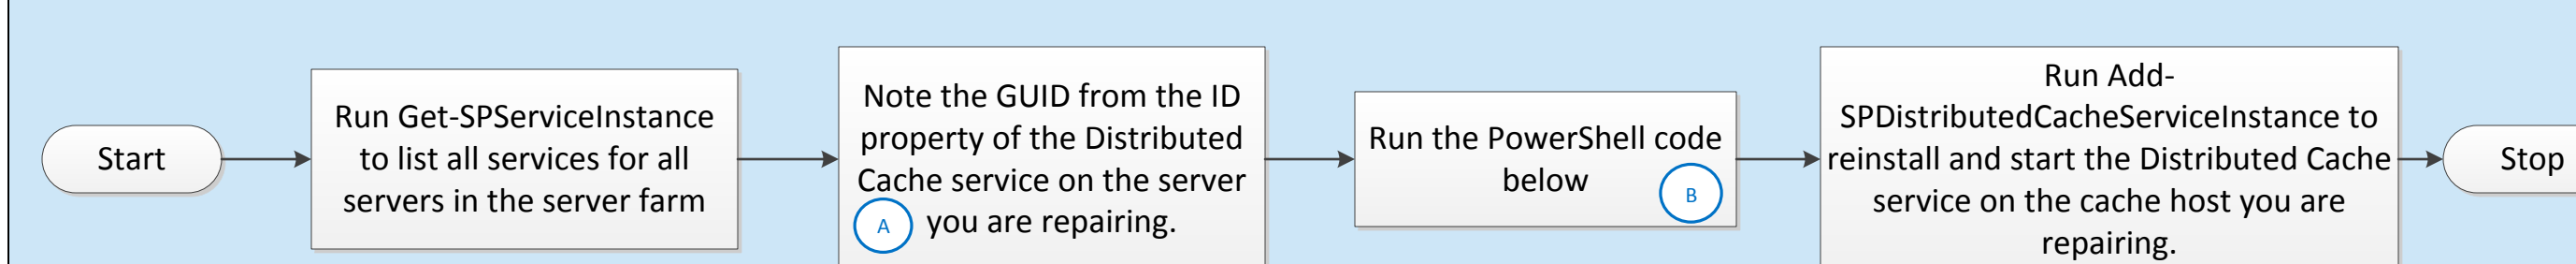
CHANGING A CACHE HOST TO A NON-CACHE HOST

Use this procedure if you want to change a cache host into a non-cache host.



REPAIRING A CACHE HOST

During installation, configuration, or maintenance activities, the Distributed Cache service may enter a non-functioning state. Evidence of a malfunctioning Distributed Cache service will appear in Health Rules in Central Administration or when users use features in SharePoint Server 2013 that rely on the Distributed Cache. For example, the Newsfeed on a user's My Site will start reporting errors. Use the following procedure to restore a non-functioning Distributed Cache host.



Legend

- A The Get-SPServiceInstance cmdlet lists all services on all servers in the server farm. Ensure you note the correct GUID from the server you are trying to repair.
- B \$s = get-sp-serviceinstance GUID \$s.delete() where GUID is the GUID of the Distributed Cache service running on the server being repaired.

RECONFIGURING THE MEMORY ALLOCATION

Use this procedure to adjust the cache size of the Distributed Cache service. Use this procedure during initial configuration of the Distributed Cache service, when changes are made to the total memory on the server, or when a new cache host is added to the server farm. Note: A default of 10% of total memory is assigned to the Distributed Cache service at installation time. You may want to increase this memory allocation depending on your requirements.

