

# Google Search Appliance

Planning for Search Appliance Installation

**Google Search Appliance software version 7.4**



**Google, Inc.**  
1600 Amphitheatre Parkway  
Mountain View, CA 94043  
[www.google.com](http://www.google.com)

GSA-PLAN\_200.03  
March 2015

© Copyright 2015 Google, Inc. All rights reserved.

Google and the Google logo are, registered trademarks or service marks of Google, Inc. All other trademarks are the property of their respective owners.

Use of any Google solution is governed by the license agreement included in your original contract. Any intellectual property rights relating to the Google services are and shall remain the exclusive property of Google, Inc. and/or its subsidiaries ("Google"). You may not attempt to decipher, decompile, or develop source code for any Google product or service offering, or knowingly allow others to do so.

Google documentation may not be sold, resold, licensed or sublicensed and may not be transferred without the prior written consent of Google. Your right to copy this manual is limited by copyright law. Making copies, adaptations, or compilation works, without prior written authorization of Google, is prohibited by law and constitutes a punishable violation of the law. No part of this manual may be reproduced in whole or in part without the express written consent of Google. Copyright © by Google, Inc.

# Contents

<b>Planning for Search Appliance Installation .....</b>	<b>5</b>
About This Document	5
How Does the Search Appliance Work?	6
Installation	6
Crawl	6
Traversal	7
Feeds	7
Indexing	7
Serving	8
About the End User License Agreement	8
How Do I Plan My Installation?	8
For Sites or Locations with New Content	8
For Sites or Locations with Existing Content	9
What Character Encoding Should Content Files and Feeds Use?	9
What Hardware and Software Do I Need?	9
What Does the Google Search Appliance Shipping Box Contain?	10
What File Types Can Be Indexed?	10
What File Sizes Can Be Indexed?	11
What Content Locations Can Be Crawled or Traversed?	11
How Many URLs Can Be Crawled?	12
How Do I Control Security?	12
How Can My Security Model Improve Performance?	12
Can the Search Appliance Use a Dedicated Network Interface Card for Administration?	13
What Ports Does the Search Appliance Use?	13
Ports Used at All Times	14
Additional Ports Used When the Dedicated Administrative Network Interface Card is Enabled	16
What User Accounts Do I Need?	16
How are Administration Accounts Authenticated?	17
How Do I Obtain Technical Support?	17
Support Requirements	17
How is Power Supplied to the Search Appliance?	17
How is Data Destroyed on a Returned Search Appliance?	18
What Values Do I Need for the Installation Process?	18
Required Values	18
Optional Values	20

What Tasks Do I Need to Perform Before I Install?	20
Required Tasks	21
Optional Tasks	22
Electrical and Other Technical Requirements	22

# Planning for Search Appliance Installation

This document provides the information you need to plan a Google Search Appliance installation. The guide contains an overview and checklists of the values you must determine and decisions you must make before you set up your network and the content files and perform the Google Search Appliance installation process. After you complete the installation process, the search appliance can crawl and index the content files. When the crawling and indexing processes are complete, end users can search the content files.

For information about specific feature limitations, see [Specifications and Usage Limits](#).

For planning information for other software versions and search appliance models, see the Archive page (<https://support.google.com/gsa/answer/2812980>), which contains links to previous versions of the search appliance documentation.

## About This Document

---

The information in this document applies to the Google Search Appliance models GB-7007, GB-9009, G100, and G500.

This document contains basic information about how the Google Search Appliance works. This document is for you if you are a network, web site, or content management system administrator, or if you install or configure the Google Search Appliance.

If you are installing a Google Search Appliance, you need some knowledge of networking concepts. These concepts include IP addresses, routers, dynamic host configuration protocol (DHCP), and ports.

If you are configuring the software, you'll need to know how your web site or intranet is structured and how the content you want to index and serve is structured.

If you are configuring a search appliance and a connector to index content in a content management system, you'll need to know about object types and properties in the content management system and about how the content management system's software is configured.

In GSA release 7.4, the on-board connector manager and connectors are deprecated. They will be removed in a future release. If you have configured on-board connectors for your GSA, install and configure an off-board Google Connector. For more information, see the documentation that is available from the [Connector Documentation page](#).

# How Does the Search Appliance Work?

---

The Google Search Appliance is a one-stop search and index solution for businesses of all sizes. Using a search appliance, you can quickly deploy search within an enterprise. By default, a search appliance can index and serve content located on a file system or a web server. You can also configure the Google Search Appliance to use a connector manager and a connector to index and serve content located in a content management system such as EMC Documentum or Microsoft SharePoint.

The search appliance comes with Google software installed on powerful hardware, simplifying the planning process because you do not need to choose a hardware platform.

The Google Search Appliance model GB-7007 can be licensed for up to 10 million documents and the Google Search Appliance model GB-9009 for up to 30 million documents.

The Google Search Appliance model G100 can be licensed for up to 20 million documents and the Google Search Appliance model G500 for up to 100 million documents.

This section contains an introduction to the basic operations of the Google Search Appliance and descriptions of the preinstallation planning process.

## Installation

Before an intranet, web site, or content repository can be indexed, you must install the search appliance on your network and set up the software on the appliance. Installing the search appliance requires physically attaching it to the network and then starting the search appliance.

Setting up the software on a search appliance includes the following tasks:

- Ensuring that the correct ports are available on your network.
- Providing correct network settings, so that the search appliance can communicate with the computers on the network
- Providing email and time settings
- Assigning a password to the administrator account
- Ensuring that the search appliance has access to the file system or web servers where the content files are located.
- Configuring the initial crawl of your file system or web servers

If you are indexing content in a content management system, you must also install a connector manager and the connector for the particular content management system. Review the documentation set for the correct connector software version (<https://support.google.com/gsa/topic/4566684>), which provides information on preinstallation tasks, required software, and required hardware for the connector manager and connectors.

## Crawl

Crawl is the process by which the Google Search Appliance locates content to be indexed. Crawl is a pull process, where the search appliance pulls content from the content location. The search appliance can also crawl a relational database to obtain metadata.

When you configure the software for crawling, you define three sets of URLs, which can be in HTTP or server message block (SMB) format:

- Start URLs, which control where the crawl begins. All content must be reachable by following links from one or more start URLs.
- Follow Patterns, which set the patterns of URLs that are crawled. Use follow and crawl URLs to define the paths to pages and files you want crawled. If a URL in a crawled document links to a document whose URL does not match a pattern defined as a follow and crawl URL, that document is not crawled.
- Do Not Follow Patterns, which designate paths to pages and files you do not want crawled and file types you do not want crawled.

If the search appliance is crawling a web site, the crawl software issues HTTP requests to retrieve content files in the locations defined by the URLs and to retrieve files from links discovered in crawled content. If the search appliance is crawling a file share, the crawl software uses the SMB or common Internet file system (CIFS) protocol to locate and retrieve the content files. For more information on crawl, see *Administering Crawl*, which also includes checklists of crawl-related tasks in the “Crawl Quick Reference.”

## Traversal

Traversal is the process by which the Google Search Appliance locates content to be indexed in a content repository such as SharePoint or Lotus Notes. Traversal is a process in which the connector issues queries to the repository to retrieve document data to feed to the Google Search Appliance for indexing.

## Feeds

Feeding is the process by which you direct content to the Google Search Appliance instead of having the search appliance locate content. Feeding is a push process, in which the content files are pushed to the Google Search Appliance. You can feed several types of content to a Google Search Appliance:

- A list of URLs  
The crawl software fetches documents listed in the URLs.
- Content files  
The files and their URLs are fed to the search appliance.
- External metadata that is not stored in a relational database or where it is difficult to map the metadata to the content file

For more information on feeding, see the *Feeds Protocol Developer's Guide* and *External Metadata Indexing Guide*.

## Indexing

Indexing is the process of adding the content from the crawled documents to the index.

After a file is retrieved by the crawl, the file is converted to an HTML file and submitted for indexing. The indexing process extracts the full text from each content file, breaks down the text, and adds both the text and information such as date and page rank to the index so that users' search requests can be satisfied. The index and the HTML versions of each indexed file are stored on the search appliance.

## Serving

Users submit search requests to the Google Search Appliance a web page similar to the search page at Google.com. A user types a search term into the search box and the request is transmitted to the serving software. The search appliance locates results in the index. The search appliance then returns the results to the user's browser as a series of links. When the user clicks a link in the results, the content file is displayed.

You can customize the behavior and appearance of the search page from the Admin Console, which you use to administer and configure the search appliance. For complete information on customizing the search page and other aspects of the user experience, see *Creating the Search Experience*.

## About the End User License Agreement

---

During the initial search appliance configuration process, you must accept an End User License Agreement. Google recommends that you copy the license agreement when you view it. After the configuration process is complete, you cannot view the End User License Agreement again.

## How Do I Plan My Installation?

---

Before you install the Google Search Appliance, follow one of the high-level preinstallation workflows below to ensure that the installation goes smoothly.

### For Sites or Locations with New Content

1. Determine the physical location of the search appliance.
  - Will the search appliance be installed in a data center or in your office?  
  
If you install the search appliance in an office, place it in an area where any noise produced by the cooling fan in the search appliance will not be disturbing.
  - Does the location meet the electrical and temperature requirements described in "Electrical and Other Technical Requirements" on page 22.
2. Analyze your business's content and decide which files you want indexed.
3. Decide whether to use a database feed.  
  
Use a database feed to associate metadata with a corresponding content file and include the metadata in the index. If you are indexing a content management system, the connector automatically associates metadata from the repository with the appropriate content file.
4. Determine which files are public and can be viewed by any person performing a search of the content.
5. Determine how to provide the appropriate security for files that are not fully public. For example, store confidential files in locations that are not crawled or use a security model that requires authorization to view content that you want to protect from public view. For more information, see *Managing Search for Controlled-Access Content*.
6. Design a directory structure for the web site or intranet that supports the desired security model.



7. Implement any authorization and authentication requirements.
8. Deploy the content files to your web site or intranet.
9. Complete the tasks and collect the information and values described in the tables in “What Values Do I Need for the Installation Process?” on page 18 and “What Tasks Do I Need to Perform Before I Install?” on page 20.

## For Sites or Locations with Existing Content

1. Determine the physical location of the search appliance.
  - Will the search appliance be installed in a data center or in your office?  
  
If you install the search appliance in an office, place it in an area where any noise produced by the cooling fan in the search appliance will not be disturbing.
  - Does the location meet the electrical and temperature requirements described in “Electrical and Other Technical Requirements” on page 22?
2. Decide which directories and files you want indexed.
3. Decide whether to use a database feed.  
  
Use a database feed to associate metadata with a corresponding content file and include the metadata in the index. If you are indexing a content management system, the connector automatically associates metadata from the repository with the appropriate content file.
4. Determine which files are public and can be viewed by any person performing a search of the content.
5. Determine how the security model used to protect your content files can be reflected by the security models on the search appliance. For example, store confidential files in locations that are not crawled or use a security model that requires authorization to view content that you want to protect from public view. For more information, see *Managing Search for Controlled-Access Content*.
6. Complete the tasks and collect the information and values described in the tables in “What Values Do I Need for the Installation Process?” on page 18 and “What Tasks Do I Need to Perform Before I Install?” on page 20.

## What Character Encoding Should Content Files and Feeds Use?

---

Google strongly recommends that you use the UTF-8 character encoding for feeds and for documents that will be in the index.

## What Hardware and Software Do I Need?

---

You need the following hardware and software to install and support the search appliance:

- A laptop or desktop computer that can be physically connected to the search appliance using wired Ethernet. The laptop or desktop computer can be a Windows or Macintosh computer. The necessary cables are included with the search appliance.

- A web browser, which must be a current version of Internet Explorer, Firefox, or Chrome. The *Guide to Software Release 7.4* lists the supported browsers.
- An uninterruptible power supply to provide electricity to the search appliance if there is a power failure. The index data cannot be backed up, and the crawl, index, and serve functions will be interrupted if there is a power failure. If you require the search appliance to be highly available, consider using a second search appliance as a hot backup unit.

In some circumstances, Google for Work Technical Support may ask you to attach a USB keyboard and monitor directly to the search appliance so that you can manually restart the search appliance.

## What Does the Google Search Appliance Shipping Box Contain?

---

The Google Search Appliance shipping box contains the following:

- One Google Search Appliance model GB-7007, GB-9009, G100, or G500
- Four cables
  - One orange Ethernet cable
  - One yellow Ethernet cable
  - Two power cords, localized for your region
- Rail kit
- Rack installation guide
- Bezel key
- Product information guide
- A printed welcome letter

## What File Types Can Be Indexed?

---

The Google Search Appliance can crawl and index more than 200 different file formats, including:

- HTML
- Portable Document Format (PDF)
- Text files
- Common word-processing and spreadsheet files

The exact file formats and versions depend on the software version installed on a particular search appliance. For a complete list, see *Indexable File Formats*.

A search appliance can also index metadata associated with content files. The metadata can be in HTML meta tags. Metadata can be fed from a database and then indexed.

The Google Search Appliance cannot index text contained in graphic file formats, such as JPEG, GIF, or TIFF. When a file in a graphic format is submitted for indexing, text embedded in the graphic is not indexed. However, the file name is indexed. If any metadata is associated with the graphic in an HTML meta tag that metadata is indexed.

Certain file formats are excluded from the crawl by default on the search appliance Admin Console. When you configure the crawl, ensure that the field for excluded URLs and file formats correctly reflects the file types you do not want crawled and indexed.

## What File Sizes Can Be Indexed?

---

By default, the search appliance indexes up to 2.5MB of each text or HTML document, including documents that have been truncated or converted to HTML. After indexing, the search appliance caches the indexed portion of the document and discards the rest. You can change the default by entering a new amount of up to 10MB.

To change the default amount, use the **Index > Index Settings** page in the Admin Console.

## What Content Locations Can Be Crawled or Traversed?

---

The Google Search Appliance can crawl files located on an intranet or a web site.

If you install a connector, the Google Search Appliance can also traverse content located in a content repository such as FileNet or Documentum. For more information, see the [Google Search Appliance Connector documentation](#).

Content on a web site is crawled using the HTTP or HTTPS protocol.

Content on an intranet is crawled using the SMB or CIFS protocol. Intranet files are typically stored in a Windows shared directory or in a web-enabled virtual directory. See the Windows Help system for information on creating a shared directory. You can create a virtual directory in several ways:

- By using the Virtual Directory Creation Wizard of Internet Information Services (IIS)
- By importing a configuration file
- By using the `lisvdir.vbs` script
- By using the Apache web server to enable directory browsing

For more information on creating virtual directories, see the Windows Help system.

Content files can also be located on Macintosh, UNIX, or Linux computers on an intranet. On Macintosh computers, use the CIFS protocol. On UNIX or Linux computers, you can web-enable the file locations and use HTTP or HTTPS for crawling, or you can use the SMB protocol without web-enabling the locations.

If a file is in a location that requires a password for access, whether on an intranet or a web site, you must provide a user ID and password for the location on the Crawler Access page of the Admin Console.

## How Many URLs Can Be Crawled?

The number of URLs that your search appliance can crawl depends on the model and license limit. The following table lists the maximum number of URLs matching the crawl patterns you define that the search appliance can crawl.

Search Appliance Model	Maximum License Limit	Maximum Number of URLs that Match Crawl Patterns
GB-7007	10 million	~ 13.6 million
GB-9009	30 million	~ 40 million
G100	20 million	~26 million
G500	100 million	~133 million

## How Do I Control Security?

Your business may require you to restrict access to certain enterprise content. You might want to restrict what content is crawled and indexed, and you might want to restrict which users have access to particular content. The Google Search Appliance supports various security models:

- You can exclude content from the index by storing the content in locations that are not crawled.
- You can exclude content from the index by using a robots.txt file to prevent particular locations from being crawled.
- You can require the search appliance to provide credentials before crawling particular locations.
- You can design an authentication model under which users who cannot be authenticated are not able to see particular content.
- You can design an authorization model that defines which users are authorized to perform certain functions on particular documents.

The search appliance supports a range of authentication and authorization methods, including HTTP Basic, Windows NT LAN Manager Authentication (NTLM), HTML forms-based authentication, certificate authentication, lightweight directory access protocol (LDAP) directory servers, Authentication and Authorization SPI.

For information on how to configure crawl for your security model, see *Administering Crawl*. For information on how to integrate your search appliance with different authentication and authorization models, see *Managing Search for Controlled-Access Content*.

## How Can My Security Model Improve Performance?

Using policy ACLs and per-URL ACLs to control which users have access to content located in particular URLs speeds up the process of authorization and improves search appliance performance. For more information on ACLs, see *Managing Search for Controlled-Access Content*.

## Can the Search Appliance Use a Dedicated Network Interface Card for Administration?

---

You can optionally configure your search appliance to use a dedicated network interface card (NIC) for administrative functions.

Use this option when you have two or more search appliances that are typically accessed through a load balancer. Search appliances in this configuration are reached at the same IP address and there is no way to connect to a specific search appliance in the configuration.

Using a dedicated network interface card for administrative purposes enables you to ensure that you are connecting to the correct search appliance when you need access to the Admin Console of that particular appliance. This option is available only on search appliances that have four Ethernet ports. Older search appliances with two Ethernet ports cannot use a dedicated administrative network interface card.

To use the dedicated NIC, you must specifically assign an IP address, a subnet mask, and a gateway to the NIC. The IP address and subnet must be different from the primary search port (LAN1) and the orange port (LAN3, 192.168.255.0/24). In other words, all three GSA network interfaces must be in different subnets.

## What Ports Does the Search Appliance Use?

---

The search appliances use many ports to send and accept requests. The following sections describe the outbound and inbound ports that are used depending on whether you enable the dedicated administrative network interface card on the search appliance.

## Ports Used at All Times

The following table lists the outbound search appliance ports.

Outbound Ports	Function
25	Sends SMTP requests
51	AH protocol of IPsec, which is used for communication among search appliances in multi-node configurations
53	Sends DNS (UDP) requests
80	Sends HTTP crawl and search requests
123	Sends NTP requests
139	Sends NETBIOS requests for SMB crawling
389	Sends LDAP requests.
443	Crawls secure content and connects to the SupportCall server.
445	Sends Microsoft CIFS requests for SMB crawling
500	UDP IPsec IKE key exchange protocol port. IPsec is used for communications among search appliances in multi-node configurations.
514	Sends SYSLOG requests
7885	Used for exporting logs for the preinstalled connector manager at the URL <code>https://search_appliance_location:7885/connector-manager/getConnectorLogs/ALL</code>
8080	Default port for requests to the connector manager when a connector manager is installed on an external host. Configurable when a connector manager is installed.

The following table lists the inbound search appliance ports.

Inbound Ports	Function	Protocol	When Open
22	Communication with the search appliance for Support use.	SSH	When enabled
51	Communication among search appliances in multi-node configurations	AH protocol of IPsec	Always open
80	Accepts search requests	HTTP	Always open
161	Accepts SNMP requests, both TCP and UDP	SNMP	Always open
443	Accepts search requests  This is the default port for HTTPS access. HTTPS requests sent to port 443 are automatically routed by the search appliance to the correct service. We recommend that you access HTTPS services using port 443.	HTTPS	Always open
4430	Used for secure connections. Also, the search appliance forwards traffic from port 443 to 4430.	HTTPS	Always open
500	IPsec IKE key exchange protocol port	UDP	Always open
7843	Connector manager and security manager	HTTPS	Always open
7886	Connector manager and security manager	HTTP	Always open
8000	Accepts requests for the Admin Console, the search appliance's administrative interface	HTTP	Always open
8443	Accepts requests for the Admin Console, the search appliance's administrative interface	HTTPS	Always open
9941	Accepts requests to the search appliance's Version Manager utility	HTTP	Always open
9942	Accepts requests to the search appliance's Version Manager utility	HTTPS	Always open
19900	Accepts HTTP POST for XML feeds, including from connectors	HTTP/XML	Always open
19902	Accepts HTTPS POST for XML feeds	HTTPS/XML	Always open

## Additional Ports Used When the Dedicated Administrative Network Interface Card is Enabled

When the dedicated network interface card is enabled on a search appliance, the ports in the following table are visible *only* on the dedicated interface card except for SSH (22). Port 22 is always available on both the main (yellow) NIC and on the dedicated administrative NIC. All other port usage remains as detailed in the section “Ports Used at All Times” on page 14.

Port	Function	Protocol	When Open
22	Communication with the search appliance for Support use.	SSH	When enabled
161/162	Accepts SNMP requests, both TCP and UDP	SNMP	Always open
514	Sends SYSLOG requests, UDP	SYSLOG	
8000	Accepts requests for the Admin Console, the search appliance’s administrative interface	HTTP	Always open
8443	Accepts requests for the Admin Console, the search appliance’s administrative interface	HTTPS	Always open
9941	Accepts requests to the search appliance’s Version Manager utility	HTTP	Always open
9942	Accepts requests to the search appliance’s Version Manager utility	HTTPS	Always open

## What User Accounts Do I Need?

You need the following accounts to use with the search appliance:

- One or more administration accounts on the search appliance itself

The default administration account has the user name admin and the password that you assigned during initial installation. You can create additional administration accounts after you install the search appliance, with two different levels of user privileges. These accounts are administrator or manager accounts. The administration account with the user name admin must be used to run the network configuration wizard during the initial search appliance configuration process and to connect to the Version Manager, which is used to update the search appliance software.

- User accounts required for access to content files that you want the search appliance to crawl

If the content files you want crawled and indexed are in a location that requires a login, create a special user account on your network for the search appliance. When you configure crawl on the Admin Console, provide the user name and password for that account. The search appliance will present those credentials before crawling files in that location.



## How are Administration Accounts Authenticated?

---

During search appliance installation, you choose among different means of authenticating administration users.

- Under **Local Authentication**, administrators and managers are authenticated using credentials you enter directly on the Admin Console.
- Under **LDAP Authentication**, administrators and managers are authenticated against an LDAP server. To use this option, you must initially connect to the Admin Console using the admin account and the password you assign the account during configuration, then provide settings for the LDAP administrator group and the LDAP server itself. After you save the LDAP information, LDAP authentication for administrators and managers takes effect.
- Under **Local and LDAP** authentication, the search appliance attempts to authenticate administrators and managers against both the local credentials and the LDAP server. If an account can be authenticated against either the local credentials or the LDAP server, the login attempt succeeds.

## How Do I Obtain Technical Support?

---

For complete information on obtaining technical support, refer to the web page at [http://www.google.com/support/enterprise/go/gsa\\_support](http://www.google.com/support/enterprise/go/gsa_support).

## Support Requirements

Under the terms of the Support Agreements for the Google Search Appliance, Google for Work Support requires direct access to your search appliance to provide some types of support. For example, direct access is needed to determine whether your search appliance is eligible to be returned to Google and exchanged for a new search appliance. Different access methods have different requirements. The requirements for remote access are discussed in *Remote access methods for technical support* (<https://support.google.com/gsa/answer/2644822>).

## How is Power Supplied to the Search Appliance?

---

The Google Search Appliance models GB-7007, GB-9009, G100, and G500 are provided with two redundant power supplies.

# How is Data Destroyed on a Returned Search Appliance?

When a Google Search Appliance is returned to Google, these precautions are taken to remove customer data:

- The data on each drive is removed using software that conforms with the DOD 5220.22-M standard.
- Defective drives are physically destroyed by being crushed.

## What Values Do I Need for the Installation Process?

The following tables describe the values you need before you install the Google Search Appliance. If you are indexing a content repository, refer to the connector documentation for more information on values you need before installing the connector manager and a connector.

### Required Values

Before you install and configure the Google Search Appliance, obtain the following required values and write them in the column labeled **Your Value**. Most of these values will be provided by your network administrator.

Value	Definition	Your Value
A static IP address for the search appliance (IPv4 and IPv6)	<p>The static IP address identifies the permanent network location of the search appliance. Both IPv4 and IPv6 addresses are valid. A search appliance cannot use DHCP to obtain static IP addresses directly from the network. You cannot assign a static IP address to the search appliance that is in the range 192.168.255.[0-255]. The search appliance must not be on the same subnet as 192.168.255.[0-255] and cannot directly communicate with hosts that are assigned IP addresses in that range.</p> <p>You can assign a host name to the search appliance in addition to a static IP address. If you use a host name to access to the search appliance, you have more flexibility in moving the physical location of the search appliance or changing the IP address of the search appliance.</p>	
The subnet mask for the subnet on which the search appliance is located (IPv4 configuration only)	The subnet mask identifies the subnet on which the search appliance is located. It is used to determine whether the search appliance and other computers are on the same network.	
Prefix length (IPv6 configuration only)	The length of the address used for search appliance administration.	
The IP address of the default gateway or router (IPv4 and IPv6)	This IP address identifies the router to which the search appliance routes network traffic directed to any host outside the local subnet. The IP address must be on the same subnet as the search appliance.	

Value	Definition	Your Value
The IP address or addresses of network time protocol (NTP) servers	These IP addresses identify servers that synchronize computer times on the internet. The search appliances require accurate time settings to record correct time stamps in logs, track license expirations, and crawl or recrawl documents at the correct times. It is best to identify at least three accessible NTP servers for the search appliance to use. The NTP servers can be public or private. Do not attempt to operate a search appliance without identifying at least one NTP server. For more information, refer to <a href="http://support.ntp.org/bin/view/Servers/WebHome">http://support.ntp.org/bin/view/Servers/WebHome</a> .	
The user names, passwords, and email addresses for administrative users	These identify the users who access and administer the search appliance. The accounts are configured on the Admin Console. During the installation process, you must provide a password for the default account, which has the user ID admin.	
The IP address of one or more domain name system (DNS) servers	These IP addresses identify DNS servers used to resolve host names. Identifying DNS servers enables the use of host names, rather than IP addresses, in crawl URLs when the search appliance crawls an intranet. The search appliance will not operate correctly without functioning DNS servers.	
The DNS suffix, which is also called the DNS search path	The DNS suffix provides possible alternative expansions for host names when a fully-qualified domain name is not used in a URL. For example, if the DNS suffix is mydomain.org and a host name is myhost, the DNS suffix is used to example myhost to myhost.mydomain.org. You can enter NULL during the configuration process, which means that no value has been set for the DNX suffix.	
The email addresses of users who will receive notifications sent by the search appliance	The search appliance sends messages containing status reports and problem reports. A single email address can receive both types of reports or different email addresses can be identified for the two types of reports. A mailing list or mail alias can also be used.	
The email address used to send email from the search appliance	This account is used to send email messages and alerts from the search appliance to administrators or end users. The default value is nobody@localhost.	

## Optional Values

Depending on how your network is configured and on your administration needs, you can optionally obtain these values to use when you install and configure the search appliance.

Value	Definition	For More Information, Contact
The fully-qualified name of a simple mail transfer protocol (SMTP) server on the network	The fully-qualified name identifies the mail server used by the search appliance to send email. During installation, you can provide an invalid name and installation will continue normally. If you provide an invalid name, the search appliance will function normally, but you will not receive email notifications and you will not be able use the "Forgot Your Password?" feature to reset your password if you forget it. It is best to provide the search appliance with this information either during or shortly after installation. Google strongly recommends that you supply the name of an SMTP server.	Consult your network administrator
Logins and passwords needed for access to content locations	When content files are in directories or on devices that require logins and passwords for access, provide the logins and passwords required for access. The logins and passwords are entered on the Admin Console after you run the configuration wizard.	Consult your network administrator
The host name of the search appliance	A host name identifies the search appliance on the network. If you use a host name to access the search appliance, you have more flexibility in moving the physical location of the search appliance or changing the IP address of the search appliance.	Consult your network administrator
To use a dedicated administrative network interface card, an additional IP address and subnet mask	The IP address and subnet mask identify the dedicated network interface card	
Subnet mask to be used by IPMI	The subnet mask identifies the subnet on which IPMI is located.	
IP address of the default gateway to be used by IPMI	This IP address identifies the router used when network traffic is directed to any host outside the local subnet.	

## What Tasks Do I Need to Perform Before I Install?

The following tables describe required and optional tasks to perform before you install a search appliance. If you are indexing a content repository, refer to the connector documentation for more information on tasks to perform before installing the connector manager and a connector.

## Required Tasks

Before you install and configure the Google Search Appliance, perform the following required tasks.

Task	Description	For More Information
Ensure that the search appliance host name is configured in the network's DNS.	If you are using a host name as well as IP address to identify the Google Search Appliance on your network, the name must be defined in the network's DNS.	Consult your network administrator
Ensure that the search appliance can crawl content files located anywhere on the network.	Content on your network might be located on more than one subnet. The search appliance must be able to crawl content on all subnets where the content is located. If content is on subnets other than the subnet on which the search appliance is located, an incorrect router setup might block the crawl. This occurs when access control lists on routers block the search appliance or when routing tables on the routers do not allow the search appliance to reach other subnets.	Consult your network administrator
Mount the search appliance on a rack or otherwise place it in the desired location.	You can mount the search appliance on a rack in a data center or keep it on a flat surface in your office. If you keep it in your office, choose a location that has good sound isolation from work areas.	Consult your hardware administrator
Create an account with Google for Work Technical Support.	A Support account enables you to receive technical support.	The Welcome email you received when you purchased the search appliance or the Welcome letter enclosed in the box with the search appliance.
Ensure that a computer is available from which to run the configuration program and that a web browser is installed on the computer.	<p>You need a laptop or desktop computer that has physical proximity to the search appliance and can be attached to the search appliance with a cable. You can use a computer running Windows or a Macintosh. There are no restrictions on the browser used.</p> <p>If you have firewall software running on the computer, ensure that the firewall is configured so that you can open the network configuration wizard on the search appliance at <a href="http://192.168.255.1:1111/">http://192.168.255.1:1111/</a>.</p>	Consult your hardware administrator
Ensure that a backup electrical source is available to supply electricity to the search appliance if there is an electrical failure.	Electricity can be provided by an uninterruptible power supply (UPS) or a gas- or diesel-powered generator.	Consult your network administrator

Task	Description	For More Information
Decide whether the search appliance will autonegotiate network speed and duplex settings with the router or switch to which it is connected.	The search appliance can autonegotiate network speed and duplex settings.	Consult your network administrator
If you purchased more than one GB-9009, ensure that you pair the processing unit with the correct storage unit.	The processing unit label is printed with a service tag and a number in the format U1-xxxxxxxxxxxx. The storage unit that must be used with that processing unit has a label with its service tag and the same number that is on the processing unit.	

## Optional Tasks

The following table describes optional tasks you can complete before installing the search appliance.

Task	Description	Accomplished?
Configure proxy servers.	If the search appliance must access content through a proxy server, set up the proxies.	
If you plan to enable remote access using secure shell (SSH) for Google Support, arrange for network port 22 to be opened.	Google for Work Technical Support can use port 22, which is reserved for SSH remote login, for direct access to the search appliance, simplifying the support process.	

## Electrical and Other Technical Requirements

Your Google Search Appliance must be installed in a location meeting the temperature, electrical, refrigeration, and other requirements shown in the following tables. The configuration totals are valid at 110 AC input voltage.

The cooling fan on both the G100 and G500 operates at a higher speed than previous models. Therefore, you might notice more noise produced by the fan on these models.

The G100 and G500 complies with the following ASHRAE guidelines:

- Class A2
- Class A3 (fresh air configuration)
- Class A4 (fresh air configuration, up to 45°C for a maximum of 1% of its annual operating hours)

The following table shows requirements for the G100 and G500.

Requirement	Google Search Appliance G100 (T5)	Google Search Appliance G100 (T4)	Google Search Appliance G500 (U3)
Typical Thermal Dissipation	1228.4 BTU/hr 1975.8 BTU/hr maximum	836 BTU/hr	1504 BTU/hr
Operating Temperature Range	10° C to 35° C (50° F to 90° F) with a maximum temperature gradation of 10°C per hour. Note: For altitudes above 2950 feet, the maximum operating temperature is derated 1°F/550 ft.	10° C to 35° C (50° F to 90° F) with a maximum temperature gradation of 10°C per hour. Note: For altitudes above 2950 feet, the maximum operating temperature is derated 1°F/550 ft.	10° to 35°C (50° to 95°F) with a maximum temperature gradation of 10°C per hour. Note: For altitudes above 2950 feet, the maximum operating temperature is derated 1°F/550 ft.
Storage Temperature Range	-40° C to 65° C (-40° F to 149° F) with a maximum temperature gradation of 20° C per hour.	-40° C to 65° C (-40° F to 149° F) with a maximum temperature gradation of 20° C per hour.	-40° to 65°C (-40° to 149°F) with a maximum temperature gradation of 20°C per hour.
Operating Relative Humidity Range	10% to 80% (noncondensing), with a maximum dew point of 29°C (84.2°F).	20% to 80% (noncondensing), with a maximum humidity gradation of 10% per hour.	20% to 80% (noncondensing), with a maximum humidity gradation of 10% per hour.
Storage Relative Humidity Range	5% to 95% (noncondensing), with a maximum dew point of 33°C (91°F). Atmosphere must be non-condensing at all times.	5% to 95% (noncondensing), with a maximum humidity gradation of 10% per hour.	5% to 95% (noncondensing), with a maximum humidity gradation of 10% per hour.
Maximum System Power Consumption	579 W	356 W	585 W
Input Voltage (AC)	100-240 vAC, auto-ranging	100-240 vAC, auto-ranging	100-240 vAC, auto-ranging
Frequency	47~63 Hz	47~63 Hz	47~63 Hz
Total Current	2.4 Amps @ 110 vAC 1.6 Amps @ 220 vAC	3.3 Amps @ 110 vAC 1.6 Amps @ 220 vAC	5.3 Amps @ 110 vAC 2.6 Amps @ 220 vAC
Weight	65 pounds (29.5 kg) at maximum configuration	49.5 pounds (22.5 kg) at maximum configuration	61 pounds (23.30 Kg) at maximum configuration
Physical Dimensions	17.49" (W) x 26.92" (D) x 3.44" (H)	19" (W with Flange) or 17.5" (W without Flange) x 29.9" (D with handle) or 28.4" (D without handle) x 3.4" (H)	19" (W with Flange) or 17.5" (W without Flange) x 29.9" (D with handle) or 28.4" (D without handle) x 3.4" (H)
Industry Rack Height	2U	2U	2U
Airflow Rate	10.2 l/s 21.6 CFM	—	—
Sound Power Level	6.2 bels	—	—

The following table shows requirements for the GB-7007 and GB-9009.

<b>Requirement</b>	<b>Google Search Appliance GB-7007 (T3 Series)</b>	<b>Google Search Appliance GB-9009 Processing Unit (U2)</b>	<b>Google Search Appliance GB-9009 Storage Unit (U2)</b>
Typical Thermal Dissipation	1139.7 BTU/hr	1467.2 BTU/hr	790.9 BTU/hr
Operating Temperature Range	10° to 35°C (50° to 95°F) with a maximum temperature gradation of 10°C per hour Note: For altitudes above 2950 feet, the maximum operating temperature is derated 1°F/550 ft.	10° to 35°C (50° to 95°F) with a maximum temperature gradation of 10°C per hour Note: For altitudes above 2950 feet, the maximum operating temperature is derated 1°F/550 ft.	10° to 35°C (50° to 95°F) with a maximum temperature gradation of 10°C per hour Note: For altitudes above 2950 feet, the maximum operating temperature is de-rated 1°F/550 ft.
Storage Temperature Range	-40° to 65°C (-40° to 149°F) with a maximum temperature gradation of 20°C per hour	-40° to 65°C (-40° to 149°F) with a maximum temperature gradation of 20°C per hour	-40° to 65°C (-40° to 149°F) with a maximum temperature gradation of 20°C per hour
Operating Relative Humidity Range	20% to 80% (noncondensing), with a maximum humidity gradation of 10% per hour.	20% to 80% (noncondensing), with a maximum humidity gradation of 10% per hour.	20% to 80% (noncondensing), with a maximum humidity gradation of 10% per hour.
Storage Relative Humidity Range	5% to 95% (noncondensing), with a maximum humidity gradation of 10% per hour.	5% to 95% (noncondensing), with a maximum humidity gradation of 10% per hour.	5% to 95% (noncondensing), with a maximum humidity gradation of 10% per hour.
Maximum System Power Consumption	281 W	417 W	344.9 W
Input Voltage (AC)	90~264 vAC, auto-ranging	90~264 vAC, auto-ranging	90~264 vAC, auto-ranging
Frequency	47~63 Hz	47~63 Hz	47~63 Hz
Total Current	2.6 Amps @ 110 vAC 1.3 Amps @ 220 vAC	3.9 Amps @ 110 vAC 1.9 Amps @ 220 vAC	3.2 Amps @ 110 vAC 1.6 Amps @ 220 vAC
Weight	51.35 pounds (23.30 Kg) at maximum configuration including bezel and power supplies	52.15 pounds (23.65 Kg) at maximum configuration including bezel and power supplies	65.05 pounds (29.50 Kg) at maximum configuration including 1U cover and bezel
Physical Dimensions	17.44"(W) x 27.13"(D) x 3.38"(H): no rear handles or bezel	17.44"(W) x 27.13"(D) x 3.38"(H): no rear handles or bezel	17.62"(W) x 20.13"(D) x 5.06"(H): no rear handles or bezel
Industry Rack Height	2U	2U	3U