System Requirements

The following list describes the hardware, operating system, and network requirements for the Google Earth Enterprise system:

One server or two?
Google Fusion and Google Earth can be run on the same server or on different servers. In a lab environment or in a proof of concept system, Fusion and Earth Server should be run on a single server for ease of use and to reduce the amount of hardware needed. In a larger deployment or in a production environment, they should be run on separate servers. The Fusion server is often working at 100% CPU for extended periods, and this can impact the ability of the clients to receive the data they are requesting in a timely manner. Also, running on separate servers provides some fail-over capability.

If you have one modern, powerful server and one less capable server, use the most powerful server as the Fusion server and the less capable server as the Earth Server.

Hardware: For the Fusion and Earth servers, a Dell 2950 or equivalent is recommended. The Fusion server is more disk and CPU intensive than the Earth Server. The minimum recommended server configuration is:

- Dual Intel 3.0 GHz or Dual AMD Opteron 248 CPUs
- Minimum 1 GB RAM per core is required; 2 GB RAM per core is recommended.
- Minimum 500 GB of total hard disk storage (or more, based on the amount of data you have that you intend to process)
- DVD drive
- NVIDIA GeForce4 or higher graphics card - 64MB (for Google Earth Enterprise Fusion GUI only). This is for optimal performance of the Fusion GUI. Most server graphics cards will work well enough.

Operating System: One of the following operating systems needs to be installed. The 32-bit or the 64-bit version can be installed.

- SuSE Linux Enterprise Server (SLES 9 or 10) and the most current security patches available. You should have prior knowledge of how to install and upgrade SuSE for Linux. If installing SLES9, the latest update is SLES 9.3. You need to boot from the SLES 9.3 upgrade CD to achieve a successful installation. Do not install the base SLES 9 and then try to upgrade using the three 9.3 CD's.
- Red Hat Enterprise Linux RHELv.4 (both AS and ES) and the most current security patches available. RHEL 5 is supported at this time. The latest version of RHEL 5 is Release 5.2 that is supported.

Extra Packages: The following additional packages must be installed on the servers.

- Java Runtime Environment - full JDK from Sun (version 1.6 or later)
- OpenSSL (version 0.9.8e or the latest version available)

- Python (latest version available for your platform)
- X11 libraries, including Mesa (for disconnected production servers only)
- Google Earth Enterprise Server does not require an X server to be running or installed.

Network

The network needs to be functioning and stable prior to installing the Earth Enterprise software. Once the software is installed, it is difficult to change the hostnames of the servers. Settings such as hostname should be the full DNS name of your destination server. For example, myserver.mydomainname.com. This is necessary for the Google Earth Enterprise Fusion software to connect to the Google Earth Enterprise Fusion daemon. You can verify the hostname of your workstation by typing "hostname –f" at a shell prompt, and you can verify network connection by using the ping command to reach other hosts in the same network.

The network should be gigabit if possible, especially the connection between the Fusion server and the Earth server, and connections to network storage.

The users of the system need to be able to connect to the Google Earth server on port 80 (or 443 if SSL encryption is used). Firewalls must be configured to allow HTTP or HTTPS traffic from the clients to the Earth Server. No other ports or protocols are required for user access to the Earth Server.

The Fusion server must be able to communicate with the Earth Server over port 80. If possible, the firewall restrictions between the Earth Server and the Fusion server should be minimized. Fusion and Earth Servers can be connected across the enterprise network or locally with a crossover cable or switch. Most servers today have multiple network interfaces so a separate local network is easy to establish.

Storage

Note: Once your storage is configured, please copy all of the data that came on the external USB hard disks onto your servers prior to the Google engineer arriving. Copy the contents of the disks into the /gevol/src directory.

Three types of storage can be used with Google Earth. The Fusion server tends to be CPU limited, so fast and expensive storage is not as critical as it may be for other software applications. Also, the Fusion server does not need to be a high availability server. Once the Fusion server has processed a set of data and copied it to the Earth Server, Fusion can be shut down. The Google Earth server must be available to the users at all times. The advantages and disadvantages of different storage types are listed in the table below:

Storage Type	**Advantages**	**Disadvantages**
Network Attached Storage (NAS)	Multiple servers can read and write to the same	More overhead than SAN or direct attached

	storage. NAS is required for a Grid system. Reasonably priced and easy to administer.	May be slow if multiple servers are accessing the same storage
Direct Attached Storage	Very inexpensive Easy to administer	Only one server can access Does not scale as well as NAS or SAN
Storage Area Network (NAS)	Fast and reliable Works well for large datacenters	Expensive Complex

For small systems, direct attached storage is recommended for low expense and ease of use. For larger systems, NAS is preferred but, if a datacenter already uses a SAN, then SAN can also be used.

The Fusion server generally needs three to four times as much storage as the Google Earth server. The total amount of storage needed for the Fusion server and the Earth Server varies widely depending on how much data is being added to the globe. Imagery takes by far the largest amount of disk space. Terrain and vector data take much less. Please work with the Google Earth team to help estimate your disk space needs.

The following table has some estimates for storage space:

Data Published	**Fusion Server**	**Earth Server**
Google base data set – 15 meter USA plus 500 meter globe.	300 Gigs	100 Gigs
Google base plus 15 meter imagery for the globe	3.5 TB	1 TB
1-foot imagery of Washington DC area	70 GB additional	20 GB additional

When configuring a Linux workstation, Google recommends that you use the following mount point naming conventions:

- Single drive
  - Mount the single drive to slash (/). All data (/gevol/assets, /gevol/src, and /gevol/published_dbs) resides on that drive with the local path defined using the /gevol naming convention.

- Two drives – a small system drive and a larger data drive
  - Mount the small system drive to slash (/). Mount the larger data drive to /gevol/. Source and asset data volumes can then be defined as /gevol/assets and /gevol/src.

- Three drives – a small system drive and two larger data drives
    - Mount the small system drive to slash (/). Mount the first large data drive to /gevol/assets. Mount the second large data drive to /gevol/src.

- NAS or SAN
    - Install the operating system on the internal drive mounted to slash (/). Work with your installation engineer prior to arrival to decide how to partition and mount the network storage.

Grid Configuration

If you are installing a Fusion Grid, the following additional prerequisites apply.

1) Ensure that all of the computers in the grid have the same version of Fusion installed. System administration will be easier if all computers are running the same version of the Operating System, but this is not a requirement.

2) Configure each of the computers with Network Time Protocol (NTP) and synchronize all computers to the same time server.

3) Test that all of the computers are in the hosts table, DNS or other naming service so they can ping one another by hostname. All servers must agree on one another's hostnames and IP addresses.

4) The Fusion software runs as an unprivileged user named "gefusionuser". This user must have the same user ID and group membership and group ID across all of the computers in the grid, including the NFS server. The user must be included in a group called "gegroup".

If you are using Google Earth Enterprise over a network with at least two workstations, storage devices, and/or servers, you can use a centralized network authentication system, such as LDAP, NIS, or one of the many commercially available systems available. If you use a centralized network authentication system, you must add the following users to your authentication system's user list:

gefusionuser
geapacheuser
getomcatuser
gepguser

The primary group for all of these users is gegroup.

5) Host-based firewalls should be configured to allow unrestricted communications between the computers in the grid. The selinux system should be turned off if possible.

6) Fusion assigns jobs to servers alphabetically. All CPU's will be tasked on one server first before the next server's CPUs are tasked. Accordingly, faster servers should have their hostnames lower in the alphabet. Task rules can change this behavior.

7) The storage must be carefully considered before the fusion process begins in production. The recommended storage configuration begins with an NFS exported directory mounted at /gevol/assets (or the asset root of your choosing) on all computers, with additional volumes mounted at /khvol/src or other mount points. Any number of other volumes can be mounted as more storage is needed. Fusion volumes should represent entire system volumes (filesystems). The should not be nested if possible. Fusion doesn't enforce the non-nested rule, but the ability to manage free space may be restricted if nested volumes are defined. All computers should have the same volumes mounted in the same places. All of the computers in the grid need to see the same files that are created under /gevol/assets/.config, so the asset root cannot be a local filesystem.

8) A Linux server should NOT be used as the NFS server. It caches small files and breaks the NFS protocol. Use a commercial NFS solution. Do not use a SMB file system as it is not fully tested or supported.

9) For NFS installations, automount should be used if possible to mount the volumes to the grid computers. Static hard or soft mounts should not be used because they can hang the whole grid if one NFS server goes down. A soft mount seems like a good idea because it can tolerate a read operation failure and not kill the system. However, many 3rd party libraries do not check their read returns so a soft mount will be a silent failure causing garbage data. NFS options = rw,no_wdelay,sync,no_root_squash Need no root squash to allow permissions changing. Protocol tcp is by default now. Default read and write size are OK, or set them to 8192.