

The GroupStation in Defense Imaging

Alan Dare
SGI
2/8/00

Introduction

This document describes the benefits of applying the SGI™ GroupStation concept to defense imaging applications.

Current Architectures

With the increased power of commercial off-the-shelf (COTS) servers and workstations, many facilities have created image analysis centers based upon these products. Figure 1 depicts the common application of this COTS solution. For convenience this networked topology will be referred to as the digital image network (DIN). The image server typically provides services such as compression, decompression, reduced resolution data set (RRDS) generation, image storage, cataloging, mass storage administration, and serving images to clients. The image server may provide downloading of full frame imagery or serve subsections of

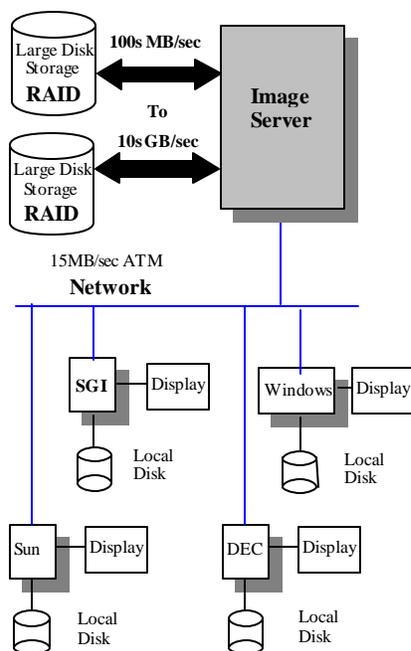


Figure 1

images or tiles on demand. Examples of these types of image services are Lockheed-Martin's Display Broker for tiling services and GDE's PPS/OPS for full frame imagery serving.

Connected to the image server are high-performance RAID subsystems with bandwidths from hundreds of megabytes per second to tens of gigabytes per second. The amount of storage varies, but online storage for several days to several weeks is common. Not shown in Figure 1 is the mass storage subsystem that usually accompanies the system. Mass storage devices provide the near-line storage for imagery archived from the RAID disk subsystems.

Various workstations are connected to the image server via an FDDI or ATM network and use standard network protocols such as TCP/IP and NFS™. The workstations have local storage subsystems to hold their tiles or full-frame imagery for softcopy exploitation, mapping, charting, geodesy, and reporting. Most workstations require large local disk storage to contain their product for analyses, exploitation, and reporting purposes. Along with this, some form of graphics acceleration

is usually present and coupled with a memory footprint between 512MB and 2GB. Thus, most of these workstations connected to the image server are generous in size. The DIN architecture is COTS based and provides flexibility and cost savings over previous proprietary systems. However, it has limitations.

Architecture Limitations

The DIN architecture described in Figure 1 has several limitations that affect performance, quality, reliability, and cost. The specific limitations of this system include:

- Download time for large imagery to the desktop
- Limited performance of low-cost graphics subsystems in workstations
- Quality limitations of low-cost graphics subsystems in workstations
- Duplication of resources on each workstation
- Problems in maintainability and reliability
- Dedicated resources

Although the bandwidth of modern COTS servers and disk subsystems is high, the networks connecting to these systems can provide only a fraction of this bandwidth. Thus, trying to respond to a potential critical situation by analysis of a broad area search (BAS) becomes very difficult. The network becomes a bottleneck. This situation is exacerbated if several image analysts (IA) are accessing imagery over the same network. Consider the problem of downloading a 70GB mosaic image to a workstation for analysis. At an ATM peak rate of 155Mb per second, it would take one hour to transfer the image at full bandwidth on a dedicated network connection. Rarely would an IA be able to obtain a dedicated network for a single transfer in this architecture. Networks based upon FDDI or 100Base-T exhibit worse performance. Although individual analysis centers may not routinely access mosaics of this size, the problem remains the same for centers trying to serve large amounts of tiles. As the number of IAs increases, their requests for tiles increase. Thus, the required bandwidth increases. The result is network saturation and a backlog of imagery. A significant portion of an IA's time is spent waiting for imagery downloads.

After an image has been downloaded to an individual workstation, analysis begins. Due to the high cost of placing individual workstations on each IA's desk, a tradeoff is made among cost, quality, and performance. This tradeoff is apparent when performing electronic light table (ELT) functions. Roaming, zooming, rotating, and comparison can become sluggish at best. This is caused by lower performance graphics subsystems that must be duplicated in each workstation. The quality of these graphics subsystems is also affected. The large proprietary systems used in the past for softcopy exploitation provide high-speed roaming and zooming with high quality. These systems provide bicubic interpolation and large convolution matrices. This of course comes at increased cost per IA workstation. The lower cost workstation provides bilinear interpolation and 3x3 convolutions. Thus, workstation softcopy performance and image quality are worse.

Other limitations include duplication of resources. Each workstation is equipped with a substantial amount of disk storage. The imagery is already stored on a disk subsystem on the server. However, a copy is sent to the workstation and placed on its disk for analysis. To reduce costs it would be advantageous to use a network protocol such as NFS to eliminate the need for large local disks on each workstation. Here again network performance, with large imagery, becomes a bottleneck for use in a production environment. Furthermore, the duplicated disk drives become cumbersome to use as a shared resource. If an IA is out for a month, how do the other users of the system utilize the quiescent disk asset? It becomes a difficult proposition and a potential waste of resources. Another resource that is dedicated and difficult to distribute as a system resource is the CPU power within each workstation. If the center is not staffed with IAs on a 24-hour cycle, the power of the CPUs in each workstation becomes an untapped resource. It would be useful to apply this untapped resource to provide decompression and RRDS generation during unused periods. Contemporary models using distributed resources have improved with applications of MPI and other libraries. However, distributed models are difficult to implement and rarely come close to achieving results equivalent to a pooled resource. Centers using a distributed model on a full-time basis to make use of the most potential of all resources run the risk of poor response time on active systems. If an IA is using a system and an RRDS is started on his system, sluggish performance will frustrate the user and reduce efficiency.

Management of this system is also difficult. The network must be set up, tuned, and watched on a constant basis. System administration of a large number of workstations on this network must be maintained.

Upgrading the operating systems and utilities of the workstations in a large site can become a harrowing experience. Clearly this architecture has many advantages over proprietary systems. However, as previously discussed, it does have limitations.

The GroupStation Architecture

An alternative to the DIN architecture described in Figure 1 is the GroupStation concept shown in Figure 2. The GroupStation is built upon the scalable architecture of the SGI™ Origin™ and Silicon Graphics® Onyx2® family of computer systems. Origin servers are expandable from two to 128 processors. An Origin server with a graphics subsystem is called an Onyx2 system. With the exception of the graphics subsystem, the systems are identical. The GroupStation concept described in its simplest terms includes high-end graphics subsystems directly in the image server. In Figure 2 a network of workstations as described in the DIN model is encapsulated within a dotted line. The GroupStation architecture doesn't preclude a DIN, and existing DINs can be upgraded to a GroupStation without sacrificing previous workstation investments. However, the power behind a GroupStation is the advantage that is obtained by connecting powerful graphics subsystems directly to the image server. The GroupStation architecture provides the following benefits over the previous DSN architecture:

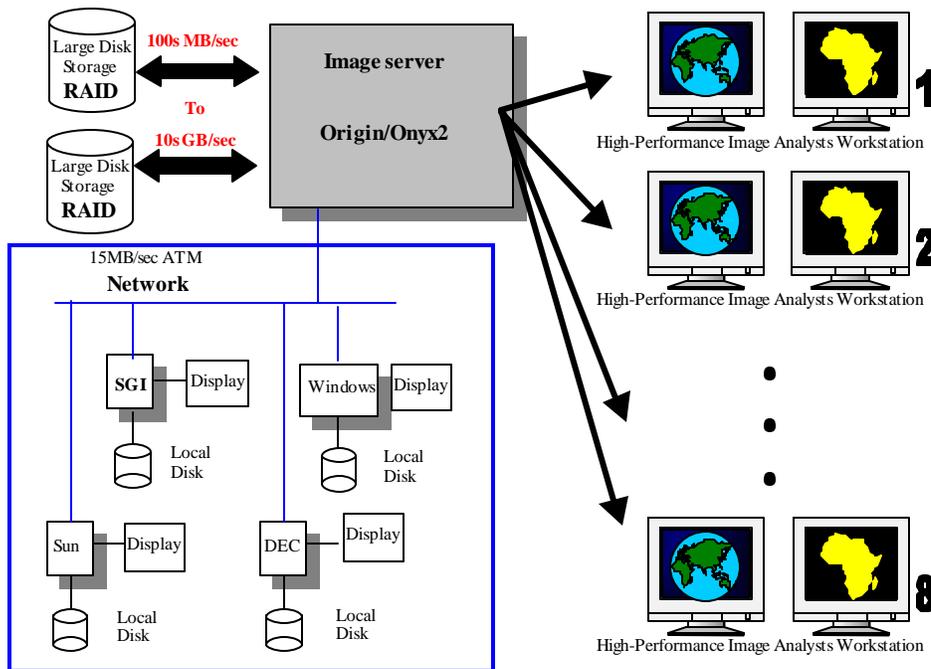


Figure 2

- No image downloads
- Scalable server
- Shared components
- Highest quality imagery
- Highest resolution
- High-performance graphics
- Stereo in a window
- High-performance 3D visualization
- Programmable video output
- Single system to maintain
- Elimination of complicated workflow
- COTS

GroupStation Advantages

The addition of Silicon Graphics InfiniteReality® graphics subsystems directly into the server provides increased performance that cannot be matched with the DSN model. Specifically, no image downloads are required for viewing imagery on a GroupStation terminal. There is no need to download data. The GroupStation terminal is part of the image server. Using a GroupStation terminal, SGI has demonstrated accessing and displaying a 67GB image within two seconds of initiating an ELT program. On DIN architecture that same image would require a minimum of an hour transfer time on a dedicated ATM network before an IA would have access to it.

The InfiniteReality graphics subsystems provide other benefits. The InfiniteReality graphics system is one of the highest performance and highest quality COTS graphics systems in the world. InfiniteReality graphics can display up to 896 million pixels per second. This provides smooth roaming and zooming over large images. The InfiniteReality graphics subsystem can maintain a pixel depth of 12 bits RGBA or 16 bits luminance alpha throughout the graphics subsystem. The InfiniteReality graphics subsystem provides 7x7 convolutions and bicubic resampling in hardware for sharper images in warping and zooming. Since this is in hardware, performance is not sacrificed for quality. This performance and this quality were previously obtainable only by proprietary systems, and they are not obtainable with other COTS workstations. The InfiniteReality graphics system is the first truly COTS system to provide IDEX performance, quality, and resolution.

This same graphics system is used for 3D-scene generation and has the ability to generate 11 million polygons per second. This provides smooth visualization of imagery draped over digital elevation models (DEM). Thus, the same GroupStation terminal can be employed for softcopy exploitation, mapping, charting, geodesy, and mission rehearsal. For example, the InfiniteReality graphics subsystems are used by the DoD's TopScene and PowerScene for mission rehearsal. Other important aspects applicable to defense imaging are the ability to connect two HDTV resolution monitors to each graphics subsystem. This provides an IA a larger viewing area than other graphics systems. The video output is also programmable. This allows virtually any monitor to be connected to the InfiniteReality subsystem. Stereo viewing in a single window is also supported. Other workstations that support stereo analysis force the entire screen to switch into stereo. This creates blurring of user interfaces that are not contained in the 3D-viewing window. InfiniteReality supports stereo in a single window. This allows your 3D window to be in stereo and objects outside this window to be viewed correctly as nonstereo items.

The Origin/Onyx2 architecture containing the InfiniteReality graphics subsystems is one of the most powerful and scalable computer systems available. These systems can grow CPU power from two to 128 processors, memory from 512MB to 256GB, I/O bandwidth from 5.1GB per second to 82GB per second, memory bandwidths from 2.8GB per second to 44.5GB per second, and online storage to 74TB and one to eight InfiniteReality graphics subsystems. This expandability doesn't require box swaps, long down time, or multiple instances of the operating system. The server can simply be expanded as needed. The GroupStation concept saves duplication of resources by combining the resources in the server. Thus, the CPU power and memory reside on the server under a single instance of the operating systems. This makes the CPU and memory utilization for the center much higher because it is easier to access these resources. The resources are not spread out across a network of workstations. It saves costly duplication of disk resources on each workstation. The imagery resides on the server and doesn't need to be copied. A single server versus a host of workstations connected to a server becomes easier to administer and maintain. Users wishing to deploy their GroupStation terminals away from the server can purchase COTS fiber-optic keyboard, mouse, and display extenders. This allows the GroupStation terminal to be up to a mile away from the actual server.

As described previously, an individual server can have up to eight InfiniteReality graphics subsystems. For sites needing more than eight individual stations there are options. Figure 3 shows one of the configurations containing 16 GroupStation terminals. This system architecture is composed of two servers connected to dual attach RAID disk subsystems. This dual-server configuration provides graphics expandability and fault tolerance. In configurations with eight or fewer GroupStation terminals, if significant fault tolerance is required it may be desirable to configure as a multiple server architecture. The DIN architecture has single points of failure in the server and the network. The GroupStation described in Figure 2 has a single point of

failure in the server. The GroupStation described in Figure 3 has no single point of failure. In all three cases the RAID disk system was ignored as a single point of failure. RAID systems are redundant by design, and newer systems contain redundant power supplies, fans, and controllers. Redundant servers can be employed without duplicating or equally splitting resources between them. The secondary server could be a smaller server containing enough CPU and memory to allow functionality at a reduced capacity if the primary server fails. When the primary server is repaired, full throughput is reestablished without a complete shutdown of the facility.

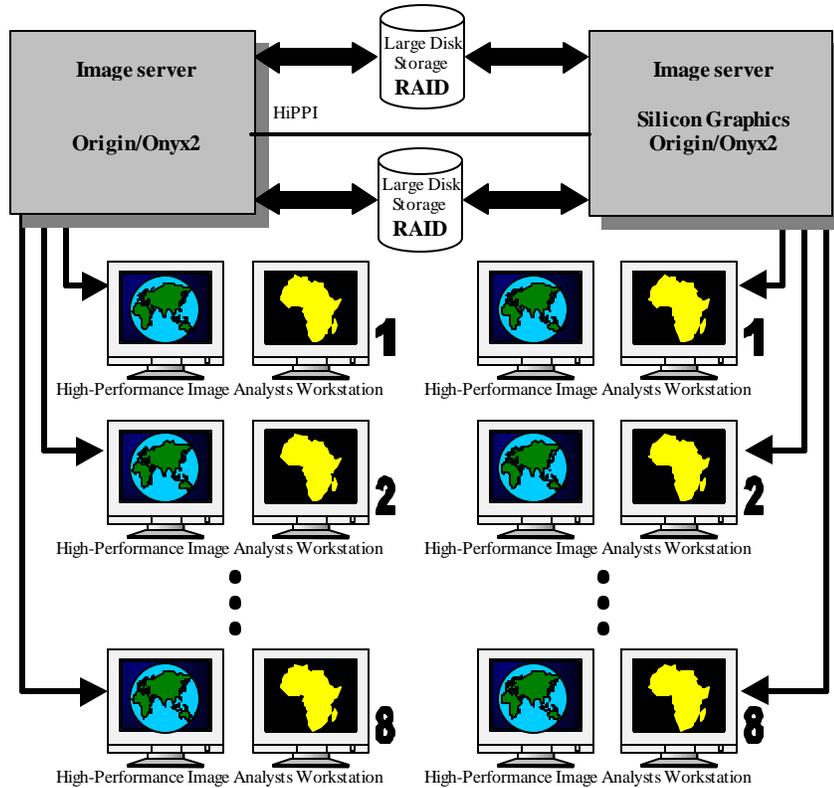


Figure 3

The final analysis of the GroupStation depends on cost. With the significant increase in performance and quality that InfiniteReality graphics subsystems provide, one might believe the GroupStation is significantly higher in cost than creating a DIN architecture. However, by elimination of redundant workstation components and economies of scale, system costs for the DIN architecture compared to those of the GroupStation are roughly equivalent.

Silicon Graphics has invested in research and development to deliver truly scalable servers and high performance graphics systems. The commercial market demand for high-performance and high-reliability systems supported the development of the state-of-the-art Origin and Onyx2 systems. These COTS systems provide the DoD the ability to replace proprietary systems with flexible, cost-effective solutions at equal or greater performance. The GroupStation in defense imaging leverages Silicon Graphics investments in commercial accounts to provide the most cost-effective, high-performance COTS system available.

Copyright © 2000 Silicon Graphics, Inc. All rights reserved.