Metropolitan Area Sun Ray™ Services

By Lars Persson–Sun Professional Services

Sun BluePrints™ OnLine - May 2002

`http://www.sun.com/blueprints`

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95045 USA
650 960-1300

Part No. 816-4818-10
Revision 01 May 2002

Metropolitan Area Sun Ray Services (MASS)

The IT costs for companies, organizations, and government bodies is a modern plague. Costs frequently grow out of control and are hard to calculate. IT operations cost money, but not all costs are visible. Computer problems, primarily at the user level, generate productivity problems on a massive scale, and with productivity loss, your income decreases as expenses increase. Deploying Sun Ray™ appliances in your network environment is a powerful weapon against such hidden costs, and because of the modest pricing, the cost of deployment is reasonable.

The desktop computational device can be obsoleted by deploying Sun Ray appliances without loss of computational power to the end user. Metropolitan Area Sun Ray Services (MASS) is a method to get the good things that Sun Rays bring forth, out from the small LAN and into the wide open spaces of metropolitan area networks (MANs, see “What is a Metropolitan Area Network?” on page 11). This technology also opens a market for entire new types of operations, as described in this article.

Sun Ray appliances have a great potential to be a viable alternative to the common workstation. By using true broadband technologies, deploying Sun Ray appliances on a massive scale in MAN environments is possible. The advantage is that centralized data centers can be run by a few specialists. This method eliminates production losses and incompatibility problems at the user level, reduces overall costs of operation, increases data integrity, and practically eliminates the risk of virus attacks. The MAN could be a CityNet, and the end users could be anybody from the common citizen (surfing the Internet, retrieving mail, or home office use), to small firms, franchising companies, corporations, government bodies, health care, or other organizations.

Today, all IT operations consist of various servers, but on the user side, we are still bound to the desktop computer. The purpose of this paper is to describe a viable thin client alternative, the functions of the components that form it, and how to massively deploy such clients over broadband technology in MAN-type networks.

This article is intended for people working with IT architecture or IT strategies with the intent to rationalize IT operations, especially if these operations are spread over a large geographic area.

The Vision of Sun Ray Appliances in a MAN-type Network

The immediate possibilities of MASS can be appreciated at once. MASS can be viewed as an innovative method of deploying Sun Ray appliances in various quantities, at different locations, and still connect them to a centralized data center in a robust way. This concept can be sufficient for the individual company or organization. MASS is scalable too. Instead of using a Sun Ray backnet (a private Sun Ray network) and its Sun Ray server, you could consider a Sun Ray data center hosting a large number of Sun Ray domains, each with its own backnet.

MASS could grow to become a valuable asset for large MANs like a CityNet. With MASS, the owners of a CityNet can offer its users maintenance-free access to the Internet while providing an appliance, complete with surfing capabilities, mail, and an excellent office package—the StarOffice™ software.

Because the Sun Ray servers and application servers are situated in a data center, user data can be backed up centrally, thus offering data protection to the individual user. Sun Ray appliances deployed to ordinary users can run processes in user mode on the servers. This feature makes them practically immune to the virus plague that constantly rages the Internet these days.

Taking one step further, it is possible to facilitate easy and inexpensive outsourcing for everything from small firms to franchising companies, and large corporations for mail, web access, office, and portal carrying applications.

Although Sun Ray appliances use a special private network, it is possible to configure the broadband/Ethernet switch (the proxy switch for which the Sun Ray appliance is connected) so that the appliance has access to other networks. This configuration enables standard workstations to coexist with Sun Ray appliances on the same switch (using port VLANs), should the need arise. Proxy switches could also be deployed in residential areas using one or more fibers and common narrowband techniques. Users could have access to 10 Mbps or 100 Mbps connections even at home. The Sun Ray appliance could function as an internet "surf board", office carrier, and mail tool. The field is open for net-bound applications on portals to be instantly accessible by Sun Ray users without locking out workstations or other types of appliances or computers. The maintenance-free surf board could be an option in an Internet subscription.

Other uses for MASS could be in educational services. Instead of maintaining hundreds of workstations, plug-and-play Sun Ray appliances could be set up. If a Sun Ray appliance malfunctions, the only thing that has to be done is to replace it with another unit. No reinstallation or upgrading is needed, and possible licensing costs could be incorporated in the fee for access to the MASS network. Servers could either be outsourced or be placed at a location within the educational service's premises. Sun Ray appliances could be spread to schools over the MAN, regardless of distance within the MAN. Hospitals, government bodies, and so on are also potential MASS customers.

The main advantage for MASS users in a MAN area is that no unexpected costs will occur. Hardware upgrades every six months are no longer necessary, and companies and firms do not have to keep an expensive IT department anymore. MASS is easy to establish. It can expand as needed, and it is easy to maintain.

Establishing MASS is a "win-win" situation for everybody involved, and it will encourage growth and development of new companies and ventures that do not exist today. This article describes how all of this can be achieved.

What is a Sun Ray Thin Client?

The general idea of a *thin client* is to provide zero administration at the user level. Sun Ray appliances do not have a local operating system, and they do not need to be preconfigured to function. The work involved when installing a Sun Ray appliance is limited to connecting it to its keyboard, mouse, screen, and the network. After the Sun Ray appliance is powered up, it is instantly accessible to a user.

Sun Ray appliances have a very small footprint, and they are extremely rugged and dependable. Because they do not need forced cooling, Sun Ray appliances lack fans, making them silent. Also, there are several versions of Sun Ray appliances. Some are built into flat screens, while others can use modern multi-sync VGA-type video display units.

Sun Ray appliances are centrally administered by very powerful Sun Ray control software, making them suitable as a highly cost-efficient alternative to the traditional workstation.

An extremely useful feature of the Sun Ray appliance is its smart card function. When using smart cards, the session displayed on the appliance is tied to the identification number of the smart card, not to the appliance itself. This feature makes it possible to move the current session to another Sun Ray unit simply by removing the card and inserting it into another Sun Ray appliance. If a smart card is not available, the Sun Ray appliance can be used anyway, but the session will be tied to the physical Sun Ray unit. The smart card function makes the Sun Ray appliance both session persistent and session independent at the same time. In the Sun Ray Server software version 1.3, there is even a non-smart-card mobility feature, enabling user sessions to be hot-desked without a smart card.

How a Sun Ray Appliance Works

A Sun Ray appliance can be described as a a terminal device with no local computing environment, and requires a connection to a Sun Ray server. Sun Ray appliances connect to their servers through a dedicated Sun Ray network using standard TCP/IP. In this article, we refer to that network as a *backnet*.

Sun Ray appliances communicate using UDP and need a high quality network. Thus, Sun Ray networks must be practically error free (less than 0.1% packet loss), cannot be overloaded, and must have low latency and sufficient bandwidth, or the Sun Ray appliance might lose its connection to the server and stop working.

A Sun Ray server can, in principle, be any reasonably modern Sun computer with sufficient capacity. When choosing and sizing servers, we recommend following the directions given by the Sun Ray development team. However, modest resources can be used to form a Sun Ray server if you use other servers as application carriers.

At boot time, a Sun Ray appliance establishes a link to the Ethernet, and then uses DHCP to get its network and Sun Ray vendor specific parameters. When the Sun Ray appliance is up and connected to its Sun Ray Server, the server perceives it as yet another video card (and additional hardware such as a mouse, keyboard and so forth).

In a minimal environment, the Sun Ray server is also the application carrier. Sun Ray users log in to run their programs, and displays them on their local screen. In larger deployments, the Sun Ray server has a connection to a *public network*, or *frontnet,* such as a corporate LAN segment. Sun Ray users can have their application servers elsewhere (as usual in a true client/server environment).

When traversing the frontnet, the Sun Ray appliance is seen as originating from the IP address of the Sun Ray servers frontnet interface and not from the real IP address of the Sun Ray appliance (no routing occurs between the backnet and frontnet). The Sun Ray appliance is identified through the display number issued by the Sun Ray server; thus the display data can be sent to the correct device.

Because Sun Ray appliances are not running anything but microcode locally, the processes they open on the servers are not dependent on the state of the Sun Ray device. You can use several methods to determine what process belongs to which Sun Ray appliance, but these methods are beyond the scope of this article.

Sun Ray appliances have smart card readers. Using them is optional; users can usually log in on an application server without using smart cards. If smart cards are used, the display number is tied to the smart card identification code. Removing the smart card will free that particular Sun Ray appliance for others to use. Inserting the card again will display the open session accosted with the card. Removing the card and inserting it in another Sun Ray appliance causes the session to be displayed there. The purpose of this feature is twofold. You can remove the card without logging out, and your session is available, without logging in, the next time you need to use the application server. The other benefit is that your session is not bound to a geographical location, but instead it can be moved around (flexible office, meetings, and so forth).

Until now, a limitation of the Sun Ray technology has been the modest scalability and inability to effectively function over long distances while retaining a high degree of redundancy. Naturally, there is a demand for these functions, and by using broadband networking techniques, these demands can be met today.

The Sun Ray Environment Backnet Dilemma

Sun Ray appliances place, as previously mentioned, high demands on the quality of the backnet. These demands can be summarized as follows:

- Sufficient bandwidth (application-dependent)
- Low latency (less than 50 ms round-trip between Sun Ray server and DTU)
- Ensured in-order packet delivery
- Near error-free networks (less than 0.1% packet drop)

Note – Errors can occur, but they must be insignificant.

Errors in an Ethernet network are usually caused by over utilization, misconfiguration, or physical defects on a network segment, causing packets to be dropped.

Over utilization causes collisions. The collision detect (CD) function of Ethernet causes communication to cease for a certain time frame (9 ms + *dt*). Collisions destroy, fragment, or misalign packets sent over the wire at the moment of collision. The CD function causes latency. In severely affected segments, communication becomes impossible (collision storms). Collisions are avoided by using good quality Ethernet switches that can operate in duplex mode, and are designed and implemented in such a way that they can handle the load caused by the network traffic.

Typical misconfigurations causing errors on the network are either ARP/RARP storms or duplicate IP addresses. Both can be determined by using network analysis tools.

Physical defects might be introduced by terminating equipment (hosts and other Ethernet-connected devices) or transfer equipment (switches and so on), but most likely, defects stem from faulty cabling or connectors. The most common culprits are faulty patch cables. Physical errors are found by containment and measurement of the suspected equipment.

Sufficient Bandwidth

The trade-off is ensuring that there is sufficient bandwidth most of the time, or guaranteeing sufficient bandwidth all of the time. The latter is done by ensuring that the sum of all possible traffic will never exceed the capacity of the network. This is, however, very costly. Usually devices never operate at their potential peak load at once, so the capacity of a network designed this way is usually grossly over-dimensioned.

You can ensure adequate bandwidth by testing and calculating the actual peak loads, average loads, and minimum loads to deduce the capacity needed for the network components. This method can lead to a theoretical oversubscription of the network. It is common to see oversubscription numbers in the order of 2:1, 3:1 and 4:1. It is not uncommon to see numbers in the order of 10:1.

Note – In an oversubscribed network, the theoretical need for bandwidth (the sum of the maximum capacity of all network attached devices) supersedes the actual capacity of the network. It is common and quite safe to oversubscribe networks if proper investigation has been carried out. Oversubscription is done to satisfy the actual need for bandwidth, while keeping equipment costs down.

Example:

When dealing with Sun Ray appliances, we measured the actual need for bandwidth and found that usually a Sun Ray appliance has very low bandwidth demands, but can peak up to approximately 25 Mbps when running streaming video display programs. Based on this, we could say (bandwidth-wise) that if Sun Ray appliances are attached through 100 Mbps networks, and are all running heavy streaming video programs all the time, it would be safe to have an oversubscription rate of 3:1. This rate would give us a guarantee that the Sun Ray appliances would work at all times. The other extreme occurs when users do not run streaming video at all. In such cases, the oversubscription rate could be 10:1 on a 10 Mbps network. A single Sun Ray appliance often uses less than 10 Mbps, and most of the time the appliance transmits user-generated traffic. These examples are not a recommendation, but examples of how you could reason. Measurements and risk assessments give the correct oversubscription rate for a particular configuration.

Low Latency

Latency can crudely be described as equipment handling time. In Sun Ray environments, low latency is a priority because the performance of the Sun Ray appliance, in effect, is dependent on the network performance. Latency is caused by collisions, as mentioned previously, but also inter-LAN links can cause latency, as can VLAN technologies.

When deploying Sun Ray appliances in a network, the “golden rule” is to make the architecture as uncomplicated as possible. This reasoning is seemingly the opposite of the mere thought of deploying Sun Ray appliances in a complex environment such as a MAN or a CityNet. We will show, later, that it is possible to combine MANs and low latency.

In-Order Packet Delivery

Because Sun Ray appliances use UDP as an information carrier, there is no way for the network to determine if the packets arrive in the correct order. This determination has to be handled in the application layer. When it comes to Sun Ray appliances, it is a design choice to not wait, but instead, consider out-of-order packets as lost packets. This design choice causes a problem when it comes to load sharing over multiple network trunks, as shown in FIGURE 1.

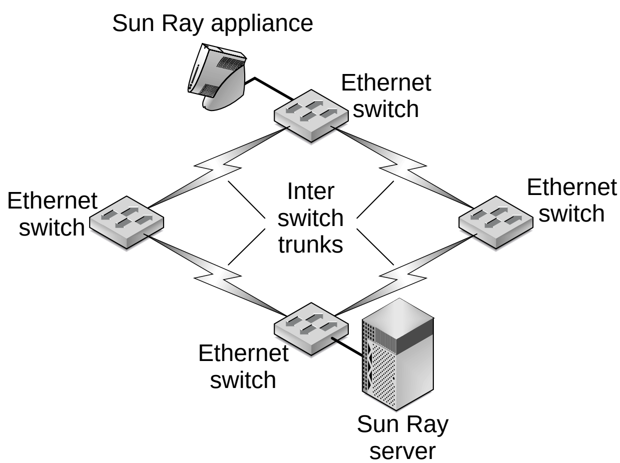

FIGURE 1 Load Sharing Over Multiple Network Trunks

Sun Ray appliances can tolerate occasional out-of-order packets, but if this happens too frequently, Sun Ray appliances might not work in that environment. In this theoretical model, Ethernet switches are set up as a load-sharing set. One packet can travel over the one interswitch trunk leg, and the next packet over the other one, depending on the load situation for the particular devices involved.

The latency for the components depends on a fixed constant for the hardware, plus a variable based on the load of the device in question. The two legs are not dependent on each other. A packet leaving the server routed to one leg can arrive later than a packet leaving the server afterwards routed over the other leg. The Sun Ray appliance has limited means of determining if this happens.

More About the Sun Ray LAN Backnet

As mentioned previously, a Sun Ray backnet must be almost error-free, have sufficient capacity, low latency, and ensure that packets arrive in the same order as they were sent. In principle, this criteria rule out shared media (such as coaxial cables, hubs and fan-outs), even though they might be supported.

To ensure that the demands are met, we recommend using modern Ethernet switches with duplex capabilities. We also recommend connecting only one Sun Ray appliance per port. The sum of all traffic generated by the clients should not exceed the capacity of the network interface of the Sun Ray server (see“Sufficient Bandwidth” on page 6). Today, many switches use VLAN technology, and it is important for the continuity of this article that you are aware of the functions and features of VLANs.

VLANs From a Network Point of View

A VLAN-capable switch can determine, based on certain configuration criteria, where traffic through it belongs. When the criteria are met, traffic is sent out through one or more ports. If the criteria are not met, the switch prevents the traffic from exiting the ports.

The three most common types of VLANs are:

- Per-protocol based VLANs
- Tagged VLANs (IEEE 801.2Q)
- Per-port VLANs

In the first case, you can set up a switch to separate different types of networking protocols (such as TCP/IP, Appletalk, XMS and so forth). Because we deal solely with TCP/IP, we do not discuss the other protocols.

In the tagging mode, the switch can inject (and discover) LAN identifiers in the header of an Ethernet packet, and route the packet to the correct location. The typical use for tagged VLANs is to be able to use an interswitch trunk to transfer several VLANs over the same line. Because the switch must open every packet going through a tagged port, latency could increase. However, high performance switches can easily handle this load, even under Gigabit wire speed. For large networks where Sun Ray appliances are just a part of the entire network traffic, and that span a large number of switches, the compound complexity could be a problem.

The last type of VLANs are the most important for this paper. Per-port VLANs are a way to partition a switch into separate LANs. These LANs are logically separated, but physically belong to the same switch. With per-port VLANs, the physical and the logical architecture differ.

FIGURE 2 shows two LANs (A and B) using the same switch. Traffic does not flow from A to B inside the switch. It is separated as though A and B were two separate LANs. The example could represent a minor Sun Ray deployment where A could be the frontnet and B the backnet. A Sun Ray server should, in this case, have the frontnet interface connected to an A port (for example port 1) and the backnet interface connected to a B port (such as port 2).

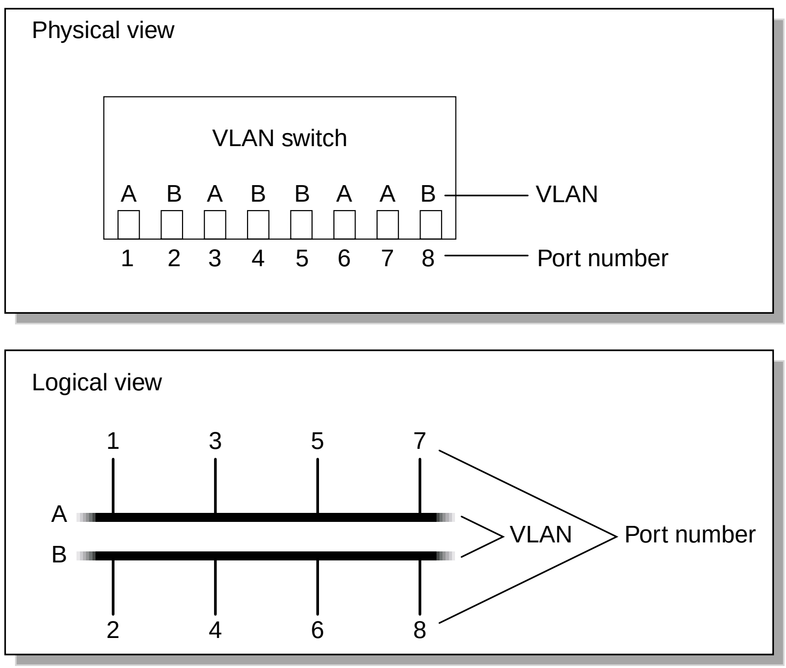

FIGURE 2 Physical and Logical Views of a VLAN Configuration on an Ethernet Switch

Long Distance Backnets

The distance between a Sun Ray appliance and its server is the sum of the segment lengths between them. For twisted pair, such a segment length can be approximately 100 meters. If the Sun Ray appliances must be at long distances, other media must be used, or there will be many needless interswitch trunk hops that increase latency. In cases such as these, fiber technologies are used. Two types of fibers are used for data communications, multimode fibers for ranges up to 2 km, and single-mode fibers for ranges up to (and, in some cases, well exceeding) 60 km. The fiber is not bound to run TCP/IP as such, but rather, it is a general media.

A fiber segment has no significant added latency, even when run over great distances, because it operates with laser generated light. Light in glass is sufficiently fast to enable us to disregard latency generated by the distances for which they are designed. You can safely assume that the distance-related latency for 60 km is zero.

Summary of Sun Ray Networking

Great care must be taken when designing Sun Ray backnets to ensure that they have very few errors and low latency, deliver packets in the correct order, and have sufficient capacity. Per-port VLANs can be used to separate a Sun Ray backnet from other networks, and with care, tagged VLANs could be used to tie switches together within a building or compound. If you fulfill the design criteria, Sun Ray appliances can be deployed over longer distances using fiber technology.

Broadband Versus Narrowband

We frequently regard a broadband network as a “high speed connection” to the Internet run over Ethernet network interface cards (NICs). This is not a correct point of view. We must differentiate between broadband and narrowband networks in that broadband networks can be used for many different traffic characteristics, while narrowband networks are used for one. Broadband and narrowband networks have nothing to do with the capacity of asymmetrical digital subscribe line (ADSL) or cable modems. These are access points to a broadband network.

Narrowband technologies are Ethernet, Token Ring, FDDI, and certain types of WAN links. Narrowband technologies can be adapted to run over broadband technologies, but it is not possible to run broadband technologies over narrowband. Broadband technologies, on the other hand, are not limited to a single link protocol (like Ethernet) or to data communication only. Broadband networks use frame, or frame trains, and special methods to encapsulate and control the information sent through it. The information does not have to be computer related at all. For example, voice -over-IP is not a broadband technology because the audio is digitized and sent as IP packets, mostly over Ethernet, whereas on broadband, voice can be sent as a separate channel without encapsulating it into a computer communication protocol.

The three major types of broadband communication are:

- Constant bit rate (time critical)—typically video streams
- Available bit rate (in-order packet delivery)—typically phone systems
- Unspecified bit rate (error-free)—typically data communications

Broadband networks can be adapted to form the physical layer for LAN technologies (broadband ISDN or ATM), or to run subsets of TCP/IP directly (such as PPP over SDH).

The three major groups of broadband technologies are:

- Synchronous optical network (SONET, North America)
- Synchronous digital hierarchy (SDH, Europe)
- Broadband ISDN (ATM)

ATM can function by itself or run over SONET or SDH. The latter is more common.

What is a Metropolitan Area Network?

A MAN resembles a large corporate network where different LANs are connected through campus-type transfer networks into a core-type backbone. The difference is that the typical corporate network is often a narrowband network, while a MAN, often, at least partly, relies on broadband technology.

A MAN also spans a geographical area and can, from a LAN point of view, traverse long distances. In a corporate network, security is implemented at the edges, policies at the transfer nets, and the core is unrestricted.

In a MAN, the edges are often owned by subscribers to the MAN services, and are not under the MAN owner’s direct control; whereas the core and transfer nets are under a very high degree of control.

MANs are often owned by telcos, major ISPs or government agencies, even though individual companies or corporations can build their own MAN-type networks. A MAN owner has an advantage if the MAN can deliver various types of services, not only data communications. To effectively use the available MAN bandwidth, MAN owners want to eliminate overhead as much as possible because bandwidth is, in the end, what generates a profit.

Principles for a Sun Ray Environment Over Broadband Networks

First, some seemingly prohibiting facts:

- A Sun Ray appliance cannot connect directly to a broadband network.
- A Sun Ray appliance cannot be connected to a pair of single mode fibers for long distance communication.
- A Sun Ray appliance is a LAN-bound narrowband device, connecting only to 10 Mbps and 100 Mbps Ethernet networks over RJ45 connectors (twisted pair).

The solution is to connect the Sun Ray appliance to a broadband-capable switch. Such a switch is frequently called an *edge device* or *proxy switch*. The proxy switch could be equipped with one (or several) fiber interfaces, enabling long distance connections. The Sun Ray appliance can "talk" Ethernet to the proxy switch; the proxy switch then "talks" TCP/IP through the broadband network.

In the computer center, the Sun Ray server can be equipped with one or more broadband interface cards (BICs). Because we have control over the broadband network characteristics, we can set up channels through the network, connecting the proxy switch and the Sun Ray server. The proxy switch is configured with per-port VLANs, and the channels to the broadband interface are included in the VLAN group.

The Sun Ray appliance sees all traffic going to and from the Sun Ray server. The channel through the MAN can be disregarded because it can be configured in such a way that it does not add any significant amount of latency, is error-free, and delivers packets in the correct order. Other proxy switches can be set up elsewhere, and channels can be set up to these switches as well.

It is now irrelevant where the Sun Ray appliances are located geographically. They are all logically on the same LAN. Because we deal with per-port VLANs on the proxies, we can set up a proxy to contain several VLANs to different Sun Ray domains (residing on the same server or on different servers). Thus, one VLAN on a proxy can consist of one Ethernet port and one broadband channel, while the rest of the Ethernet ports plus another channel can be set up as a different VLAN.

FIGURE 3 illustrates a MAN that has a data center with four Sun Ray domains, geographically spread over the MAN into smaller or larger Sun Ray sites. This figure shows that although a Sun Ray backnet is a LAN, the LAN can be spread over a geographical area.

The figure also shows that it is possible to have several different Sun Ray LANs separated from each other on the broadband-based MAN at the same time. There is no limitation to the number of LANs it is possible to have.

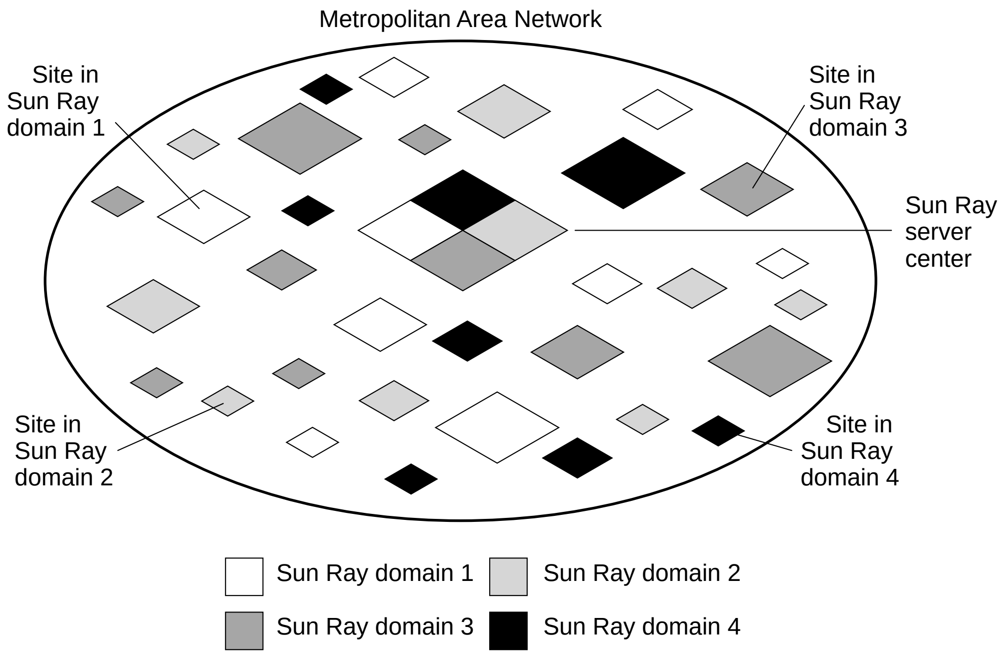

FIGURE 3 Metropolitan Area Network

The Sun Ray Environment Using Broadband Technology

Note – One of the major manufacturers of broadband ISDN and ATM equipment when this article was written, is Marconi (formerly Fore Systems). The following sections refer to Marconi equipment when describing LAN emulation hardware and BICs. Even though Sun Microsystems makes their own BICs, the Marconi BICs are more suitable to the Sun Ray environment over a broadband network. All work done with broadband technologies over ISDN, so far, has been done using Marconi equipment, even though talks with other vendors are held continuously.

Setting up, configuring, and maintaining broadband channels through a MAN is a complicated task, especially in large to very large networks. By superimposing broadband ISDN on top of SONET or SDH and using LAN emulation, the task can be simplified.

LAN emulation is easy to configure and maintain, and the necessary channels are automatically set up. Load sharing, redundancy, and failover are features built into the network and need not be configured separately. An error-free network is also automatically made available as well as in-order packet delivery.

On the server side, several BICs can be trunked together to form what can only be described as an *interface conglomerate*. Alternatively, they can be used as several separate interfaces. Useful BICs have bandwidths of either 155 Mbps (OC3) or 622 Mbps (OC12). One to four BICs can form one interface conglomerate, yielding a maximum total capacity per conglomerate of approximately 2.5 Gbps (OC48).

BICs, in an interface conglomerate, do not have to reside on the same bus inside the server, thus preventing bus overload. It is also possible to have several interface conglomerates on one server. One interface conglomerate can hold several emulated LAN interfaces. Each of these emulated LAN interfaces is perceived by the machine as an Ethernet interface and handled as such with, for instance, the `ifconfig` command. An emulated Ethernet interface can be used as a Sun Ray backnet.

The frontnet can be run over a narrowband interface (10/100/1000 Ethernet, FDDI, and so on) or over an emulated Ethernet interface on one of the interface conglomerates.

MASS Superimposed Over Narrowband Networks

If a CityNet is run over narrowband networks, it can still be used as a carrier with additional equipment such as dense wavelength division multiplexing - DWDM (preferably from Nortel Networks or Marconi because they provide ring-type DWDM equipment).

Standard fiber technology uses monochrome light bursts as an information carrier over the media. DWDM uses polychrome light so the fiber appears to be several fibers (one for each color). Each one of these virtual fibers has the same capacity as the original fiber. By introducing DWDM between the fiber media and the equipment using it, you have more fibers. When running the Sun Ray broadband networks over one color, the other narrowband traffic is not disturbed, provided that it runs over a different color.

MASS Pilot Deployment and Tests

MASS has been deployed as a pilot project with mission critical applications running on it. In addition to this, practical tests have been carried out in the Sun facilities at Watchmoor 2, in London, UK. The rest of this article contains the report from these tests.

Tests in the Watchmoor Facilities

It is possible to use a broadband network as a carrier for the Sun Ray protocol. The tests described were carried out by the Sun Professional Services Network Team in cooperation with the Sun geographic sales office (see TABLE 1).

Testing clearly demonstrates that a Sun Ray server, directly attached to a broadband network, with the use of broadband-to-Ethernet proxy switch technology (media gateway), provides several beneficial features, such as scalability, high availability, redundancy, and long distance coverage. The tests described in this article are, however, aimed at exploring the functionality and behavior of a Sun Ray environment carried over a combination of broadband and Ethernet technologies. The tests do not cover scalability, general performance, or performance under load.

Background

Due to a concern for slow connections, networks for Sun Ray appliances are very strictly defined.

Deployment of Sun Ray appliances is limited to small office environments (30–50 Sun Ray appliances per server). This stipulation has, however, not hindered speculation about using Sun Ray appliances as a maintenance-free alternative to PCs for the average person in urban areas where MANs or CityNets already exist. To verify both the limitations and the speculation, the first series of tests were carried out 18 months ago at an actual customer site in Sweden. This customer site currently uses broadband technology as their backbone.

Sun Ray appliances were deployed in different campuses with one Sun Ray server (a Sun Enterprise™ 250 server) placed in a computer center on one of the campuses. Cable lengths between the campuses were up to 15 km. The Sun Ray appliances were attached to Ethernet/broadband proxy switches over 10 Mbps and 100 Mbps networks. To this day, the customer runs this configuration in production and is very pleased, even though the configuration is not formally supported.

To demonstrate the technology in a controlled environment, the two person team (Lars Persson and Jane Lundstrom) who performed the work on Sun Ray appliances over broadband networks in Sweden, were asked to set up a miniature replica of the customer's Sun Ray environment.

Equipment

The broadband equipment used in the tests was manufactured by Marconi. The computers and Sun Ray appliances were made by Sun Microsystems.

- One ES3810 proxy switch with dual OC3 uplinks, two 100 Mbps Ethernet connectors, and 12+24 10 Mbps Ethernet connectors
- Two ASX200 BX backbone switches with 3X4 OC3 and 1X1 OC12 per switch. These ports are both UTP and SC fiber, and the fiber ports are both single-mode and multimode.
- One Ultra Enterprise™ 10 workstation functioning as a Sun Ray server with a dual OC3 BIC
- Two Ultra 1 workstations functioning as a web server and a video/audio TRADER with a single OC3 BIC
- One AVA-300 video/audio digitizer
- Two Sun Ray appliances

FIGURE 4 shows how the equipment was set up.

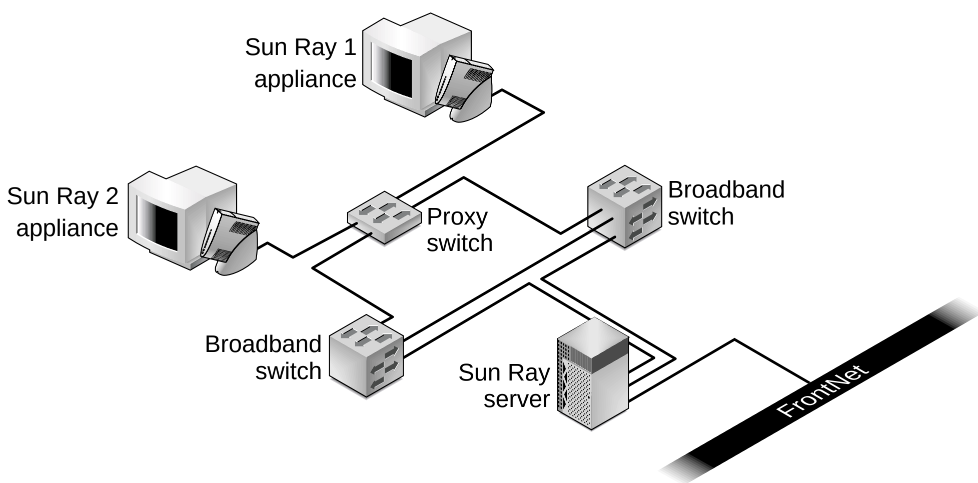

FIGURE 4 Physical View of Sun Ray Replica Environment

Broadband Technology Used and Why

Based on the experience of the test team and hardware availability, LANE 2.0 over-broadband ISDN was chosen. Broadband ISDN runs on top of both SONET and SDH. In these tests, SONET was chosen. There is no reason why Sun Ray appliances should not run over raw SONET or SDH (PPP), but currently this technology is limited to broadband networks using Ethernet proxy switches.

TABLE 1 Broadband Tests in the Watchmoor Facilities

Test No.	Test	Result
1.	Connectivity test—Load a Sun Ray appliance through a proxy switch over one broadband switch to a broadband-attached Sun Ray server with one BIC.	A Sun Ray appliance boots over 10 Mbps and 100 Mbps networks from a broadband-attached Sun Ray server. It is possible to switch between 100Mbps and 10Mbps without having the Sun Ray appliance lose connection to the Sun Ray server. The session is maintained even though the Sun Ray appliance is moved to another port on the proxy switch, both on the same card, and to a different card. It is also possible to move between 10 Mbps and 100 Mbps and get a renegotiation of interface speed. As a result, the Sun Ray server can be moved to different ports on the broadband switch without loss of connectivity. On the Sun Ray appliance, the session froze for up to 15 seconds with each move, but the session was never lost.
2.	Same as Test 1, but use dual uplinks on the proxy switch and test failover.	Unplugging one interface caused the switch to fail over to the next BIC. Reconnecting the interface and unplugging the other one made the server fail over to the first BIC again. On the Sun Ray appliance, the session froze for 10 to 15 seconds, but the session was never lost.
3.	Same as Test 2, but over two broadband switches with the server connected to one broadband switch, and the proxy to the other broadband switch (one interlink hop).	Same results as observed in Test 1.
4.	Same as Test 3, but using one proxy uplink on one broadband switch, and the other proxy uplink to the second broadband switch.	Same results as observed in Test 2 for failover.

TABLE 1 Broadband Tests in the Watchmoor Facilities

Test No.	Test	Result
5.	Same as Test 4, but using dual BICs on the server, each connected to its own broadband switch. Tested for the ability for the Sun Ray session to survive the failure of one of the switches (power cycling), and when the broadband switch returns, it should survive power cycling of the other broadband switch.	Failure of an entire broadband switch does not affect the Sun Ray appliance session, except for the time delay described previously. Failure of the second switch, when the first switch was restored, yielded exactly the same results.
6.	Server attached to one broadband switch, proxy switch attached to the other broadband switch, and the inter-broadband switch link consisting of a 35 km single mode fiber.	Because of the hardware limitations of the broadband switches, the entire test was not possible to fulfill. It was, however, possible to implement a work around by using one quad board on one of the broadband switches and forcing the traffic on one port and back in on the other, thus enabling us to boot and run a Sun Ray appliance at a distance of 35 km from the server. There were no observable delays in the Sun Ray appliance response or general behavior.
7.	Introduce digitized PAL video from a camcorder (or similar) to an audio/video unit (AVA) handled by an OC3 attached TRADER and display on a Sun Ray appliance through a 10 Mbps and 100 Mbps Ethernet network.	It is possible to transmit PAL or NTSC video streams through the broadband network and out to the Sun Ray appliance on both 10 Mbps and 100 Mbps. The image quality is good, but to get a good streaming ability, the video transmit streams must be fine-tuned. Because performance is not within the scope for these tests, such tuning was not done. The video stream is currently good enough for less sophisticated video conferences. Currently, we do not know the quality of streaming video we can achieve. No audio streams were tested due to lack of audio equipment, but the AVA handles video and audio steams in a similar way. Such a test would not, at this stage, have yielded anything more than the one performed.
8.	Establish a method of inter-connectivity between the test environment and service delivery network (previously known as Architecture.COM).	Inter-connectivity to the service delivery network was accomplished through a per-port VLAN on a proxy switch. The VLAN connects to a service delivery network through a 100 Mbps Ethernet network. Inter-connectivity was tested using a minimal service delivery network with a web server. It was possible to connect to the web server from the Sun Ray appliance through HTTP.

ELANS and VLANS

Two ELAN/VLAN configurations were established, one called `security` and another called `london`.

The `security` configuration was the ELAN/VLAN configuration for the private Sun Ray network and `london` was the ELAN/VLAN configuration for the public network. In addition to this, a third ELAN/VLAN configuration called `mgmt` was configured and used for out-of-band management purposes. All ELANS were set up as anycast services on both switches.

The global topology file for the broadband environment was set up as a DLE set (Distributed LAN Emulation) on both switches, as is shown in the following code example.

CODE EXAMPLE 1 Global Topology File for Broadband Environment

```
# LECS.CFG
# Date: 12/24/00 19:57
# Revision date: 2001-03-03
# TFTP-host(s): Sun Ray01-m
# User: Sun Microsystems
# Revisor: Lars Persson, Sun PS
# LECS in asx21, asx41
#
# The search ordering of elan names
#
Match.Ordering: london, security, \
                mgmt
#
# Parameters for elans
#
.Multicast_Send_VCC_Type:        Best Effort
.Maximum_Unknown_Frame_Time:     1
.LAN_Type:       Ethernet/IEEE 802.3
.Maximum_Unknown_Frame_Count:    1
.VCC_TimeOut_Period:     1200
.Forward_Delay_Time:     15
.Maximum_Frame_Size:     1516
.Expected_LE_ARP_Response_Time:  1
.Path_Switching_Delay:   6
.Aging_Time:     300
.Control_TimeOut:        120
.Connection_Complete_Timer:      4
.Flush_TimeOut:  4
.Maximum_Retry_Count:    1
#
# Parameters for DLE elan: london
```

CODE EXAMPLE 1 Global Topology File for Broadband Environment *(Continued)*

```
# LES/BUS on asx21, asx41
#
london.Address:                   c5100000aaaa0000aaaa0000aaaa0000aaaa0192
london.Accept:                XX.XXXX.XXXX.XXXX.XXXX.XXXX.XXXX.XXXXXXXXXXXX.XX
#
# Parameters for DLE elan: mgmt
# LES/BUS on asx21, asx41
#
mgmt.Address:                     c5500000aaaa0000aaaa0000aaaa0000aaaa019a
mgmt.Accept:                  XX.XXXX.XXXX.XXXX.XXXX.XXXX.XXXX.XXXXXXXXXXXX.XX
#
# Parameters for DLE elan: security
# LES/BUS on asx21, asx41
#
security.Accept:                 XX.XXXX.XXXX.XXXX.XXXX.XXXX.XXXX.XXXXXXXXXXXX.XX
security.Address:                 c5600000aaaa0000aaaa0000aaaa0000aaaa019c
#
# entries that the VLAN Manager does not parse at this time
#
LECS.Reload_Period: 30

All equipment used was set up with an ELAN instance in the mgmt net so it could be
reached through the broadband network. The IP plan on the management network
was as follows:
Network> 10.1.0.0 / Netmask> 255.255.254.0 / Broadcast> 10.1.1.255
asx1           10.1.1.21
asx2           10.1.1.41
es3810 (proxy switch)        10.1.101
Sun Ray server                10.1.1.244

On the Sun Ray server networks were configured as follows:

lo0: flags=1000849<UP,LOOPBACK,RUNNING,MULTICAST,IPv4> mtu 8232 index 1
        inet 127.0.0.1 netmask ff000000
fa0: flags=1000842<BROADCAST,RUNNING,MULTICAST,IPv4> mtu 9188 index 5
        inet 0.0.0.0 netmask 0
        ether 0:20:48:2e:1f:c6
fa1: flags=1000842<BROADCAST,RUNNING,MULTICAST,IPv4> mtu 9188 index 6
        inet 0.0.0.0 netmask 0
        ether 0:20:48:2e:33:e
el1: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index 7
        inet 10.1.1.244 netmask fffffe00 broadcast 10.1.1.255
        ether 0:20:48:2e:1f:c6
el0: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index 8
        inet 192.168.128.1 netmask ffffff00 broadcast 192.168.128.255
        ether 2:20:48:2e:1f:c6
```

CODE EXAMPLE 1 Global Topology File for Broadband Environment *(Continued)*

```
el2: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index 9
        inet 193.182.63.94 netmask ffffff00 broadcast 192.182.63.185
        ether 6:20:48:2e:1f:c6
```

In this example, `security` is represented by `el0`, `london` by `el2` and `mgmt` by `el1`. The interfaces `fa0` and `fa1` are the actual BICs. The `el` interfaces emulate Ethernet. The proxy switch has two ELAN / VLAN combinations (`london` and `security`) and one ELAN (`mgmt`).

The `security` VLAN was used to connect the Sun Ray appliances to and `london` provides Ethernet access to the public network. The `mgmt` ELAN instance was for administrative purposes. The `london` and `security` have both 10 Mbps and 100 Mbps Ethernet interfaces. Finally, the ASX `en` has one instance in each the `mgmt` ELAN for administrative purposes. The AVA resides on the `london` network and so does the AVA TRADER.

Conclusion

With limitations in the video broadcasting test and some reservations in accordance to the distance test, it can be seen clearly that Sun Ray appliances do work in a broadband environment, and that the infrastructure itself offers several beneficial features normally provided by computers—load sharing, robust and persistent session handling, failover and so on. With long distance single mode fibers it is possible to have a network radius of at least 120 km (one interlink hop).

Author’s Biography

Lars Persson, Senior Systems Specialist—Networking

Lars works as a Senior Systems Specialist in Sun Professional Services in Denmark. He has been with Sun for six years. During this time, he worked with various network technologies, with emphasis in Broadband ISDN and Gigabit Ethernet. His expertise is used in designing large corporate network architectures, and he also works on telco projects. Lars is also a team leader for Denmark’s Networking Focus Area Team and he was awarded Sun™ Star 2000.