Network requirements for overseas Virtual Desktop Deployment

Background

This document shows extracts of a Network design that would allow a Virtual Desktop environment to be deployed overseas. The design is an example of a solution that could be deployed within an overseas campus environment.

The scope was to allow offshore 3rd pary employees access to systems within two UK Data Centres (London and Swindon) via a terminal services environment. The solution also had to provide connectivity to the same infrastructure for company employees both from within the workplace and over the internet.

The product chosen was Sun Global Desktop (SGD) with Wyse terminals deployed at 3rd party offices in India using two factor authentication. The network within the 3rd party was initially seperated by vlan trunks and a firewall but later was physically seperated and became an extension of the UK Corporate network.

The 3rd party clients had access to relevant applications dependant on job roles, the permissions set within SGD by the administrator determined the level of access and therefore number of applications visible on the desktop.

The traffic between the Terminal and the UK SGD servers utilised the Secure Socket Layer (SSL) protocol with AES 256bit encryption. No hard discs were installled in the Wyse terminals and USB ports were disabled. User id and RSA Tokens were used for the 2 factor authentication.

The WAN links were originally leased lines which were replaced by an MPLS IPVPNQoS network which allowed IP Telephony and IPVideo to be incorporated into the solution at a later point.

The MPLS Network gave a round trip delay (RTD) of 180-200 msec which was found to be acceptable for Terminal services and Voip and IP Video. Issues were seen when the optimum route could not be taken and the MPLS traffic had to follow an alternate path back to the UK. RTD of >250msec saw keystroke echo being experienced on the terminal screens within the 3rd party India locations and >300msec saw Voip calls affected by latency and jitter.

The above issues may have been resolved as this solution was implemented in 2007-2008 but worth noting if setting up SLA's with Telcos for Overseas connectivity. Most Telcos offer SLAs based on availability of the network but worth investigating what RTD the availability is set against.

1. SGD Architecture

SGD (Sun Secure Global Desktop) servers provide a web front-end for user access to applications. The SGD servers are placed in a DMZ environment in the London and Swindon datacentres. Load balancers will be used to balance traffic across pools of SGD servers. Network links will be provided to both London and Swindon which will work alongside Global Server Load Balancing (GSLB) to deliver a high availability SGD solution across both locations. The SGD server servers will be split up into three pools (internal, Offshore & Internet). The internal pool will be reached though internal-facing load balancers. The offshore and internet pools will be reached through external-facing load balancers.

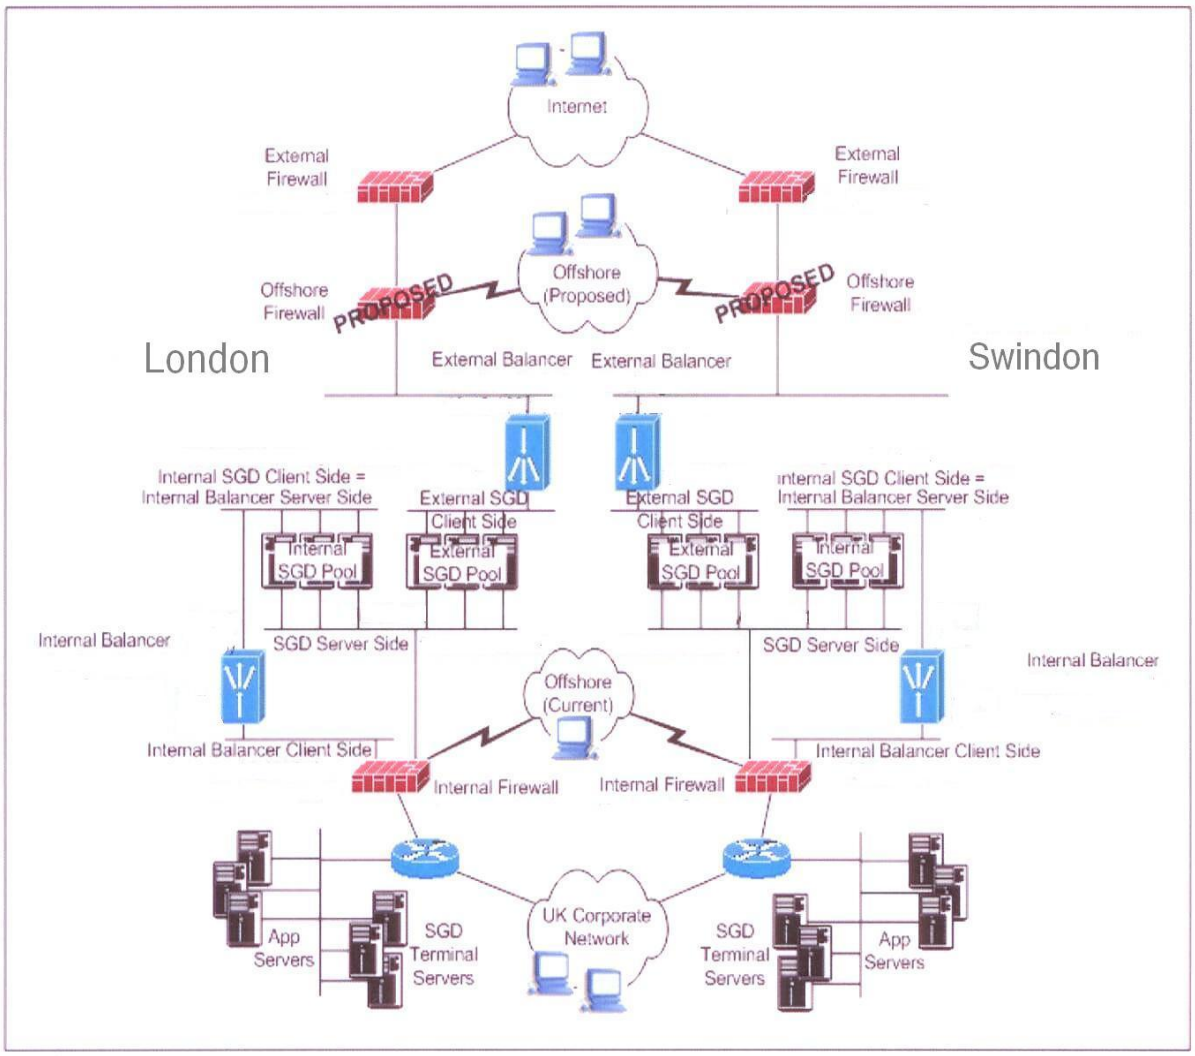

1.1 Current Infrastructure

In the current architecture (shown in figure 1) the offshore provider uses thick clients and hence a firewall connection is required to their local office network. The offshore LAN is delivered using the local provider's LAN switch infrastructure.

Figure 1

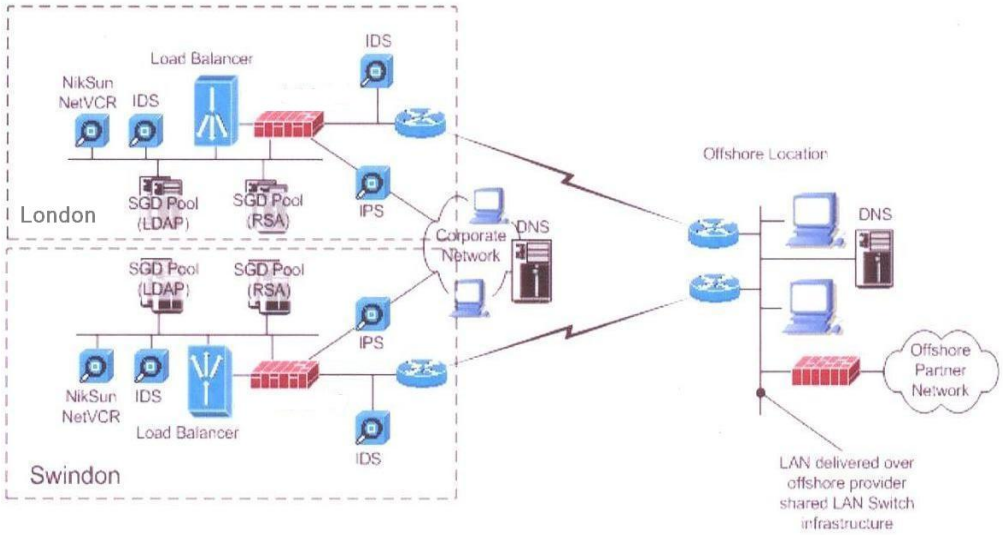

1.2 Proposed Infrastructure

In the proposed arrangement (shown in figure 2) the offshore network will connect off a separate dedicated firewall. In addition, the proposed use of thin clients at offshore locations will allow the removal of the firewall link to the local partner office network.

It is proposed that a separate LAN switch environment will be provided offshore to overcome the need to use shared existing infrastructure.

Figure 2

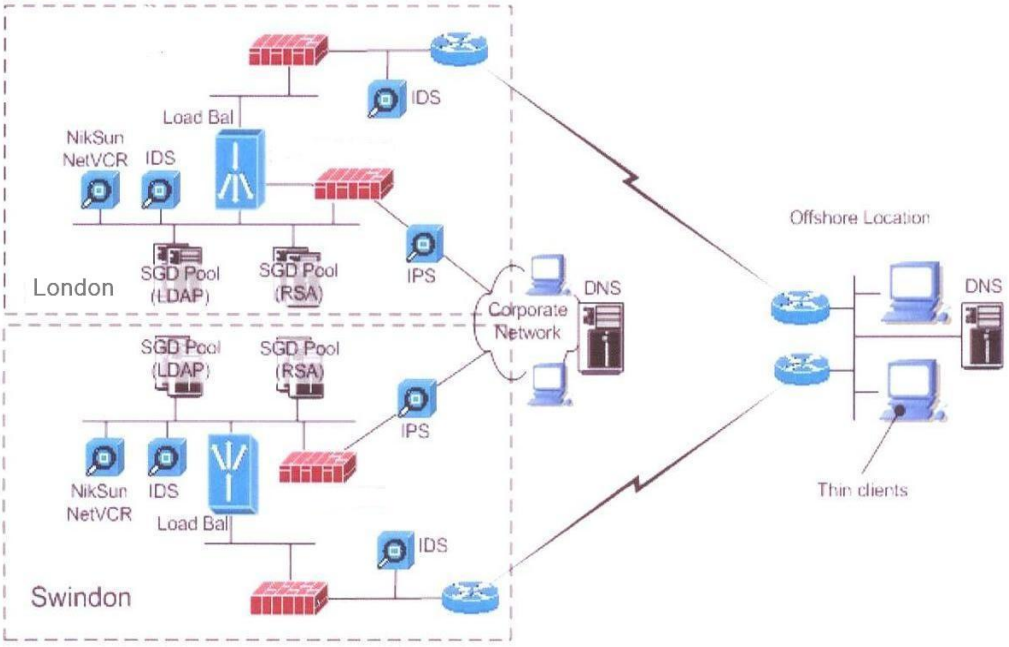

1.3 Balancer Topology

Four Nortel Application Switch 2208 units (running OS ver 23.2.1.1) have been purchased and will be installed to achieve the following topology:

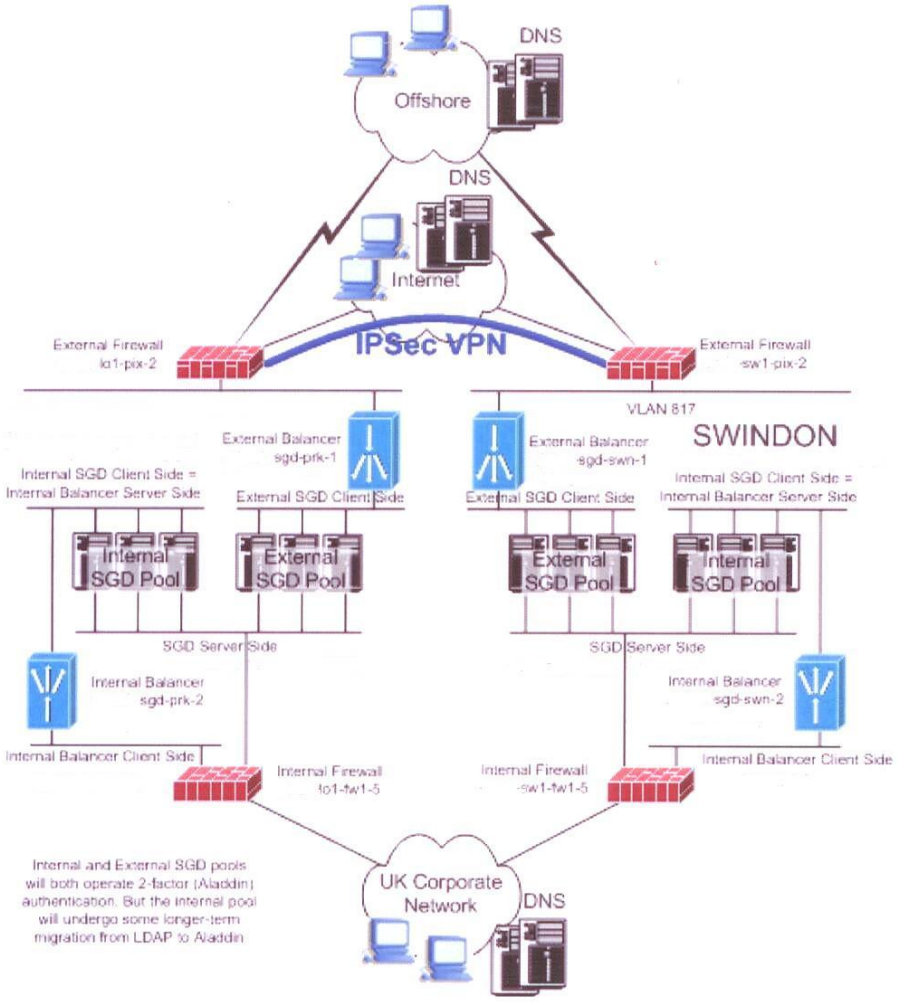

The balancers will spread load across their local pool of SGD servers. Global Server Load Balancing (GSLB) will also be used to extend that pool across both locations.

1.3 Global Server Load Balancing (GSLB)

GSLB uses DNS to share traffic destined for a particular name across a number of geographically separate IP addresses. It operates as follows:

- The SGD service will be reached using a URL for example http://sgdpool1.zone.com/
- Within the intra.com domain SGDPool1 is "delegated" to the SLBs. To achieve this a pair of NS records exist in the domain.com zone file that point SGDPool1.domain.com at the two balancers **London** & Swindon).
- When the client requests a DNS resolution for SGDPool1.domain.com the DNS server asks the two SLBs for an IP address for SGDPool1.domain.com.
- Depending on the configured load balancing plan, the **London** balancer will respond with a DNS A-Record containing either it's own IP address or the Swindon balancer IP address. Swindon will respond in the same way.
- The two balancers send GSLB updates (DSSP) between themselves to work out the current balance of traffic. In the event that an SLB can't reach it's partner, it only offers it's own IP address in it's DNS response.
- The client's local DNS server then sends the A record on to the client thereby directing the client to the appropriate SLB.

Local SLB uses "Least Connections" as the default server selection metric. We will configure GSLB so that it also uses "Least Connections".

FAILOVER RULE

A working DSSP connection will result in both balancers offering eachother as a viable target for clients. If a client's network path to one balancer is down but DSSP survives then the available balancer will offer the client the unavailable balancer. So, if a path to a balancer is down then DSSP must also be down.

If DSSP follows the path to the client then the above is satisfied. But since there are many clients and only one DSSP path then we have a problem achieving high availability.
This issue is at it's worst where there are two very separate user communities. For example the external balancer is serving the internet and offshore communities. If the DSSP path were to follow one path then the other would not achieve high availability. **Our solution** to this is to increase network redundancy for offshore by building an IPSec tunnel through the internet to link the balancers. This tunnel would carry DSSP so failover protection is provided for the internet community.

Clients on the external side connect via at least two separate network paths (the internet and the offshore network). The GSLB updates can only flow via one of these paths and are therefore unable to detect a failure on the other path. The end result would be that the available balancer will be advising clients that they should connect to the unavailable balancer. The solution is to build a direct IPSec VPN connection between London and Swindon to provide a failover path in the event that one of the offshore WAN links fail.
The following diagram shows how this failover works:

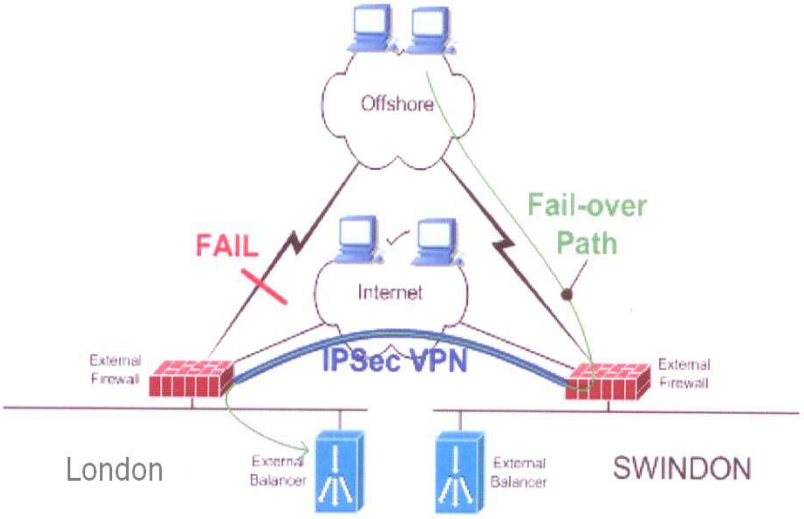

A GRE tunnel will be carried inside the IPSec VPN between the WAN routers. A GRE tunnel will allow us to carry dynamic routing information across the VPN link thereby providing the mechanism that allows the VPN to come into operation automatically in the event that an offshore WAN link fails. The Balancer's GSLB update traffic (DSSP packets) will also be carried in the IPSec tunnel.

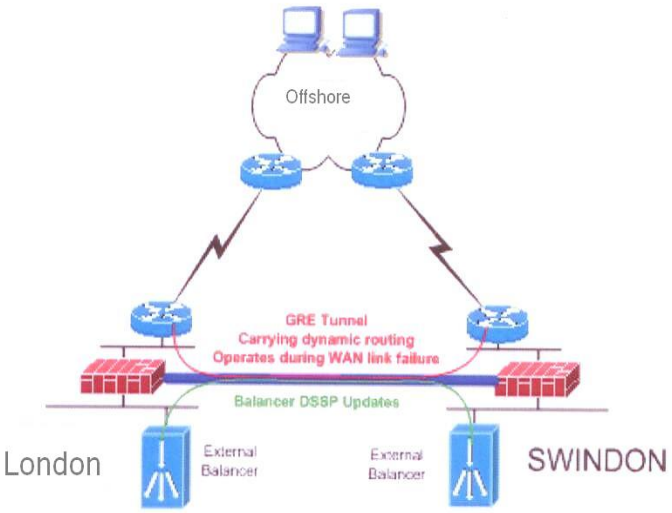

On the internal side there is only one network path to the internal clients. The GSLB updates can be sent along this path and can be used to check for failure on that path (no IPSec VPN needed). This is shown in the following diagram:

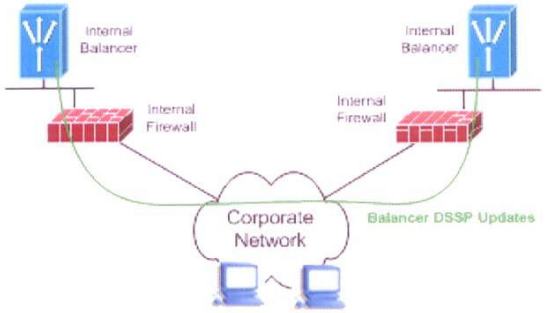

1.5 Domain Names & SGD Server Balancing Pools

On the external balancer there will be two pools relating to

- Offshore = desktopRSA. intra.com
- Internet = desktop. .com

Two separate pools are needed on the external balancer because the same SGD server cannot run two URLs in parallel. Please note however, that we plan to run two SGD server instances on each Sun server. Each instance will have it's own IP address.

On the internal balancer there will be just one pool relating to

- desktop.hosts.plc. intra.com

It has been agreed that the entire internal user community will use 2-factor authentication. However, it is likely that there will be a period where single-factor and 2-factor is in use in parallel thereby requiring two internal pools and to internal URLs.

1.6 GSLB DNS Config

At the client location, the delegated DNS server zone file configuration would look something like the following. We are recommending reducing the DNS TTL value to 10 minutes (600 secs) but it should be remembered that the DNS resolver does not pass this value to the application and IE will has it's own TTL of 30mins. We will need to accept that browsers may get stuck on an unavailable balancer for up tot 30 mins.
We assume here that the balancer issues only one A record (and not a weighted list of A records).

On the offshore DNS systems, locate the intra zone file. This starts with the following:

```
@          IN      SOA      intra.com.
```

Modify the TTL value (the last value in the parenthesis) to be 600 sec

1.7 Traffic Flows and Static Routes

The client-side and server-side traffic flows are shown in the following diagram (showing one site only).

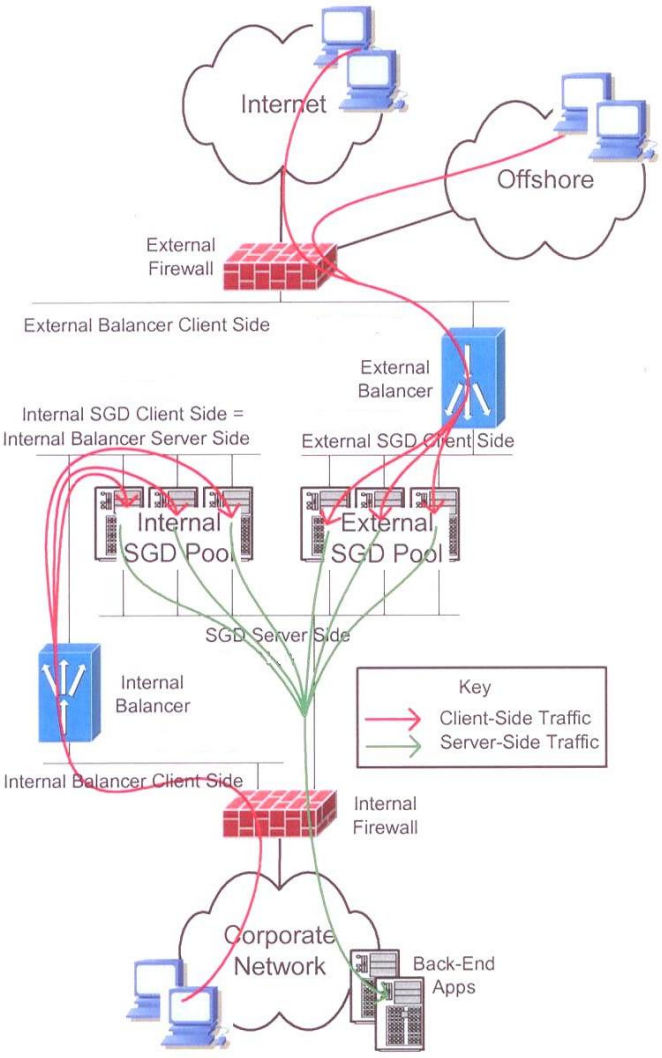

Please note that we will need to add static routing to the SGD servers and move away from the more simple default gateway configuration. The SGD servers will have a default route to direct client-side traffic to the respective balancer. A set of specific static routes will be required on the SGD servers to carry server-side in the opposite direction. Care must be taken to avoid the risk that internal client IP address ranges are not accidentally included in the server-site static routes.