Introduction to UUDynamics iSTAR

2005-6-6

UUDynamics.com

Introduction to iSTAR

Computer and networking technologies dominate many facets of our lives today. The Internet has grown into a universal, ubiquitous medium that continues to transform our society in fundamental ways. As the mainstream adoption of Internet technology increases, gaining efficiency and moving towards progress, others utilize the technology to cause disruption and to facilitate criminal activity. The growing security risk that naturally accompanies the increasing amount of critical information entrusted to the Internet has heightened our concerns. We are all painfully aware of the high frequency of hacker incidents and disastrous security breaches that can even occur in supposedly secure network environments.

The Internet facilitates a ubiquitous and global means of communication, and its related technology evolves at break-neck speed. New challenges and advances can render popular solutions obsolete overnight. In this paper we will attempt to illustrate these points by outlining the previous generations of VPN technology. We will then introduce UUDynamics' iSTAR technology that has emerged out of the growing necessity for a new solution for the increasingly dynamic needs of modern businesses and their increasingly mobile work forces.

Security protection is not an absolute concept. It is a measure taken to protect valuable resources. It must constantly evolve in order to provide efficient protection for the resources that it is assigned to. It is relative to the resources that it protects and to the level of protection required for the environment in which it is deployed. It is relative to identified threats, vulnerabilities and potential attacks and must provide sufficient defense against them. iSTAR provides a powerful platform for secure, any to any application connectivity. It provides a multitude of security tools so that our customers can formulate their security policies and tradeoffs at the correct level. It even supports low-level Layer 2 IPSec-like connectivity for site to site when it is appropriate. Its friendliness and plug-n-play model of global deployment make it a next-generation VPN product, outdating both IPSec and SSL VPN. iSTAR can successfully satisfy all of an enterprise's Intranet needs and is particularly suited for extranet activities, where business partners must carry out electronic transactions in unplanned or ad hoc situations and be able to build secure channels bridging multiple disparate administrative domains.

Evolution of Internet security

From private network to Internet-based VPN

Historically all private networking is based on leased or dedicated lines. The migration from this infrastructure to an Internet based VPN is inevitable for the following reasons:

- Cost of leased or dedicated lines is high.
- The Internet offers the most robust and universal connectivity. Internet backbone has plenty of bandwidth. It is far more cost effective to upgrade an enterprise’s on ramp, which affords connectivity to many different networks, than to lease a private line for a dedicated destination.
- Through an Internet Load Balancer, an enterprise can gain bigger bandwidth and built-in redundancy; this HA (High Availability) option is simply not available to private networks based on leased lines.
- As the world is on its way to becoming an integrated whole, connectivity must cover a much larger area than before. Private networks based on dedicated lines have very limited geographical coverage and will not be able to catch up with the increasingly interconnected world.

Although the motivation for migrating to Internet-based VPN stems mainly from the considerations of cost, reliability, and coverage, network administrators must realize the potential vulnerability and information leakage between the end points, even for private lines. The service provider from whom the private lines are leased from is the classic “man in the middle” in the jargon of security professionals, and must be guarded against as a matter of security principle. Even network constructed solely with raw leased or private lines still must institute security system. Ideally that security system should protect both its private and public lines, and assist the migration from private lines to Internet-based VPN.

Firewall for perimeter protection

In the last two decades, with the increasing popularity of the Internet and the dawn of universal accessibility, security threats and malicious attacks have grown proportionally. Around 1994, heightened security concerns brought about the concept of **firewall**, which is a simple and intuitive approach to encircle corporate resources with a wall on the perimeter, allowing the flow of traffic only through a heavily guarded **gateway.** With increasingly sophisticated tools the guard at the gateway checks not only the source or destination of the traffic, but also the content of application packets in real time. The assumption is that the perimeter of the facility can be clearly demarcated and there is sufficient processing power and big enough iron at the gate. In smaller enterprises where the big iron is not affordable or there isn’t enough incoming traffic to justify a big box, a NAT (Network Address Translator) is typically utilized to protect the facility and to share the increasingly scarce public IP addresses.

IPSec

Firewall provides the guard at the gate, but very little protection for corporate information in transit out in the public Internet. Leased or dedicated lines solve part of the problem, but customers find public Internet a better and more robust transmission medium. The desire for secure and robust communication over public Internet brought in the IPSec standard around 1997, based on which a multi-billion dollar industry was born. One can guess from its name that IPSec is based on IP protocol at the very low level of network architecture, technically known as the IP stack of 7 layers, with IPSec sitting between layer 2 and 3. It assumes the availability of public IP addresses and is very static in its deployment. The complete IP address plan must be known beforehand and made consistent among connecting parties, with the web of connection paths identified. Security policies are largely unspecified by the standard. Since it is situated at a very low level of the stack and therefore very far away from applications where the security demands are originated, a professional maintenance crew must be employed to translate the high-level requirements to low-level IP "filters". It is very difficult to bridge extranet partners when communication must be carried out across incompatible planes of IP addresses and security policies. Traveling or mobile employees frequently find it difficult to connect to their home office to fetch information, as they are blocked by many levels of firewalls and NAT devices. Network administrators find the process of translating high-level security requirements into low-level IPSec filters very complex, tedious, and prone to error. The need to install IPSec client software also precludes remote access from shared computers in Internet cafés or kiosks, or from a PDA device for which the necessary client software may not be available.

SSL VPN

SSL VPN grew out of the need to solve difficulties encountered when working with IPSec-based infrastructure, and first appeared on the scene around 2001. It is primarily a remote access solution, adopting the dynamic transport protocol of SSL and is delivered through a browser-based "clientless" client. It is closer to application-level requirements, and offers more extensive and user-friendly authentication tools. It made it easier for individual extranet partners to work together and share business resources with application support. However, SSL VPN is still just a client to gateway remote access solution. It is not a "networked" solution and is designed only to remedy the defects of IPSec. It is very far from being able to offer "extranet-to-extranet" exchanges, or application-to-application interfaces.

Attempts to integrate IPSec and SSL VPN

SSL VPN is a relatively recent development, which arrived on the scene to patch up the defects of IPSec. IPSec is still the mainstay of perimeter-based protection, with SSL VPN appliances scattered around in selected spots for more dynamic applications in large enterprises. The deployment of SSL VPN is almost always a retrofit to existing IPSec infrastructure.

In a clearly emerging trend, VPN vendors are attempting to "integrate" the SSL VPN into IPSec, disregarding just how incompatible those two solutions are. This forced integration begets complexity that is the number one enemy of security.

The need for a new architecture

The disappearing perimeter

The traditional public-IP based solution is not end-to-end, and is problematic in today's increasingly wireless and mobile world. It is no longer possible to clearly demarcate corporate perimeters and gateways where traditional VPN protection is concentrated.

As diagramed below, under the onslaught of 3G cell phones, various versions of 802.11 Wi-Fi, and many other types of firewall-penetrating software and hacking tools, the perimeter defense is no longer meaningful. It brings to mind the futile attempt to protect a vast territory by building the "great wall of China", which is time consuming, costly to build, and useless when airplanes arrive on the scene. Perimeter defense is a losing battle. The enterprise IT environment is more like an open territory than a clearly demarcated aggregate. The writing is on the wall. The ultimate and correct protection must be end-to-end.

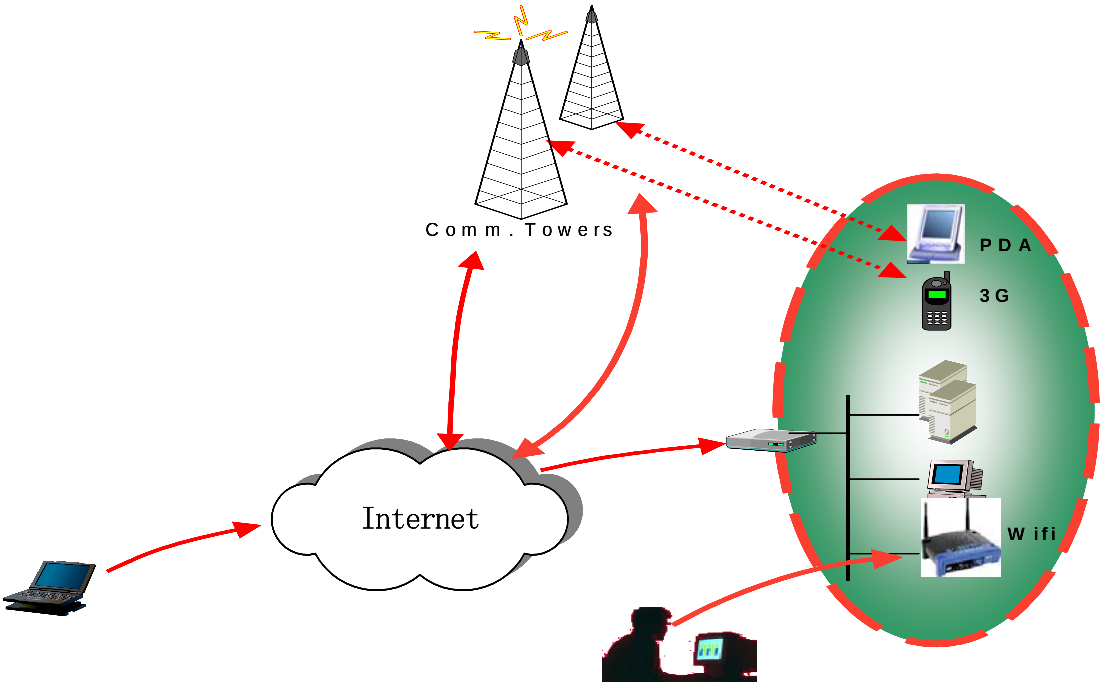

Perimeter based defense is no longer adequate

The need for a distributed solution

The traditional gateway-centric solutions are not only outdated in terms of their protection coverage, but also highly inflexible in their placement. Since their granularity of protection is corporate-wide and perimeter-based, it is rather difficult to deploy finer coverage for protecting departmental computing. Perimeter defense keeps all shareable resources in the DMZ next to front gateway, whereas ideally the departmental resources ought to be distributed to where the departments are. With this current mindset, security requirements will continue to come from the departments, but must be relayed to professional network administrators, who then translate them into low-level or machine-dependent “filters”. This procedure is indirect, inefficient, and error prone.

Due to their inflexibility of placement, traditional VPNs largely ignore the complexity inside the enterprise behind the perimeters. They are incapable of offering any kinds of meaningful structure inside the enterprise. Any additional protection to cover the real last mile, namely from the perimeter or DMZ to where the servers really are, must be designed specially to cover each situation that demands it. Those custom-designed patches add additional complexity to the enterprise and burden further the administrators.

A future-proof solution must relieve the customers from constraints imposed by the tyranny of public IP addresses, enable secure departmental computing, and extend protection to where it belongs.

The need for “any to any” secure connectivity

The revolution of Internet so far has only delivered reasonably reliable connectivity from a web client to web servers or portals. It is quite easy to reach a web-based store front or web-based resource center. However, the connectivity between two independent parties is nowhere near as straightforward, and when it is achievable, its functionality is limited to a small set of applications. Instant Messengers, in their many incarnations, fulfill the needs of “any to any” connectivity by providing a limited set of functions, like online chat, simple file transfer, and limited audio/video interaction. Connectivity through other types of applications is simply not available.

The need for a truly integrated and networked solution

The VPN market place is now presented with two disparate solutions: IPSec for static intranet, and SSL VPN for remote access. Although intersecting somewhat, each has its own distinct objectives. If an enterprise adopts both types of products, there will be two separate tools to learn, two sets of GUI, and the integration effort to join them into a workable whole. The incurred complexity increases the TCO (Total Cost of Ownership) and breeds security vulnerability in an enterprise. For those enterprises that have never had any VPN exposures, there is no alternative but to go through the same but already recognized pains of deploying IPSec first and then redress the defects with SSL VPN.

There is a clear need for a single architecture that encompasses the functions of both.

The need for route optimization

Internet-based VPN offers the ultimate robustness, universal reachability, and cost effectiveness, but it has an Achilles’ heel. A typical Internet access session from client to server must go through many Internet service providers, each of which has its own network and peering arrangement, which are not optimized for routing efficiency, but for lowest cost. Packet delay is completely outside the control of the Internet subscribers. For delay sensitive applications, such as Microsoft Terminal Services, Citrix applications, Outlook, etc., access to servers is largely a potluck experience. In situations where the delay is too long for the application to bear, there is absolutely nothing a user can do to improve it. This accounts for many frustrating incidents when a remote user wants to access his home server and discovers the ping delay is beyond the normal acceptable range, which leaves the user in despair since there is absolutely nothing he can do to improve it.

iSTAR architectural components

To meet the challenge of next-generation requirements, UUDynamics has designed from ground up a new VPN architecture. It is not IPSec or a point-to-point SSL VPN, but a single consistent architecture that delivers an all-encompassing platform intent on answering all present and future VPN needs. It is based on the tried and true "central switching" concept that has proven itself in the past 100 years by the ubiquitous telephone system. This "central switching" based new architecture is called iSTAR (instant Secure Tunnel ARchitecture).

iSTAR delivers any to any connectivity and end-to-end security through this new architecture. It consists of two parts: the central switching units, and the end objects, together they support a logical naming scheme based on UUIDs. Externally an iSTAR client may securely access any iSTAR object assigned with a UUID.

Central switches

The central switching system is comprised of **UUSwitches**, which are the central servers that manage the UUID directory and the building blocks for the switching fabric where network applications exchange their protected or tunneled data.

This central switching system provides addressability for all elements assigned with a UUID, and mediates data transmission for the end-to-end connection. The central switching system assigns and distributes all UUIDs, and verifies the authenticity of a presented UUID submitted by the end object. Since UUSwitches control and guarantee the authenticity of all holders of UUIDs, this method has the added advantage of being able to share a single SSL Certificate across the whole iSTAR deployment. All UUSwitches present themselves externally with a single DNS name and a single SSL Certificate. There is no need to purchase additional SSL Certificates for end objects like application servers or sites.

The switching fabric is for connecting end objects when there is no direct IP routing path available. The central switching system can be fully distributed in a multi-site deployment, with multiple UUSwitches in a clustered configuration at each site. A UUSwitch can also be deployed in "Slave" mode without taking part in the management of UUID and used only as a building block of the iSTAR switching fabric.

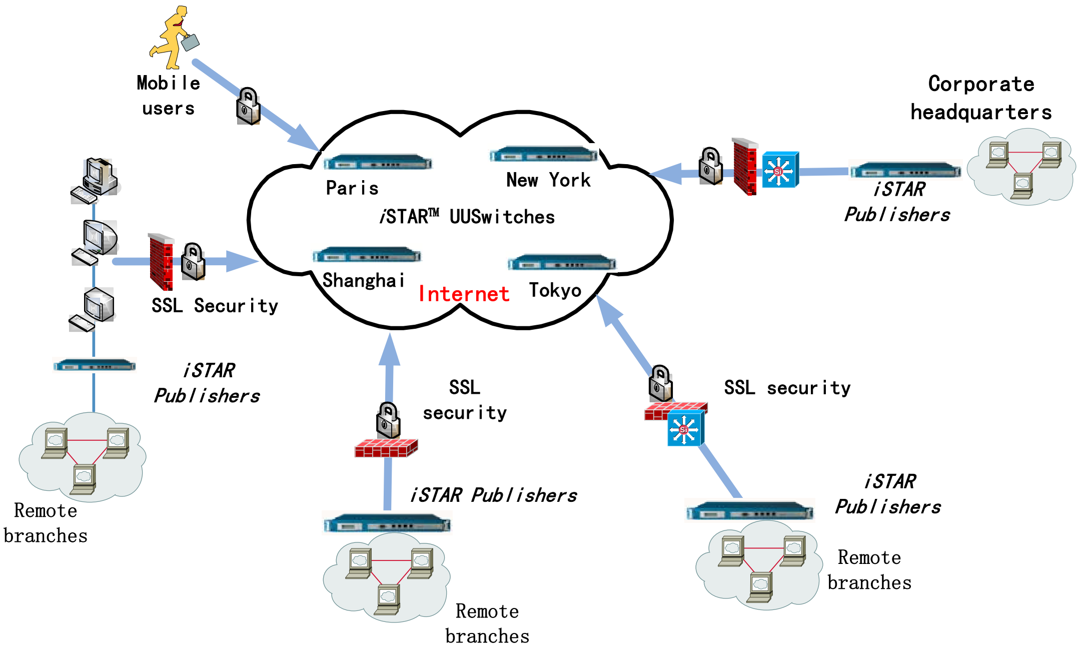

As diagramed in the above picture, all end objects connect to UUSwitches through their respective firewalls as outgoing connections. There is no need to reconfigure any firewalls or drill any holes in the firewalls in order to participate in the iSTAR community. As is also clear from the picture, the protected resources need not be situated at the gateway anymore, and can be placed where they need to be. The enterprise can connect to any end object with a UUID from behind even a NAT, which denies all incoming connections. Since there are no holes required on the perimeter, the enterprise is now largely immune to IP and port scanning, two of the most popular attack tactics.

End objects

iSTAR applications adopt the “client/server” model in a distributed processing environment.

Publisher

This is the gatekeeper that protects the server resources. The server **publishes** its protected services by allowing access from authorized **clients**.

Teleporting agent

This is the most secure and flexible iSTAR client software for qualified client/server applications. Through the **teleporting** technology, all network contexts and network calls at the client side are captured and transplanted to the server side, independent of the IP addresses. This technology offers true extranet interfaces where two independent administrative domains can securely connect to each other instantly without any extraneous concerns irrelevant to the task at hand, including the compatibility issues of IP addresses. The iSTAR protection mechanism is at the level of applications, which do away the error prone practice of translating application security requirements into low-level IP filters.

Lower level Clients

In addition to the teleporting agent mentioned previously, there are a few other clients that support lower level interactions in situations where the teleporting technique is not applicable. The **UURemote** client offers a wire level connection for highly trusted personnel. For non-Windows CE PDAs and other platforms where the teleporting client is not supported, the web browser can also be used as a "clientless client" for remote web access.

iSTAR principle of operations

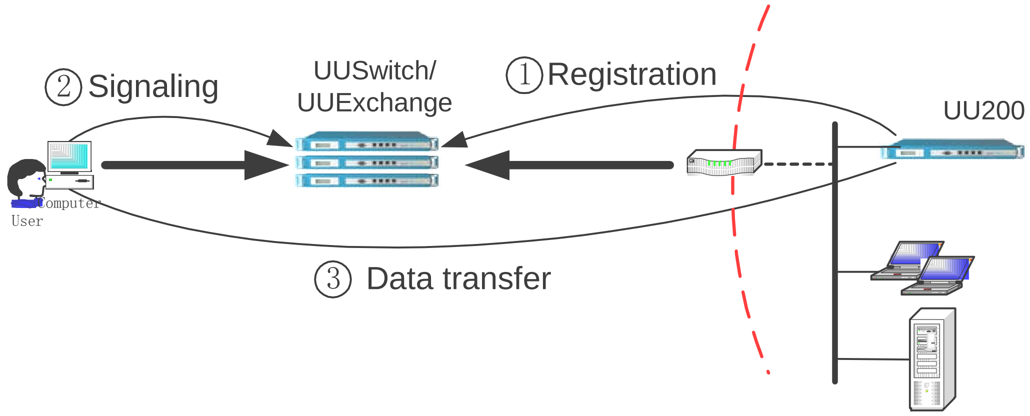

UUID

Each addressable unit in iSTAR is assigned with a unique UUID, for which a UUID directory is constructed within the central switching system for UUID lookup and authentication.

Registration

Each UUID holder on startup must register itself with the central switching system and make itself known to the iSTAR community. Through this **registration** process the authenticity of the UUID is verified, and the addressability to the owner of the UUID is dynamically established. This dynamic process of establishing the addressability of the UUID also affords the UUID owner great mobility. The owner of the UUID may move around arbitrarily and still maintain its addressability as long as it can reach the central switching system.

Connecting to servers: Signaling

An iSTAR client addresses its target server through the server's UUID. The process of contacting the central switching system, looking up the UUID, establishing the command channel to the target server, submitting the connection requests, etc. is called iSTAR **Signaling**. This is similar to the traditional telephone switching system under which a signaling process (SS7) must first take place before a phone call can be established.

Connecting to servers: Data transfer

Once the target server is identified and located (through its UUID), connection to the server must go through the switching fabric within the central switching system. The switching fabric is composed of all the UUSwitches, and UUSwitches in **slave mode**, which can be multi-site, and multi-station at each site.

URLs to published applications

With the exception of UUSoft and LAN2LAN, all other access sessions to server resources are triggered through web browsers.

Publication of services

Through the deployment of iSTAR publishers, network administrators can **publish** local resources to remote users or remote LAN. Access to published resources can be either symmetric, as in the case of LAN2LAN, or asymmetric, as in the case of a client trying to get access to server resources.

The case for symmetric LAN2LAN access is straightforward. A virtual router, constructed through two UU200 units, can join two LANs with it.

Client experience

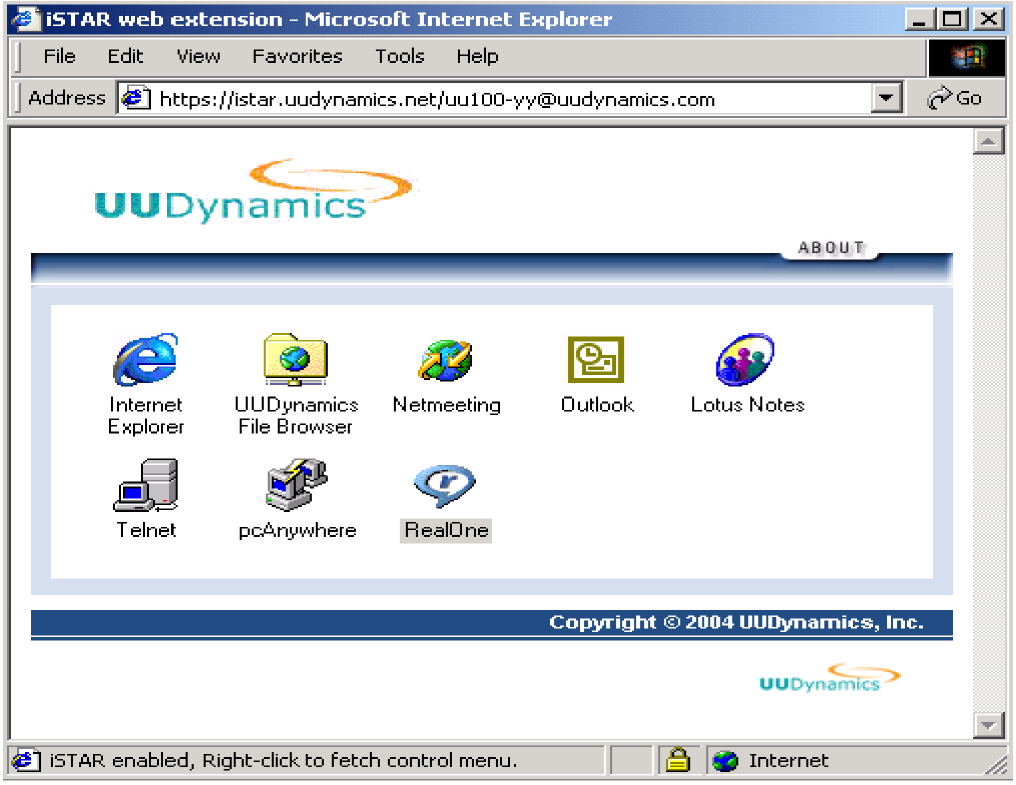

iSTAR web extension - Microsoft Internet Explorer

Address https://istar.uudynamics.net/uu100-yy@uudynamics.com

iSTAR enabled, Right-click to fetch control menu.

Displayed in the above picture is what a remote client sees when accessing a publisher. Depending on his credentials, the user is presented with a **remote desktop,** on which all authorized applications are listed. Different users may see in his version of the remote desktop a different number of application icons, depending on their assigned authority.

The appearance of those presented icons is exactly the same as those seen in the Windows desktop. iSTAR's support of applications is completely transparent to individual applications. There is no change to their look and feel and there is no need for retraining. It also offers a location independent wrapper so remote users need not know where those published services reside. Those icons are "remote short cut" to published applications.

Note that Telnet is published to the remote desktop displayed above, turning an insecure client/server application into a secure application. There is no need to purchase a separate secure version of Telnet like SSH. Similarly by publishing the FTP an administrator can instigate a secure FTP without a separate installation of SFTP. NetMeeting becomes secure NetMeeting when launched through iSTAR. This demonstrates the plug-n-play nature of the iSTAR secure application platform.

iSTAR benefits

Any to any connectivity and end to end security

It is quite clear with the UUSwitches and UUID naming scheme that iSTAR provides the much desired any to any connectivity without dependency on public IP addresses.

The power of "any to any connectivity" must be protected and secured, or it can become an "any to any" security hazard. For client to server connection, iSTAR not only connects the client to where the resources are, it also fully protects that communication from end to end. This is a major departure from traditional VPN approaches, which connect you to the gateway or perimeter of an enterprise, not to the location of the resources; they are usually deeply situated inside the enterprise away from the perimeter. iSTAR protection successfully covers the security gap between the gateway/perimeter to the resource servers, which is not protected by traditional approaches.

Route optimization

By strategically placing iSTAR central switching units, the administrators of iSTAR can construct a configuration that affords them the ability to alter and therefore control the routing over public Internet for application data. This facility for controlling routing over public Internet allows iSTAR to be deployed with many benefits available only to leased/dedicated lines, but at the very low cost level of public Internet connection.

iSTAR central switches can be hosted in any IDC, their capacity is usually quite large and is readily obtainable globally. Negotiating for IDC spaces is much simpler than other types of peering arrangement. iSTAR architecture affords a network overlay, built as a super-structure above the unpredictable lower Internet jungle, with the ability to optimize for better routes and shorter delay. In cases of intercontinental access, iSTAR frequently improves the delay characteristics to a significant degree that makes the vital difference between a successful connection and a failed connection.

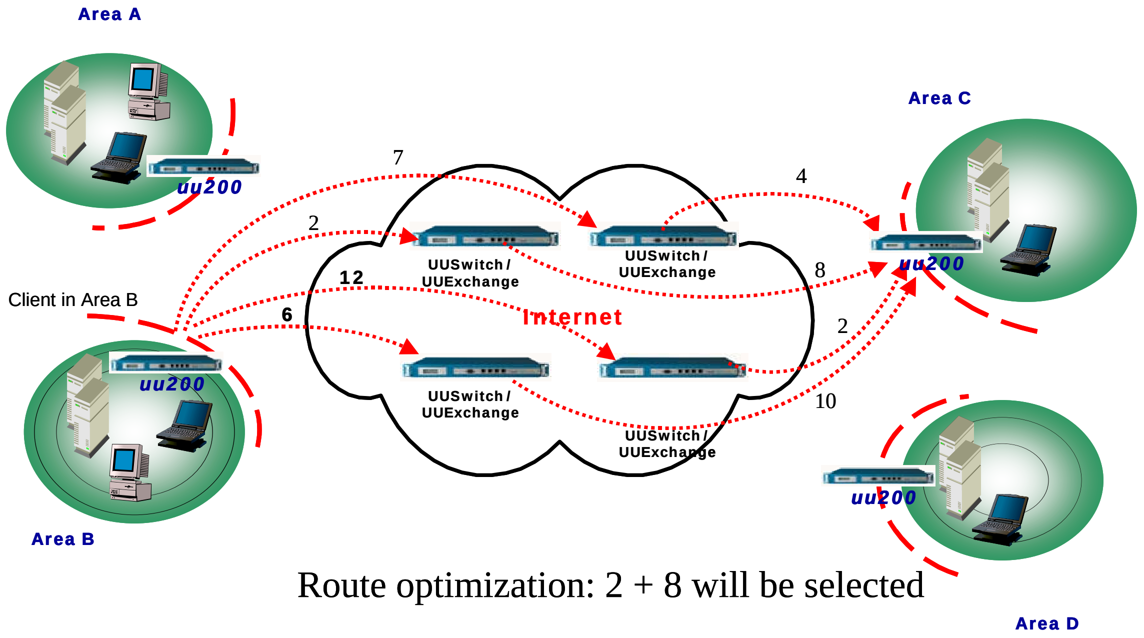

Route optimization: 2 + 8 will be selected

Support for any transmission medium, wired or wireless

End to end security not only protects communication seamlessly without information leakage in between, it also greatly simplifies security protocol. A single unified iSTAR protection covers end to end. There is no longer any need to deploy three different security schemes: one for the segment from client to the AP (Access Point), one from AP to remote gateway, and one from the gateway on the perimeter to where the servers are.

Presence aware VPN

Increasing mobility demands the easy locating and reaching of the mobile work force, where the “presence” based solution, popularized by IM (Instant Messenger), is considered the next killer application. iSTAR provides an efficient platform for exploring the benefits of presence-aware VPN.

More information

Thanks for taking time to read our presentation. For more information about iSTAR products, please visit our web site at http://www.uudynamics.com