Technology Tutorials

VIRTUAL PRIVATE NETWORKS TRANSCRIPT

Program and Presenter Opening

Bob Scarbrough: Hello and welcome to this “Cisco on Cisco” Tutorial on the Remote Access Extended Workplace program and Enterprise Class Teleworker.

My name is Bob Scarbrough, an IT Program Manager within the Cisco on Cisco team, and I’ll be your host during this session.

The theme of our session today is an overview of Cisco’s virtual private network design and our use of the Cisco Enterprise Class Teleworker or, we’ll also call it the ECT solution. You will find this an informative, a very informative and interesting topic with valuable information, so get comfortable, and let’s go ahead and get started.

It’s my pleasure to introduce the guest of today’s session: Dr. Plamen Nedeltchev, a Senior Member of Cisco IT’s Technical Staff. Plamen, thanks for coming today.

Plamen Nedeltchev: Thank you, Bob. Thank you for having me here, it’s a real pleasure to be here and to be able to present the Remote Access Extended Workplace program of Cisco IT, specifically it’s core, which is as you called it, Cisco Enterprise Class Teleworker, ECT.

Bob Scarbrough: Thank you Plamen. The program intends to help build sales lead generation, offer customers best practices that they can implement in their own network, and build some confidence with customers by offering real-world experiences from Cisco IT.

By watching this program, you’ll learn about Cisco’s worldwide deployment of the next generation Enterprise Class Teleworker solution; this is also referred to as ECT, which we will refer to many times during

the next hour.

We'll also understand the end-to-end VPN model of Cisco IT and its components. We're going to learn how to come up with a manageable total cost of ownership, which is very important to our customers, and gain an understanding of the next generation ECT, as an enabler for services, unified communications, and collaboration tools.

See how remote network access and VPN solutions have evolved, from business convenience, this is important to note, to business critical, and from technology to a service, and from a remote access solution to a remote unified communications.

Agenda

Plamen Nedeltchev: Okay, thank you, Bob.

Here is our agenda for today's program. We will begin with describing the market challenges and demands, RAEX programs and its components. And then we're gonna talk about the total cost of ownership, and how to maintain a manageable TCO.

Next, we're gonna go and discuss the next generation Enterprise Class Teleworker, which is the foundation of the remote Access network platform. In this section, we are gonna talk about end-to-end VPN model, security, connectivity, provisioning, deployment, and management.

And finally, we're gonna go to talk about the transition from Enterprise Class Teleworker, as the network as a platform, to the service-oriented architecture over VPN. We're gonna talk about QoS, IP SLA, and lessons learned.

Enterprise Class Teleworker (ECT)

Bob Scarbrough: Well Plamen, you gave us an excellent overview of what we're gonna be talking about. And I'm sure a lot of people are asking themselves: What is ECT, this Enterprise Class Teleworker? I know it's a very exciting technology that's been evolved here at Cisco, and why don't you take a few moments and explain to us what ECT is.

Plamen Nedeltchev: Thank you, Bob Scarbrough. Actually, it's a very good suggestion to start the presentation with explaining the core nature of the Enterprise Class Teleworker.

Enterprise Class Teleworker, or ECT, the way we're gonna refer to it in the future of the presentation, is taking advantage from the enormous space of the broadband deployment around the globe. It's taking advantage from the latest developments of Cisco VPN technologies, and provides enterprise-class services for home users and their offices with encrypted data, IP telephony, video, and wireless.

It's almost a full replica of the office services, offering near office or equal office user experience for home users in their home offices. And finally, ECT is about zero-touch deployment models, automated management, and manageable TCO.

Bob Scarbrough: Wow. That's exciting. So basically, you're emulating an office, a virtual office, by placing this in a person's house. Correct?

Plamen Nedeltchev: Yes, exactly.

Bob Scarbrough: Excellent.

Plamen Nedeltchev: Absolutely correct.

The Telecommunication Industry Transition and the Broadband Explosion.

The first ten years, the first decade of the new century's marking an event which most of the industry observers are referring to as the broadband explosion. The industry transition from permanent circuits to broadband connection is finally gaining speed. The lead times for permanent circuits for sales offices in the emerging markets, continue to be between three and nine months, the pricing in some cases is a showstopper.

The residential broadband offerings today are ranging from typical 1.5 megabytes per second for DSL services, to 68 megabytes per second for cable modems. Some providers are offering fiber to the offices, ranging to all the way to 50 megabytes per second per user. In the next two to three years, it is expected more ISP providers will reach 15 to 25 megabytes per second offerings in the access layer on their network.

Bob Scarbrough: Oh, that's significant. Are we talking about domestic or international, here?

Plamen Nedeltchev: Probably international and primarily in the metropolitan areas first.

Bob Scarbrough: Okay, thank you.

The Telecommunication Industry Transition and the Broadband Explosion. (Contd.)

Plamen Nedeltchev: The telecommuting staff is expected to continue to grow to up 50 million people by year 2008. And Internet over broadband continues to be a hostile environment, with almost 70 percent of the attacks coming from the Internet. Telecommuting as a trend is not only about productivity and business resiliency, it's about adding another dimension to employees' freedom to balance their personal life and business.

Based on the statistics provided by the Organization of Economy, Cooperation and Development, published this year, in the year 2006 the number of broadband subscribers globally has increased 25 percent, 26 percent, from 157 million December 2005 all the way to up to 200 million people by December 2006.

Bob Scarbrough: And this also actually aligns with our reduction of emission controls, by enabling people to work from home more frequently. Isn't that correct?

Plamen Nedeltchev: Exactly, exactly.

The Remote Access Extended Workplace Program and model is applicable for telecommuter's office, for home office, for branch office, for small and medium business, and for commercial networking. It includes for major programs.

The RAEX Model is Applicable for Telecommuter's Office, Branch, SMB, Commercial Networking

The first program is called Next Generation ECT: provides the platform for enterprise-class services, for home users, and home offices. It addresses the needs for full-time telecommuters, part-time telecommuters, and day extenders.

Bob Scarbrough: What's a day extender?

Plamen Nedeltchev: The person who works more than the expected number of hours per day. *(Laughter)*

Bob Scarbrough: Okay.

Plamen Nedeltchev: So site-to-site VPN over broadband provides the framework for the next generation site-to-site VPNs; the point-to-point connections are no longer the only option in branch-to-branch connectivity.

Teleworker QoS Program is trying to address one of the challenges of our generation of Internet, associated with enabling guaranteed Internet. The challenge here is that we're trying to offer enterprise-class services in full replica of the office environment for home users over media, which essentially does not have any guarantee in terms of services.

Bob Scarbrough: Yeah.

Plamen Nedeltchev: So we are partnering with ISP providers to help them to offer the differentiated services on the access layer on their network, and do the necessary changes and implement the necessary ecosystem components to provide guaranteed services for their customers.

Bob Scarbrough: Okay, so it's not just going to be Cisco here; it's gonna be for all ISP customers, then. Correct?

Plamen Nedeltchev: Yes, exactly.

The business resiliency management is the last program where next generation ECT is positioned as one of the major technologies for crisis management and business continuity management. We're gonna talk about that later.

NG ECT Solution Cisco IOS-Based Site-to-Site VPN

Here is the typical deployment of the next generation ECT solution for a user. You see the router, 871, on the top of the screen. His home gear, computer, PDA, camera, 7964 and 7921, 7985. We do not support yet all the older components, which are shown on this picture, but we essentially are applying a phased approach for starting deploying all these devices step-by-step.

Bob Scarbrough: Okay, so this is an illustration, then, that enables people – gives them an idea of what they could actually use at their house.

Plamen Nedeltchev: Exactly. That's a typical deployment, and today we have multiple users who are using essentially this setup.

Bob Scarbrough: All right.

Plamen Nedeltchev: Every router can – it's establishing three major tunnels: one tunnel to the management gateway, and two tunnels to the data gateways. There is a part of the router, which addresses the needs of the non-trusted users, which essentially we refer to these users as users who have no valid credentials with the corporate.

But essentially this is the – this is how every user is connected to the corporate network. And from our end at Cisco IT, it's using the management tunnel to perform particular management operations over that home's environment and that router, to be able to maintain non-compromised security, to collect data, and await images, and so on and so forth.

Cisco 870 Series Integrated Services Router

This is how the 871-router is configured today. You see on the screen four ports on the back of the router. Two of the ports are associated

with so-called trusted pool, which is /28 and the trusted antennae which is associated with trusted wireless connection. If the user connects to any of the trusted ports or is associated with a trusted wireless, his traffic is gonna be encrypted. And over the encrypted tunnel, send and exchange with the corporate environment.

If the user connects to the non-trusted pool, we associate with that – associate with the non-trusted ports, associated /24 IP subnet, which is natted IP subnet, that particular traffic is gonna be routed over the ISP's gateways and is gonna be – there's not gonna be any interaction between the trusted and non-trusted traffic.

Bob Scarbrough: So this is actually excellent for families that have, that may work in different companies. We don't want them routing them through –

Plamen Nedeltchev: Yes.

Bob Scarbrough: – Cisco's intranet.

Plamen Nedeltchev: Yes, absolutely.

Bob Scarbrough: As well as children, you know we don't want them –

Plamen Nedeltchev: Yes.

Bob Scarbrough: – gaming through the Cisco corporate network.

Plamen Nedeltchev: Exactly. And in that case, you're gonna satisfy the needs of the people who are using the residential broadband used to reach the Internet resources, and the needs of the users who have valid credentials to reach corporate resources in the same time.

But probably this is a good time, Bob Scarbrough, for me to congratulate you for the winning of the recent contest for Cisco best home network for Cisco employees. And for our viewers, Bob Scarbrough won the category of the Best Home Security Network. Can you tell us a little bit more about your home network, and how ECT fits into your home design?

Bob Scarbrough: Sure. First of all, I must say that it was an honor to win this award. It was kind of unexpected, and especially with a company the size of Cisco, with 60,000 employees. And there were several categories, so upfront here we need to indicate that the award that I won was the best home security.

And instead of actually getting into the architecture, I would like to share a little bit about the 871 because the 871-router was the core of the network itself.

Plamen Nedeltchev: Yes.

Bob Scarbrough: And building a surveillance system around it enabled me to have access of viewing both inside my house as well as externally, from anywhere in the world. And just to share a little story that was in the article on Cisco's intranet, and that was I was traveling to London and I happened to check in on the surveillance system. And I saw my wife pulling in and one of my children got out and spit some gum into the bushes.

And I called my wife and spoke to her for a while. And then I asked to speak to my son and said, "Do me a favor, next time don't discard your gum in the bushes." And he paused for a moment and said, "Dad, you're 6,000 miles away; how did you know that?" So. And it was basically built around this 871 infrastructure.

So it's just a – it's an example of some of the things that you can actually do with this technology, so.

Plamen Nedeltchev: Thank you, Bob. Thank you. That was an excellent example.

NG ECT is the RAEX enabler

Now let's take a look to the global deployment for next generation ECT today. Today, we have multiple users like Bob Scarbrough who are subscribed to the next generation ECT. In fact, we're approaching 10,000 users deployed globally, and they're located in 11 different locations. We call them hubs; we have management hubs and data hubs.

The management hubs are fully equipped hubs to maintain management traffic and data traffic. And you see on the screen the management hubs are in blue: San Jose, RTP, Amsterdam, Bangalore, and Hong Kong. And you see the data hubs, which don't have any management capabilities but they manage and they pass the traffic, the encrypted traffic, from users' offices, and they're located in: Richardson, Boxborough, Tel Aviv, Singapore, Sydney, and Tokyo. The size of the deployment is expected to exceed 30,000 people in the next three years.

Bob Scarbrough: And you were just sharing with me a few minutes ago that we're just around 10,000 employees right now that are currently using this Enterprise Class Teleworker –

Plamen Nedeltchev: Yes.

Bob Scarbrough: – Solution.

Plamen Nedeltchev: More than 10,000 subscribed users today.

Bob Scarbrough: Excellent.

Site-to-Site VPN over Broadband

Plamen Nedeltchev: Here is the site-to-site over VPN – I'm sorry, site-to-site VPN over broadband deployment. And the purpose of this particular project as part of RAEX program is to provide a fully integrated, flexible and, secure Cisco enterprise branch architecture to extend the headquarter applications in real-time for remote sites. It should allow the security architecture to integrate the security with the unified communications and mobility under the centralized security management.

This particular project holds the promise of reducing the provisioning lead times. It allows jumpstart of the branch offices, and faster penetration into emerging markets. It allows significant WAN cost and OpEx reduction, and reduces the dependencies on ISPs.

Enabling Guaranteed Internet

This is the project, if you remember the Teleworker QoS Project, which is about enabling the guaranteed Internet. You see on the bottom of the screen, and we're gonna talk about that later, that Cisco uses QoS to mark and match and prioritize the traffic which is coming from the home office router.

However, working with ISPs we expect the ISPs to honor our marking coming from the router, and carry the traffic all the way to the brass layer of their network—otherwise speaking, all the way to edge of their network—or provide a non-contention one-to-one window for our traffic so we could commit particular IP SLAs for our users, and pretty

guarantee for them enterprise-class experience in their home offices.

Bob Scarbrough: So this is what you were referring to a little bit earlier with the partnerships that we're starting to engage with, –

Plamen Nedeltchev: Yes.

Bob Scarbrough: – with service providers.

Plamen Nedeltchev: With big service providers, yes.

Bob Scarbrough: All right.

Plamen Nedeltchev: And the last project I want to talk about, as part of RAEX is associated with communication requirements for business resiliency management.

Communications Requirements for Business Resiliency Management

You see the solutions suite on the left, which represent the software VPN, which is designated to provide data connectivity and voice and data connectivity pretty much based on the best-effort services. And you see the Cisco ECT positioned to provide full office replications for – replication for the unified communications and corroboration tools for key executives, for decision makers, and business critical resources.

All these services are positioned to provide essentially three enabling technologies: VPN, Voice over IP, and conferencing. And allows Cisco, as a company, to prepare for unexpected, for manmade workforce disruption, or social distancing and any other disruption. And essentially, they offer technologies or suite of technologies for the industry to meet this type of life situations.

Bob Scarbrough: Yeah, we actually have some case studies that are posted on our Web site that describe how Enterprise Class Teleworker enabled us to continue working productively through some natural disasters, like with Katrina, the hurricane in –

Plamen Nedeltchev: Yes.

Bob Scarbrough: – North Carolina.

Plamen Nedeltchev: Yes, we have multiple examples of that. You're absolutely right, yes.

VPN Solutions have evolved from business convenience to business critical…

In a nutshell, VPN solutions evolved over the years, from the business convenience to business critical, from technology to service, from remote access solutions to remote unified communications.

The first generation Enterprise Class Teleworker was about to build a platform for the next generations. ECT has proven to be a big cost saver for Cisco IT and Cisco customers. From an industry perspective the ECT, like management security solutions are the preferred ones versus non-managed solutions, due to their specifics and advantages.

Bob Scarbrough: Now I actually recall, to interrupt here for a sec, the first generation ECT did not really have Quality of Service. And just to realize the importance of Quality of Service, I had an IP phone, a 7960, and I could make calls from my house through this hardware-based VPN solution. Although when my children were gaming, there was contention on the –

Plamen Nedeltchev: Yes.

Bob Scarbrough: – same wire and people would hear basically jitter, delays, or uh-uh-uh-uh, and I'd have to put it on mute and tell my boys to turn it off. And

now with the advent of Quality of Service –

Plamen Nedeltchev: Yes.

Bob Scarbrough: – now the boys gaming jitters and stops, –

Plamen Nedeltchev: Yes.

Bob Scarbrough: – and dad gets the voice precedence.

Plamen Nedeltchev: Yes, absolutely.

Bob Scarbrough: So excellent advancements there.

Plamen Nedeltchev: Yes, exactly.

So the network as a platform built several years ago, we pretty much evolved that system into something we call today, we're talking today, Service Oriented Architecture. Where we're essentially, not only providing encrypted data over public networks but we provide services for our users and for our customers. And even further, we're not only providing services or trying to commit IP SLAs for our users so they could have a near-office experience, even on the second phase of our moving towards the next phase of the ECT program.

And probably in the future, we're gonna be moving entirely into equal-user experience from the home offices, into building new business models, deploying new generation services, and Cisco gear enabling mobility and presence, unified communications, and collaborations.

Bob Scarbrough: Excellent.

ECT Reduces TCO

Plamen Nedeltchev: However, when you build this type of systems, you always have to keep in mind something, which is called Total Cost of Ownership. The Total Cost of Ownership, TCO, is a sum of the acquisition costs, of the operational and initial install costs, and the management costs over the lifetime of the asset. The higher the return of investment, the lower the total cost of ownership and vice versa.

Based on the research published several years ago by Sage Research, 20 percent of the TCO belongs to the acquisition costs, 35 percent of the TCO belongs to the operational costs, and 45 percent to the management costs. So, our efforts from the very beginning were concentrated to reduce the operational and management costs and altogether to increase the return of the investment.

We used several different techniques, and we're gonna talk about them during our presentation, and those are: lower the cost of the provisioning, lower the cost of the deployment. I'm gonna show you analysis which shows that IT was able to save up to 12, between 12 and 14 percent from the total cost of ownership only using zero-touch deployment technologies to deploy the CPE, the remote routers were used lower cost of management and utilizing usable components.

Bob Scarbrough: Something that's really exciting about this technology that we're not sharing in this presentation. I'd like to ask you: How many people are supporting this ECT solution that has approximately 10,000 people to date?

Plamen Nedeltchev: I probably don't have the number of the people but we have three layers of support. And from a support perspective, what we are observing from the very beginning we had significantly low number, low

percentage of support – from support, from the frontline support for our program.

Over the years, they became very efficient in their support efforts. Today the frontline supports more than – 91 percent of the cases because we are able to automate many of the routine procedures, management procedures, that allow our staff to get faster and more proficient to support the open cases and address the users' needs.

Bob Scarbrough: Okay, thank you.

NG ECT Solution: End-to-End VPN Model and TCO

Plamen Nedeltchev: Now let's talk about the end-to-end VPN model and the total cost of ownership.

NG ECT Solution: End-to-End VPN Model and TCO (cont…)

You see on the screen the next generation ECT and end-to-end VPN model. And the total cost of ownership includes four major components; we call them pillars. There's four pillars essentially contribute to something we perceive as the network as a platform.

The first pillar is end-to-end security, end-to-end connectivity, end-to-end deployment, and end-to-end management. These four pillars, they're end-to-end inside each one and they're fully integrated with others. For example, end-to-end security is based on the layers of IOS security, and we're gonna talk about that briefly in a moment.

End-to-end connectivity is associated – it's associated with the integration of the latest development of VPN technologies of Cisco and in particular, DMVPN, which is the core of deployment, and integration into the dynamic routing protocols framework of Cisco. We're gonna talk about the end-to-end deployment and end-to-end management.

Let's remember this diagram because for us, this is the key to maintain total cost of ownership. And we want to talk today less about security and connectivity; we want to talk more about the deployment and management. However all these components, all together with support model, essentially represents something we today refer to as a network as a platform.

Based on this platform, we could make the next step and start adding services to our offerings and making the next step to the service oriented architecture.

Bob Scarbrough: Excellent. Also important to note, recently there was a design guide, I believe, that was published with regards to Enterprise Class Teleworker.

Plamen Nedeltchev: Yes.

Bob Scarbrough: So for those of you that are interested in talking or taking this to a very technical level, the design guide would be a good reference for you.

Plamen Nedeltchev: Yes, you're right. Thank you.

End-to-End IOS Layered Security

End-to-end IOS layers of security, we're gonna talk briefly of some of the features but here are the major features we use or can be used as part of ECT.

RSA key loss due to password recovery, secure RSA private key, secure authentication proxy, 802.1x, Cisco IOS PKI support and PKI-AAA integration, Cisco IOS stateful firewall CBAC, Cisco IOS IPS,

network-based application recognition, NBAR.

RSA Key Loss Due to Password Recovery

The RSA key loss due to password recovery includes two major features. If someone attempts to perform a password recovery on the router, the RSA private keys will become unusable. If a user tries to change the host name of the router, the RSA private keys are permanently deleted.

In this way, we protect the router and we prevent the user to establish a VPN security session if the password recovery is performed.

We want to maintain the uninterrupted access and control over the security policies, over the audit policies, over the monitoring, and in the same time guaranteeing the user from point of view of the latest IOS, patches, protection from denial of services attack, etc., etc. This is essentially pretty much what we call managed security.

Secure ARP

Secure ARP feature is essentially when the spoke router assigns an IP address via DHCP, the entry is secure in the ARP table, and an intruder cannot just clear the ARP cache and use the IP address to gain access to Cisco network. This is one of the spoofing techniques to protect the router. However, the better way probably is gonna be to use certificates, and we're gonna talk about that in a second.

Authentication Proxy

The authentication proxy is a user authentication technique on the third layer of the network stack. If the user authenticates – if the user tries to reach the corporate resources, he's gonna be prompted for username and password. Upon successful authentication, his computer, or PC or laptop or IP address, is gonna be granted access to the network. But if we have an IP phone connected, because the IP phone does have not an effective browser, we're gonna make sure the phone is not gonna be prevented from connecting to the Call Manager, bypassing the authentication proxy.

Bob Scarbrough: And if somebody plugs a laptop or a PC into the backside of the phone, –

Plamen Nedeltchev: That's a very good question, actually.

Bob Scarbrough: – they're protected right?

Plamen Nedeltchev: Yes because the phone has an external additional – not external but additional switchboard, and that port is gonna facilitate another IP address for that particular computer. However, we're gonna challenge the owner of the computer, that browser, to provide another set of username and authentication credentials. And upon failure, that access is not gonna be granted. If that authentication is successful, that computer is not gonna bypass the authentication proxy challenge but is gonna be granted access to the corporate resources.

Bob Scarbrough: Excellent.

Plamen Nedeltchev: Authentication proxy can be implemented as a mechanism to prevent non-authorized user from accessing the corporate network. However can be implemented as a mechanism to prevent users from connecting to the corporate if, for example, a router is stolen or it's out of the region control of the user.

Bob Scarbrough: Okay.

IEEE port authentication-802.1x

Plamen Nedeltchev: The better way to approach the same user authentication paradigm would be to use the IEEE port authentication associated with 802.1x. 802.1x is essentially a technique, which simplifies significantly the configuration. And if you remember the trusted and non-trusted concept of configuring a 871-router, today we pretty much assume the user is gonna connect to the trusted port or associated with the trusted wireless.

If we have 802.1x configured, then the user can connect to any port or associate with any wireless, when that happens he's gonna be prompted for 802.1x authentication. Upon his successful authentication, he's gonna be assigned a trusted IP address. If he fails, he's gonna be assigned a non-trusted IP address.

Bob Scarbrough: So it simplifies the configuration –

Plamen Nedeltchev: Exactly.

Bob Scarbrough: – and the administration of this, as well?

Plamen Nedeltchev: Yes. And makes it very, very easy to monitor in cohesion with the whole entire concept of user authentication across the board with other segments of the network.

Bob Scarbrough: Good.

Cisco IOS Certificate Server PKI-AAA Integration

Plamen Nedeltchev: Cisco ECT is using extensively the Cisco outwards based certificate server PKI-AAA integration. Cisco PKI solution provides the classical, necessary components of PKI which are encryption, confidentiality, and non-repudiation. And PKI addresses in general the problems associated with the man in the middle (?) attack.

IOS CS supports CA, RA, and sub-CA server modes. It supports exportable and non-exportable keys, full backup, restore, and out enroll. Permanent storage of certificates on the external database and local flash, and eliminates the need of managing of certificate duplication lists, which can be very management heavy, and allows us to manage the environment much easier and much simpler way, based on the PKI-AAA integration.

Cisco IOS Firewall Features

Cisco IOS firewall features provide stateful firewall and CBAC for our routers, the firewall echo will block the non-authorized access inbound attempts from the Internet.

Bob Scarbrough: Is trusted and the un-trusted?

Plamen Nedeltchev: Actually we provide mainly the trusted. However, CBAC, since applies to the outbound interface, protects the non-trusted from unauthorized access from the Internet; and not only that, but on the non-trusted side of the router, we apply a network address translation. It essentially have double layers of – dual layers of protection for even for non-trusted users, this way protecting spouse and kids living in the same home and using the same home office just to access Internet from being attacked from the Internet.

Bob Scarbrough: Good to know.

Plamen Nedeltchev: So CBAC will open temporary some application associated ports for the return traffic, and upon expiration of the timeout those ports are gonna

be closed. Apart from the standard TCP and UDP ports, CBAC supports protocols like SIP, Skinny, SMTP, ESMTP, FTP, and others.

Network Based Application Recognition (NBAR)

The Network Based Application Recognition solution, NBAR, is an intelligent classification engine that can reorganize the applications, including Web based and client server applications, with dynamically assigned TCP and UDP port numbers.

The next generation ECT NBAR is used to match and remark the time-sensitive traffic like IP, IPT, video, IP communication, at the egress interface and the queue, and prioritize the traffic based on this marking. In such way, next generation ECT changes the status of the traffic from non-trusted to trusted, and allows the time-sensitive application to be routed in the corporate network in a cohesive way with other time-sensitive traffic.

Maybe the major advantage of NBAR is its ability to classify, mark, and match business critical applications and guarantee bandwidth for them. Overall, NBAR improves the VPN performance by ensuring identified mission critical and time-sensitive applications before they're encrypted allowing the network to apply appropriate Quality of Service controls.

Bob Scarbrough: How does it identify, just curious how it identifies and prioritizes that, the application?

Plamen Nedeltchev: Based on some indications of the payload, they're multiple ways that NBAR can essentially identify traffic. But identification is happening on the application layer instead of being – it happening by the QoS, which is matching the traffic on the second and the third layer of the OSI stack.

Bob Scarbrough: Right, okay.

NG ECT Solution: End to End Connectivity

Plamen Nedeltchev: So let's talk about now end-to-end connectivity, which is the second pillar for end-to-end VPN model.

End-to-End Connectivity

The end-to-end connectivity's built primarily on the DMVPN, and we're gonna talk about the DMVPN fundamentals. Let's talk about routing with DMVPN, DMVPN key differentiators. Let's talk about server load balancing, overall design, DMVPN and SLB design, and SLB DMVPN key advantages.

DMVPN Fundamentals

DMVPN stands for Dynamic Multipoint VPN. It's a Cisco IOS-based solution, which integrates the Cisco VPN Solutions with Cisco Dynamic Protocols Framework. It provides failover for the tunnels, for the encrypted tunnels, provides load balancing, and SLB designs, we're gonna talk about SLB in a bit, provides dynamic routing, full-mesh and partial-mesh topologies, hub-to-spoke and spoke-to-spoke tunnels, permanent and on-demand tunnels.

DMVPN includes three major components: IPSec, based on RFC2401, Next Hop Resolution Protocol, NHRP, which maintains a database of all spoke routable public interface addresses. Each spoke registers his routable address with the NHRP server, after successful negotiation of the IPsec tunnel. Spokes is gonna – they're gonna query the NHRP database for routable addresses, the destination spokes, to build direct tunnels. The third component associated multipoint GRE tunnel

interface, with allows GRE interface to support multiple IPsec tunnels.

Bob Scarbrough: So this is a significant enhancement then moving forward with NBAR.

Plamen Nedeltchev: Yes, this is significant enhancement from the typical point-to-point VPN connections, the way they were defined around 1998 and '98. DMVPN essentially expands the notion of end-to-end secure connection to the end-to-end VPN connection, which is significantly different and significant enhancement from the common understanding about point-to-point VPN encrypted connections.

Standard DMVPN Design

Here is the typical standard DMVPN design. Every router, you see all the spoke routers on the bottom of the screen, the corporate firewalls, and the DMVPN hubs. Every spoke router is building one management gateway, you see it on the center of the screen, and that gateway's gonna be built using one digital certificate.

In the same time, the second digital certificate is gonna be provided to the router to build a dual-tunnel, single-layer configuration. Primary tunnel is gonna be connected to the primary secure data gateway, and the secondary tunnel is gonna be terminated to the secondary secure data gateway.

Bob Scarbrough: These corporate firewalls here, are they in the intranet or are the in the DMZ?

Plamen Nedeltchev: On the DMZ. Pretty much associated with DMZ.

So the system is gonna work in a way that the primary and backup – they aren't gonna be primary and backup tunnels, they're gonna be primary, secondary, both inactive-active relationships. So, when the primary tunnel goes down, the secondary tunnel takes over in a matter of ten seconds. When the secondary tunnel goes down, the primary can take over in less than one ping, which is about two seconds. This particular design creates significant and predictable connectivity and performance environment for the home users, so from their perspective there's no interruption of their service whatsoever.

DMVPN: Key Differentiators

The key differentiators of DMVPN are the following. DMVPN uses crypto profiles and tunnel protection, and frees the physical interface from the crypto map, which is typically associated with a problem when you want to enforce different security policies. The management's performed over a separate VPN tunnel, independently from when the primary and secondary DMVPN data tunnels.

DMVPN allows dynamic registration of the spokes. One tunnel interface on the hub supports a single DMVPN class, eliminates the static point-to-point configurations, and reduces significantly the complexity of the configuration. DMVPN provides dynamic full- and partial-mesh capabilities providing improved support for application as voice and video.

One example about the simplicity of the configuring router with the DMVPN, if you configure a router with the standard traditional crypto maps, probably you're gonna need about 13 to 15 configuration lines per spoke router. If you assume 700 routers connected, 700 spokes connected to that router, you probably need to configure more than 10,000 configuration lines. The same configuration can be achieved with DMVPN, configuring probably something between 200 and 300 lines.

Bob Scarbrough: Oh excellent.

Server Load Balancing (SLB) Overall Design

Plamen Nedeltchev: As part of the next improvement of the DMVPN technology, Cisco offers the new design which is called server load balancing. This design applies today to the next generation ECT hubs, primarily the big ones, which are in: San Jose, RTP, Amsterdam, Bangalore.

In this design, the spokes connect to a load balancer, which is in the center of the whole design, and the load balancer connects to multiple – to a cluster of hubs typically associated with the farm. This design can be delivered in two different ways.

DMVPN and SLB Design

Design 1 is called DMVPN High Concentration Hub, and it's based on Cisco 7600 series router or Cat65, which act like a primary tunnel termination hub and perform the encryption and the encryption operations, which are typically very expensive. So to say: operations for establishing VPN tunnels. A farm of 7200 series routers are associated with the IPSec termination device and they handle all the tasks related to NHRP, mGRE, and the routing protocols.

The second design is called DMVPN IOS SLB hub. In this design, the farm device does not perform the encryption and pretty much it's standard 7206VXR router but performs the role only and solely of the load balancer, which regulates the traffic to the farm device, which has the lesser load in this particular moment when the VPN connection is initiated from the spoke end.

In the design, a farm of island set of 7200-series routers are connected to the termination to the SLB design, and they perform the functions of NHRP, mGRE, IPsec, and routing protocols.

Bob Scarbrough: Curious what a customer – why a customer would pick one design over another.

Plamen Nedeltchev: Very good question, actually. The way we position those two designs are pretty much for large scale deployments, we pretty much have a dual-head single-layer design based on the DMVPN high concentration hub. For smaller hubs, we pretty much decide to go with either DMVPN IOS SLB design, or with the standard DMVPN design which is fully capable of handling all the way to several thousand users per hub.

Bob Scarbrough: Okay, good explanation.

SLB DMVPN-Key Advantages

Plamen Nedeltchev: The key advantages of SLB are on the screen. SLB's much easier to configure and support, sees the configuration of the peer tunnel IP address is always the same no matter how large is the deployment. SLB scales better, sees the EIGRP base scalability restrictions are mitigated, and the number of tunnels is virtually limitless. SLB provides high tunnel creation rate, recovers faster when cluster mode becomes unavailable, and provide spoke-to-spoke functionality, as the standard DMVPN does.

SLB DMVPN-Key Advantages (Contd.)

SLB provides better redundancy. The standard DMVPN design provides redundancy in pairs. The dual-tunnel single-layer design, the one we were showing earlier on the screen, actually terminates the CP to two separate STGs, maintaining active-active status. In this case, you always double the number of the hubs if you need to increase the

size of the deployment with another X amount of spoke users. Everything equal in SLB design, we assume the same number of spokes per hub, the number of the hubs in the design should be adding one extra hub to the existing hubs to address another 500 or 600 users and allow them to connect to the head-end.

Bob Scarbrough: What triggers the secondary hub to connect? In other words, if one goes down is that the only that one would connect to the other server?

Plamen Nedeltchev: Actually both hubs are passing traffic. In the standard DMVPN design, both hubs are passing traffic; however, the first one is preferred based on the metrics configured on the EIGRP. When the first hub becomes unavailable, the metrics calculation is gonna show that that destination is unreachable. So that's how the failover is gonna be performed based on decisions made from the dynamic routing protocols.

Bob Scarbrough: So for an enterprise company, would they distribute these across the campus or in different locations so you don't – just so if there's a power outage or something like that it doesn't affect them?

Plamen Nedeltchev: Actually this is how the – where the DMVPN is going right now. If initial DMVPN started with connecting a dual hub and geographically collocated them, today DMVPN can be essentially deployed in two different buildings in the same campus. However, SLB design provides even fully redundant solution where the dual SLB design can connect to pair of farm hubs, which are not geographically collocated.

In other words, in extreme situations it can allow the remote routers to failover to another hub located in another part of the same campus, or located in another geographical location altogether. In this case, we could talk about fully redundant solution, if one campus goes down another one takes over.

Bob Scarbrough: Excellent.

NG ECT Solution and Low TCO- End-to-End Provisioning with Cisco Security Manager

Plamen Nedeltchev: Now after this brief introduction to the security and connectivity features of next generation ECT, let's talk a little bit more about the end-to-end provisioning with Cisco Security Manager.

If you remember we – in the TCO diagram, initial deployment and provisioning takes about 35 percent from the total cost of ownership. So, it's very important to address this particular part of the TCO if you really wanna address the TCO concerns and increase the return of investment.

Cisco Security Manager

The whole concept of building of the deployment and provisioning and management is built around several management platforms. And probably the most important one, which is performing the configuration and provisioning, it's something – it's Cisco Security Manager. Cisco Security Manager, CSM, manages devices, PIX, firewalls, ASA, firewall SMs, and Cisco IOS routers. It manages transfer mechanisms such as SSL, Telnet, HTTP, HTTPS, TMS, and Cisco Networking Services working with configuration engine 2.0.

CSM manages policies, activities, and objects. It manages site-to-site VPNs, remote access VPNs, SSL VPNs, and EZ-VPNs. CSM manages firewalls, firewall services. Manages firewall related policies and security managers, the manager that applies to the adaptive service appliances, PIX, firewalls service module installed in Catalyst 6500, and

security routers running Cisco IOS.

Bob Scarbrough: Plamen, how many CSMs do we have distributed across the enterprise?

Plamen Nedeltchev: Very good question. As part of our deployment, essentially we position one primary and one backup CSM for every single management hub. And assuming we have five management hubs, we are even now positioning ten different CSMs, Cisco Security Managers, around the globe.

Bob Scarbrough: Okay.

Cisco Security Manager (Contd.)

Plamen Nedeltchev: In general, if you look through the feature set of CSM, where it manages IPS, supports XML ____ interface, and outbound APIs to enable integration with the existing enterprise management framework. Supports fully managed services functionality to not fight administrators for non-CSM initiated configuration changes. And finally, CSM manages provisioning, manages deployment, and manages flex-configs. All together, this feature set essentially represents something, we'll call it CSM is positioned to managed security policy of ECT solution.

Configuring CSM – The Sample Device and the Security Policies

This is how the provisioning starts. Initially, we create a sample device, which can be created in four different ways: can be added from the network, can be added from the configuration file, can be added from the repository, or can be manually configured. It's very interesting feature about CSM, it's essentially intelligent to pause the configuration if you upload it from the configuration file or right from the network and you need it to populate all the policies throughout the whole CSM system.

Configuring CSM – The Sample Device and the Security Policies (cont…)

After that, you have several very simple steps to complete the initial provisioning of the CSM – of the sample device. You configure the firewall policies, including triple access rules, access rules, inspection rules, access control, authentication proxy, and inspection. Then you configure site-to-site VPNs.

In this case, we have four hubs, which are using SLB, so we configure SLB, and several other hubs, data hubs, which use standard DMVPN design, so this is where you configure the site-to-site VPN. Then you configure the Quality of Services, and then you configure NAT. And you're pretty much done with the standard set of policies for the sample device.

Bob Scarbrough: That's a nice graphical representation. Are you gonna be discussing, kinda describing the process of how it starts, as well?

Plamen Nedeltchev: Yes. Actually before I respond to your question, I want to essentially make a comment here that it's absolutely necessary and it's very, very important, too, as part of the toolset, that with CSM is something associated with a template management system, but it's called flex-config, configuration, because the managing the security is only part of the managing of next generation ECT.

Building the next generation ECT as a service-oriented network requires for us to think about deploying not only security policies but deploying services, as well. This is where the flex-config toolset becomes very handy, No. 1. No. 2, it's very important to allow Cisco customers to

integrate CSM into their existing environment, infrastructure, which essentially can differ significantly from one another.

So the flexibility of the flex-configs, allows these customers to incorporate, merge, or essentially fit the CSM into their existing environment and use the benefits and the strength of the system, adding to the existing components of their enterprise management system.

Bob Scarbrough: Interesting.

Configuring CSM – The Sample Device and the Flex Configs

Plamen Nedeltchev: Here we add multiple configs, flex-configs, to the sample device. We attach pre-pin configs, basic configs, wireless configs, IP telephony configs, video configs, and append configs. So with this, pretty much the configuration of the sample device is completed.

Cisco Security Manager 3.1 SLB Hub Configuration

Before we go any further, we configure the hub device, which is very simple to do. There are four simple steps when you configure the hub device, or several of them. You configure the SLB device, if that design requires SLB design; create hub and spoke VPN; you assign several spokes to that particular policy, and pretty much select which device is gonna perform the role of SLB.

User Request for NG ECT Service

Back to your question, Bob Scarbrough, right now let's see how the whole process works.

Bob Scarbrough: While you're taking a sip, I wanna take a moment to share something. This is a highly complex, comprehensive solution, and it's transparent to the end-user.

Plamen Nedeltchev: Correct.

Bob Scarbrough: What I like about this solution is that a person, once authorized, entitled to purchase the system or get the system, they will order the router; it'll come to their house. That's why you call it the touchless-config; it doesn't have to come to IT first.

Plamen Nedeltchev: Yes.

Bob Scarbrough: And that's how you're able to scale this out to up to 30,000 employees.

Plamen Nedeltchev: Yes.

Bob Scarbrough: It's a very simple transparent solution for the end users.

Plamen Nedeltchev: Yes.

Bob Scarbrough: And you don't have to be technical.

Plamen Nedeltchev: Yes, exactly.

Bob Scarbrough: That's –

Plamen Nedeltchev: But we're gonna talk about the deployment here in a moment.

Bob Scarbrough: Okay.

Plamen Nedeltchev: I just wanna make a step back and talk about the preparation of the deployment, which is essentially associated with the provisioning.

The user is gonna be provided with URL he wants to – the service is optional. So when the user wants to apply for ECT, he's gonna go to the Web, he's gonna apply the URL, and he's gonna see something associated with user service request page, or otherwise speaking this is the dashboard of the user.

Bob Scarbrough: Okay.

Plamen Nedeltchev: He is gonna request the service, he's gonna provide his data, his provider, his download speed, upload speed, the way he's gonna obtain IP address. Every other piece of the information is gonna be populated from the existing databases, and pretty much all the user needs to do is to click yes and request.

Eleven Steps to Provision and Deploy a Remote Router

At the moment, his service request is gonna go to the requested state. Immediately, email is gonna go to his manager for approval. EMAN will create a triple icon for the user, to facilitate in the future the PKI-AAA configuration. EMAN is gonna assign a /28, the trusted subnet for the user, for using the address management agent.

Eleven Steps to Provision and Deploy a Remote Router (Cont…)

It's gonna associate the user, is gonna create a host in the backup system, associate the host with the template management. It's gonna assign another /32 IP address for the tunnel interface of the user. Provide a TFTP address for this user; if the user requests an IP telephony service, and then is gonna move through the six single sub-steps of configuring of CSM. We're gonna get into that in a moment.

Bob Scarbrough: The one thing that you mention that's important to note here is you stated "EMAN." And that's an acronym that we use for a solution, a propriety-based system called Enterprise Management we have here –

Plamen Nedeltchev: Yes.

Bob Scarbrough: – at Cisco.

Plamen Nedeltchev: Yes. Yes, exactly. But in the same time, we could assume that every big enterprise the size of Cisco, even smaller, is gonna have some type of management system we could still refer to as the enterprise management.

Bob Scarbrough: Okay, good.

Plamen Nedeltchev: And again, the integration into any enterprise management, it's possible due to specific features of CSM associated with using the Java APIs or XML sort APIs or using flex-configs to be able to merge –

Bob Scarbrough: Integrate.

Plamen Nedeltchev: – and integrate –

Bob Scarbrough: Excellent.

Plamen Nedeltchev: – existing enterprise management system.

Bob Scarbrough: Important to know.

Plamen Nedeltchev: Yes.

The process is gonna complete with the CNS configuration – with the configuration of the user station, the CNS engine, and pretty much at that moment the user is gonna be advised to order his equipment from

the factory.

Cisco IT Implementation-CSM Integration using APIs.

Here is how CSM is positioned in the Enterprise Management System we were talking about just with you all. So if this is the Enterprise Management System, on the right you see the CSM, which essentially the EMAN system sends all these parameters, we're talking about, EMAN performs those six configuration steps, we're gonna talk about in a second, and returns response back with the status of the provisioning process.

Cisco Security Manager 3.1 6 Easy CPE Provisioning Steps

This is how the provisioning process starts. You will choose to clone a sample device. Then we set the device properties. We set the device interface rows. We set the network and host rows. We set the device properties, in terms of text objects, customization, upload speed, download speed, CNS port numbers; there are several variables to be set. At the end, we added the QoS policies to reflect the user requirements in his reports based on specifics of his broadband service. At the end, we submit and deploy the configuration.

The moment the configuration is staged on the CNS engine, the provisioning essentially is completed. You see all these jobs waiting on the CNS engine to be deployed. Meanwhile the user is gonna follow direction from IT and order his equipment, which is gonna arrive to his house, without any preps, without any specific configuration, straight from the factory.

NG ECT Solution and Low TCO- End-to-End Deployment with Cisco Security Manager

Let's talk now about the real deployment: How the user is gonna deploy his router. Everything gets provisioned, everything is staged, the user orders his router, the order receive his router. Let's see what happen next.

Bob Scarbrough: Excellent.

Conventional Provisioning and Deployment of CPE/Spoke Routers

Plamen Nedeltchev: Based on our analysis in the industry, the way the industry deploys routers today we could say that there are three major conventional ways. In-house, the one you mentioned when essentially you ship certain amount of routers for users to your IT department, they configure the router and help the user to stage it in their home.

Bob Scarbrough: Which doesn't scale too well?

Plamen Nedeltchev: Exactly.

The second one is so-called staging facility or outsourcing to ISP, when essentially the ISP performs that job on behalf of the enterprise for additional fee. And the third option would be the third-party. When the third-party company performs a manual configuration of the router, stages the router in users home, and essentially makes sure that the router is connected and the user can conduct business.

Bob Scarbrough: So I got to ask: How does Cisco IT deploy the spokes?

Plamen Nedeltchev: Actually, we don't do any of any these practices any more. Actually, those are conventional ways that traditionally exist in the industry, we

support them but we don't do that in our environment. The reason is that they – all of them, they add additional cost on the top of the provision cost, so they could make essentially the total of cost of ownership unpredictable to a certain level.

NG ECT Offers Four Deployment Options

What we do is associated with something we call "Zero-Touch" Deployment. It's a set of models, which allow the user to self-deploy himself without opening a case, without knowing anything specific, without being trained. The only thing we ask the user to essentially to be able to manage a browser and provide the username and password.

Upon authentication, the process happens on the backend, the user is just a witness. Several minutes later he is gonna be prompted to authenticate his browser, and after that he's gonna be pretty much able to connect to the corporate. How exactly that's gonna happen, in a second.

The second scenario we're gonna – we're working is called Certification Proxy. In this scenario, in certain situations and as necessary, the IT engineer to be able to obtain certificates and key material on behalf of the user, and be able to place that configuration over encrypted SSH tunnel to his router. That's gonna bring his management tunnel up and the whole process is gonna follow the ZTD model after that.

The third option: using the e-token, small USB e-tokens, which are gonna, carry the certificates and the key material, and they're gonna facilitate the building the management tunnel.

The last model I want to talk about is offline, in some special cases new implementations. Some specific pilot environments we still configure manually, but this is just for development purposes, not for deployment purposes. And regardless of the deployment options, spoke routers provisioning processes we automate it to minimize the total cost of ownership.

Bob Scarbrough: Okay, so we're offering various solutions here to customers –

Plamen Nedeltchev: Exactly.

Bob Scarbrough: – based on their needs.

Plamen Nedeltchev: Correct.

Bob Scarbrough: All right.

ECT CPE ZTD Deployment

Plamen Nedeltchev: So here's the ZTD model works. You see on the top of the screen, the home environment and the router. When the user pastes his URL, he's gonna be prompted for username and password. If he passes successfully, that upon authentication, we're gonna send to the router a small bootstrap configuration, which is gonna include all the key material and all the certificates sufficient for the user to build his management tunnel.

The moment the management tunnel comes up part of that bootstrap configuration carries a CNS agent. CNS agent is gonna be activated, it's gonna send connective upstream. When that happens, we're gonna push all the policies, all the configurations sitting on this, staged on the C-engine and they're gonna be downloaded in a matter of a less than a minute. As a result, we're gonna build either, in the standard DMVPN design, we're gonna build primary and secondary DMVPN tunnels or we're gonna build a single primary SLB tunnel to one of the farm

devices.

With that, pretty much the configuration is completed. All the user has to do is release and renew his IP address, and he could reach the corporate resources.

Bob Scarbrough: That's a good review.

ECT Architecture-Today-Auth Proxy

Plamen Nedeltchev: However, we had to address again the spouse and kids dilemma. We have a non-trusted subnet, we have a non-trusted port; we have to make sure that only users with valid credentials will be able to reach the corporate resources. For that, you see on the screen we're gonna use authentication proxy.

ROI, TCO and ZTD Cumulative Cost Savings

Let's see what is the cumulative effect from using zero-touch deployment as part of our deployment the next three years. Based on the analysis performed by our business analysts, for the next three years we're gonna be saving between about $3.6 million from the total cost of ownership which represents between 12 and 14 percent of the total cost of ownership, only by using zero-touch deployment for our technologies. Let's remember this number and we'll come back to this later.

Bob Scarbrough: And that doesn't include the reduction in mobile calls because if we're providing, enabling people to use their IP phones from home for international calls –

Plamen Nedeltchev: Exactly.

Bob Scarbrough: – that doesn't take into consideration.

Plamen Nedeltchev: Very good point. That's not even in the calculation today.

Bob Scarbrough: Okay, good to know.

Plamen Nedeltchev: However, we're committed to zero-touch technologies, and I'm gonna explain in a minute how we're gonna approach other services using this same concept we use today for deploying the remote router.

Bob Scarbrough: Okay.

The User's Experience

Plamen Nedeltchev: Let's talk about the user experience. Here is a quote from one of our users, who essentially sent us the following email: "I wanted to let you know my firsthand experience with the new ECT router and getting it setup. I dreaded the process. My last ECT router was shipped to me with a very large book of how to configure it. The new router sat in a box next to my desk for about four days because I was planning to dedicate a full weekend to the process of hooking it up and getting it configured."

Bob Scarbrough: You could tell it's a non-technical person.

Plamen Nedeltchev: Yes, that's true.

"Well much to my surprise, I hooked everything up, including my home equipment, and I had the new router configured in 15 minutes. Let me repeat that: 15 minutes. The instructions on the Web are and the printed material was easy for a non-technical person to understand. The router was setup to be configured and connected to the site; full

configuration was easy. Voila, 15 minutes later I'm back in business. It even amazes me that I was able to do that without hassle. Now back to work but from home, sincerely yours, Pat Moore, Manager, Workplace Resources."

Bob Scarbrough: Yeah. So, this is also increasing productivity of employees. I know it does for me 'cause I'm able to – our organization is global so we can make our calls to Asia and Europe either early in the morning or late at night.

Plamen Nedeltchev: Exactly.

Bob Scarbrough: I'm not sure our spouses like that too much, but –

Plamen Nedeltchev: I'm not sure either. *(Laughter)*

NG ECT Solution and Low TCO-End to End Management with Cisco Security Manager

Let's talk about now about the management and using Cisco Security Manager.

If you remember correctly from the total cost of ownership model, the management takes about 45 percent from the TCO. So addressing the management concern, day-by-day operations, and making the – reducing the burden on the head-end, managing the environment, is essential part of our effort to reduce the TCO.

TCO and Lower Costs of Management, TCO and Utilizing Reusable Components

Our efforts are in the following directions: First of all, we integrated a CSM and CNS engine into the enterprise management. If we step back and remember that we have the XML soft APS available and we have the flex-configs. That was part, which essentially allow us to migrate and work those two management platforms into the enterprise management system.

Bob Scarbrough: And customers could integrate this CSM and CNS into their management systems.

Plamen Nedeltchev: That's correct, yes.

So we are using monitoring, which is pretty much EMAN-based. We are using a significant amount of analyzing and grouping, based on static and dynamic grouping of the routers. We have the first components of the automated decision-making. And pretty much are moving this management system from monitoring concept into decision making which is to an expert level system, which in the future should be able to monitor all the operations and primarily the routine ones, significantly reducing the total cost of ownership.

In terms of deployment, we use multiple options based on this triad of platforms, again: EMAN, CSM, and CE. We use event-trigger deployments, we use scheduled deployments, we have rapid deployments when we push and pull policies and echoes to counter an attacker. We use regular deployments, once per 24 hours or 48 hours. And we manage the IOS images for our users, which essentially allows them to be free from any management burden and just conduct business. This is part of our managed security concept of Cisco IT.

Bob Scarbrough: I like the security aspect of being able to push an update to like an ACL –

Plamen Nedeltchev: Yes.

Bob Scarbrough: – to counter a threat.

Plamen Nedeltchev: Correct.

Bob Scarbrough: Whereas in the past, you'd have to – you know if you didn't have that mechanism in place, you're exposing. Your level of risk is increased –

Plamen Nedeltchev: Correct.

Bob Scarbrough: – exponentially.

Plamen Nedeltchev: Yes, but not only that, we essential – we enforce policies, we audit policies, we make sure that at no time the security can be compromised. We essentially can control if the configuration is performed from CSM or out of CSM, out-of-band, and not notify the head-end to make decisions, and so and so forth. I could probably talk a lot about all the tools but probably we don't have enough time to do it.

TCO and Automation of Routing Operations – Major Automation Wins

However, I want to emphasize something which we were able to do: analyze all the major case generators which essentially are one of the reasons to have one or another bigger management burden on the head-end and automate those, calling them routine operations. And you see them on the screen: migration from one device platform to another, connection type change, upload speed change up and down, service move from one location to another.

ISC to CS-M Migration. Platform A to Platform B Migration

If we think about the migration as one-time operation, you probably are gonna be wrong. The reality is that every enterprise migrates constantly from one management platform to another, from one device to another, from one IOS to another. Our challenge pretty much in Cisco IT is to migrate about 7,500 users globally without hiring new people, without introducing any additional burden to our IT support staff.

ZTD IPT Deployment (Home)

This is the process; you see it on the screen, it's very simple. The process essentially allows the user to apply a new service, maintain his old service for 30 days. And after his new service is connected and operational, we start a countdown which is gonna terminate his old service in a matter of 30 days. The process is fully automated. We don't expect that to create any kind of problems.

The process is very straightforward. It's very easy to understand from user perspective, and it should allow us to migrate very easy. The old management platform which was IP Solution Center into CSM, and the old generation ECT which was based on 831 router to 871 router, in a matter of the next several months.

However, if you remember we are devoted to deploying ZTD technologies for every single next service we're gonna allow users to apply as part of our production offering. The next one today as part of next generation ECT, we provide IP telephony as a standard, as a production service for our users –

Bob Scarbrough: And supported.

Plamen Nedeltchev: And supported service, exactly. The next step is gonna be providing video services and so on and so forth, so –

Bob Scarbrough: When you're saying "video services," are you talking like Cisco Unified Video Advantage?

Plamen Nedeltchev: Probably for the beginning we're gonna start with VTA. And then after that, upon increasing the capacity of our up-speed links of our ports, we're gonna probably think about the next steps of opening a video services for our users.

Bob Scarbrough: Excellent.

ZTD of IPT for Remote Access

Plamen Nedeltchev: Now, this is how the zero-touch deployment for IP telephony works. The user applies for IPT service, IP telephony service, as part of his ECT service, and upon approval orders is IP phone from the factory or installs his IP Communicator. In additional instances, the phone is gonna be configured with the user, with the directory number in the Cisco Call Manager. IPT device is gonna be shipped from the factory. The router, ECT router, is already successfully configured, up and operational.

ZTD of IPT for Remote Access (Cont…)

Bob Scarbrough: How does it know what the MAC address is, though?

Plamen Nedeltchev: We're gonna essentially intercept the MAC address from the Call Manager and match it with the user authentication from the Web. Match both, we're gonna make a connection, and we're gonna replace the random directory number configured on the Call Manager with the real username. And associate it with the form, which is associated with him during the process of applying for the service.

Bob Scarbrough: That's what I call automation.

Plamen Nedeltchev: Yes, so the process is fully automated. You see the remaining steps on the screen, essentially and it works today without creating any case or without requiring any assistance or any handholding from any of our IT engineers.

Service oriented architecture over VPN. QOS, IP SLA and Lessons learned

Now let's talk about the next step of ECT, which is pretty much associated with building the service-oriented network. When we want to commit to our users Quality of Service, we want to measure the Quality of Service using IP SLA. And let's spend some more minutes and talk about lessons learned from everything said so far.

Bob Scarbrough: Excellent.

QoS and minimum SLA Requirements

Plamen Nedeltchev: Now the Quality of Service configurations are part of every configuration today, next generation ECT router, 871-router. We classify the traffic in several different classes: time-sensitive class, which is typically associated with voice and interactive video; business-critical class, for Oracle, Sub, WebEx, and MeetingPlace; best-effort class for Internet access file transfers; and scavenger class, associated pretty much with backups, large file transfers, and streaming video. Every class, every traffic class, is gonna be associated with a service class, with a separate loss latency and jitter requirements.

IP SLA Metrics

If we want to measure the success of the IP SLA, we probably have to talk the IP SLA metrics, and you see the metrics on the screen.

Bob Scarbrough: Another quick – can you go back a slide? On the scavenger class, you said, “video streaming.” Wouldn’t you require Quality of Service there on –?

Plamen Nedeltchev: We’re gonna talk about that.

Bob Scarbrough: – video streaming?

Plamen Nedeltchev: Yes, you’re right. We’re gonna talk about in a minute. However, from point of view of the business class importance, streaming video, movies, etc., probably they’re not gonna be considered soon as a business-critical class.

Bob Scarbrough: Good to know.

Plamen Nedeltchev: That is the reason to classify this particular video as part of the scavenger class.

Plamen Nedeltchev: Thank you.

Plamen Nedeltchev: Specifically – if I wanna make one more comment. Specifically when you have the situations of natural disasters or situations of social distancing, or earthquakes, etc., you may – you want to have tools to restrict this type of streaming videos and give priority to business-critical applications, to allow key decision makers, executives, important business resources to reach the corporate resources and make some critical decisions.

Bob Scarbrough: Okay, thank you.

Plamen Nedeltchev: In terms of IP SLA metrics we’re could talk about the minimum set and the full set of IP SLA metrics. The minimum set includes latency, packet loss, delay variation, so it’s typically referred to as jitter, contracted bandwidth, throughput, contention ratio.

The full set would include parameters like availability with these two components, you see it on the screen: per-flow packet sequence preservation, which is critical to deliver Quality of Service when you’re providing Quality of Service for interactive video; emission control criteria; and ISP-supported TOS on the edge. It’s equally important every enterprise to be on the same page with the ISP providers when it comes to measuring and reporting tools, so they could be sure that their SLAs to each other can be met – can be measured and reported successfully.

Based on IP SLA metrics, we could commit to our users' particular Quality of Service. And I’m gonna give you, essentially, several rules we use in our environment to address those requirements.

IP SLA Requirements for IPT@Home and Interactive Video@Home

For IPT for home and interactive video, which are two applications typically associated with time-sensitive applications, we try to provide the following parameters in terms of IP SLA.

The loss should be no more than 1 percent, package drops and loss. One-way latency should be no more than 150 milliseconds. Jitter should be no more 30 milliseconds. Voice traffic should be classified as expedite forwarding or QoS5. Call signaling traffic should be marked as AF31, CS3, or Assured Forwarding 31.

The codec type should not be a factor, configuring IP telephony for home. However since we apply shaping and since we have a high latency environment, which typically hard in the end-to-end

environment, the preferred option would be to use G.711 which codec type manages better out-of-order packets and consecutive drops.

Interactive video traffic should be classified as Assured Forwarding 41 or marked with QoS4 or QoS2, based on the preferences of the enterprise. The minimum priority bandwidth guarantee for the LLQ queuing or Class Base Weighted Fair Queuing is the size of the videoconferencing session plus 20 percent per single video call, which is configured with 384K. On the Call Manager we have to provide and prioritize 460 kilobytes per second of guaranteed bandwidth.

Other QoS and IP SLA Requirements

Here are some other requirements and recommendations for QoS and IP SLA. The streaming video, the one we're talking about earlier, should be classified with lower priority with CS3. Loss for this particular class should be no more than 2 percent and latency should be no more than four to five seconds.

There's a group of locally defined mission-critical class, which differ from enterprise to enterprise, and then they come in two categories. The first one is transactional interactive, like client service applications, messaging applications – client server applications, interactive messaging applications, which essentially don't belong to the category of locally defined mission criticals. The other class is a bulk non-interactive application for large file transfers, e-mails, network backups, databases, etc.

The next class is the best-effort class. The general recommendation, the rule of thumb, is to allocate at least 25 percent from the WAN link bandwidth for best-effort class. Scavenger class which is less than best effort, the class which essentially can suppress the applications would have to be de-prioritized and with a lower priority than the best effort.

And finally, we configure the routing and network management class, which is optional but we use it in this QoS over VPN to prioritize the traffic like SNMP, NTP, syslogs, NFS, EIGRP, and ISAKMP and PEC.

IP SLA ProbeTypes, Radius and POP

With not only, we – here are the props we use in our environment to run and assess the quality of our services. The way the IP SLA environment is built; every router sends an IP SLA probe to a head-end, which is called IP SLA Responder, every three minutes. The responder, or a script run by the enterprise management system, essentially polls every router every five minutes and builds the graphical diagram and representation of the Quality of Services of this router.

IP SLA Statistics-Example

One of the major probes we use is on the screen. You see the statistics from "show IP SLA status" command. We measure the RTT, latency one-way, from the source to destination and from the destination to the source, we measure the jitter, we measure the packet loss. If you take a look carefully on the screen, you're gonna see that the RTT varies are average of 13 milliseconds.

The one-way latency is about 2 milliseconds to 11, and the jitter, it's very low, well under the recommended values. The packet loss is literally non-existing. No wonder that the MOS, which is the Mean Opinion Score, for this particular connection is 434, which is between good and very good.

Lessons Learned

Now let's sit back right now and try to talk about the lessons learned from this session. I hope I was able to convey some of the main lessons for our deployment so far. But if we want to comprise them into one single slide, probably this is – those are the major points I would like make as part of this lessons learned slide.

The selection of the hubs: it's very important, from point of view of providing IP SLA type of services for users, we try to maintain the locations of the hubs in a way to optimize the latency and keep it under certain threshold. In our environment, we're trying to keep that under 100 milliseconds, even we recommend 150, anticipating the increased number of congestion issues in the future with the growth of the broadband penetration around the globe.

We would recommend to start with the limited power, become familiar with the technology, and grow to 100, and then understand information, requirements, and systems, and flow and scale. Deploying the technology to multiple segments of the network will allow the IT organization, and that will happen with Cisco organization to maintain a lower TCO.

In fact, the ECT started with the pilot project of 100 to 500 users called Stealth and this is where we – that project was transformed into a production version called ECT, and today already next generation ECT, and so on and so forth. So that approach proved to be the very successful and very mature one, and that pretty much explains the success of this program internally for Cisco.

Bob Scarbrough: Also it enabled the IT to provide a lot of feedback to the business unit.

Plamen Nedeltchev: That's absolutely correct.

Bob Scarbrough: Because there were several gaps, I should say initially –

Plamen Nedeltchev: Yes.

Bob Scarbrough: – and allowed them to build a better product, for –

Plamen Nedeltchev: Absolutely.

Bob Scarbrough: – not only Cisco but for our customers.

Plamen Nedeltchev: Absolutely. Absolutely.

The next condition would be plan for phase approach for new services. Provide SLA and IP SLA for services for your users, try to work with the ISP providers to get them to support the Quality of Services on each of their network.

Use CSM CE to deploy and manage the environment. For large-scale deployments use no-band APIs to integrate these management platforms into the existing management environment, automate the routine operations, and probably finally develop a practice of monitoring and support, allow the support engineers to participate in the final, in the Pilot phase.

Bob Scarbrough: Plamen, I'd like to say thank you very much. Not only was this an informative session, but you and I go way back, we actually were involved –

Plamen Nedeltchev: Yes.

Bob Scarbrough: – in the pilot of this –

Plamen Nedeltchev: Yes, yes.

Bob Scarbrough: – several years ago. And it's been an honor to be – to working with you and it's fascinating this technology and I could see the passion in your conversing with this group here. So I just want to say thank you very much and –

Plamen Nedeltchev: Thank you, Bob Scarbrough. Thank you.

Bob Scarbrough: – I look forward to continue –

Plamen Nedeltchev: Thank you for having me here.

Bob Scarbrough: – to work with you.

Plamen Nedeltchev: And I hope that I was able to convey all the messages to our viewers and to our audience, and they're gonna find this presentation really beneficial for them moving on.

Further Resources

Bob Scarbrough: And for more information about technologies and solutions deployed at Cisco, go to the Cisco on Cisco site where you can find Case Studies with information about: what technologies we've deployed, benefits gained, lessons learned, and some operational best practices and presentations to help you learn more.

Below that, you'll see a toll-free number you can call for more information or to place an order; and you can order Cisco resources on the web from the URL at the bottom of the page.

Also, there is a plug here for Plamen's book, you can see there at the bottom and you are welcome to reference that as well.

I'd like to thank those of you watching for spending this time with us, and for being interested in what the Cisco on Cisco Technology Tutorials are all about. We hope that you've enjoyed this seminar and that it's helped answer some of your questions about virtual private networks.

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Europe Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: +31 0 800 020 0791
Fax: +31 0 20 357 1100

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Printed in USA C78-331727-01 10/06

