



data communications

Corporate Network Services of Tomorrow

Business-Aware VPNs

Authors: Daniel Kofman, CTO and Yuri Gittik, CSO

Content

Content.....	1
Introduction	2
Serving Business Customers: New VPN Requirements.....	2
Evolution of VPN Services: First 3 Generations	2
The 4 th Generation: Business Aware VPNs	4
Understanding the BA-VPNs Market Opportunity.....	6
Overall Benefits of BA-VPNs.....	6
Functional Requirements of BA-VPNs	7
Implementing BA-VPNs	7
Conclusion.....	8

Introduction

Serving Business Customers: New VPN Requirements

There is a clear trend in the telecommunications enterprise market towards increasing the value of service providers' product portfolio. Indeed, on the one hand, CIOs are looking for further reducing costs and improving productivity of company's distributed business processes (Figure 1). On the other hand, they would like to go a more disruptive way and shift their network service strategy from just enhancing productivity to value creation with new business processes and collaborative applications enabled by new communications means. Hence, emerging "new generation" communications services will have a significant and increasing impact on the enterprises and industry business processes. These services have to support communication critical applications in the changing environment of network and services convergence. Therefore VPNs, a key component of enterprise service offering, are evolving to shift the focus from multiple site connectivity to applications, collaborations, individuals and communities of individuals. The range of applications is huge, from healthcare solutions requiring very high robustness, to networked industry plants and services like online market exchanges requiring very low latency and delay variation.

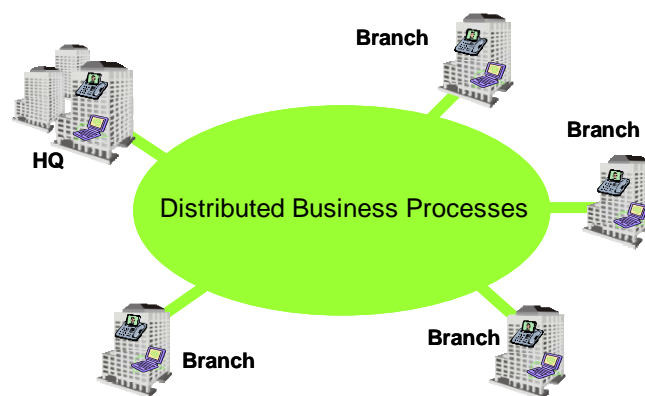


Figure 1: Distributed Business Processes

Evolution of VPN Services: First 3 Generations

The telecommunications market has been moving fast to the current 3rd generation of VPN services, and it is already facing the 4th generation. Let's briefly review the three first generations that focus on connectivity between multiple sites (Figure 2).

The 1st generation was based on circuit connectivity between sites based on TDM/PDH and then complemented with SDH network infrastructures. The lack of flexibility, coarse granularity, lack of efficient bandwidth utilization and relative high cost of these solutions, together with the readiness of new technologies, triggered a migration towards the 2nd generation, in which TDM circuits were replaced with Frame Relay and ATM connections. Flexibility was introduced in terms of provided bandwidth and classes of service; still the service provider network offered just end-to-end connectivity between sites, all the VPN

intelligence is located at customer premises. The typical topology of these VPNs was a star, centered in the company Headquarters. This solution nicely fits the system architectures and business processes used at that time.

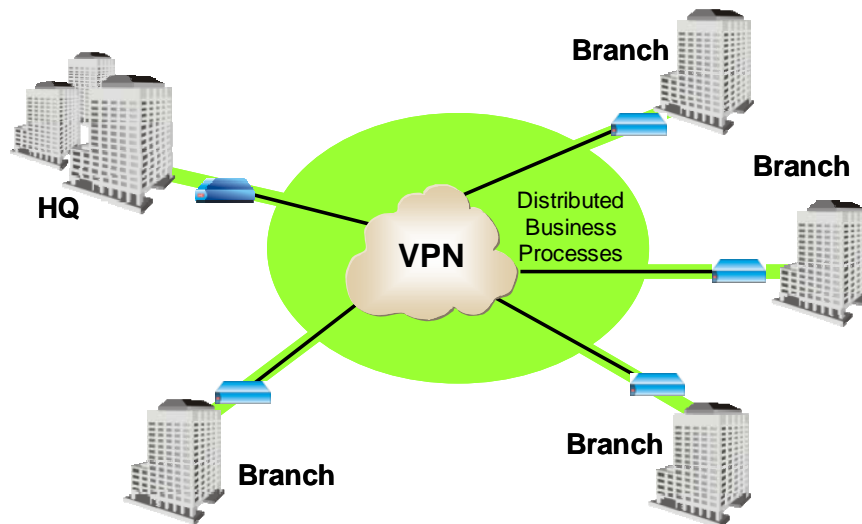


Figure 2: VPN Connectivity

Today, the market is focused on the 3rd generation, for which the first service model was the IP VPNs, typically based on an IP/MPLS infrastructure. In this 3rd generation, the VPN intelligence is managed by the service provider, by network located devices. In a simplified view, the key differentiator of this 3rd generation is that each VPN site sends all the traffic to a device controlled by the service provider and this device (usually called a Provider Edge Equipment - PE) is responsible for forwarding the traffic to the right destination (e.g. customer site) with the expected quality of service (QoS).

This service model introduces the flexibility to deploy, at a controlled cost, VPNs of any topology¹, which is well adapted to the new business processes of the enterprises. The enterprises therefore outsource the VPN intelligence to the service provider. As an example of network functionality that is outsourced, one can cite: connections termination, switching and routing.

Other important characteristics of the 3rd generation are:

- the increasing usage of Ethernet as the IP transport layer, in particular as the access technology to IP VPNs (allowing high scalability and cost reduction).
- the usage of DSL access allowing for low-cost, ubiquitous, mid-bandwidth connectivity well suitable to SMEs.
- In addition, an increasing number of enterprises, including most large enterprises, have deployed or are in the process of deploying voice over their IP VPN.

¹ In the 1st and 2nd generations, the services were built using multiple physical circuits (PDH, SDH) and virtual circuits (Frame Relay, ATM). The number of required connections depends on the number of sites and on the topology of the VPN. In the 3rd generation, a given site send the traffic to a given network device (the PE) which is in charge of forwarding the traffic to the right destination site. As a consequence, the cost of the VPN is not directly related with the topology of the VPN, there is no restriction on the possible direct traffic exchanges between all the sites (direct means here that the traffic between two sites does not transit through a third site).

More recently, layer 2 (Ethernet) VPN services started to be provided. This was facilitated by the Metro Ethernet Forum (MEF), which clearly defined the framework for various Ethernet services. These layer 2 VPNs belong to the same generation in our classification since as for layer 3 (IP) VPNs, the forwarding intelligence is deployed in network located devices, and a given site sends all his traffic to a given PE which is in charge of the forwarding towards the destination site (let us remark that, in addition, Ethernet is also being use for point to point connectivity). As for the layer 3 VPNs, these services are suitable for metro, country-wide and global networks (let us remind here that the MEF is changing his name to GEF, Global Ethernet Forum). To clarify the concept of layer 2 VN, let us remark that since most user applications are IP-based, most of the traffic over a layer 2 VPN is IP. The major difference with layer 3 VPNs is that the service provider bases the forwarding decisions on the layer 2; the layer 3 routing, when required, remains under control of the customer.

Layer 2 Ethernet VPNs are typically deployed with dedicated Ethernet network termination units (NTU) that are located at customer premises and ensure end-to-end service control, fault management (with standard Ethernet OAM) and traffic handling. In fact, in case of IP VPN, a customer-located router usually serves as the “IP NTU” that enables required termination functionality at IP layer.

A lot has been written and said regarding layer 3 and layer 2 VPNs comparison. We will not repeat this debate here, we would like just to remind that main VPN service providers have arrived to a conclusion that both VPN types will co-exist and roughly estimate that layer 2 VPNs, a market that is now growing even faster than layer 3 VPNs, will represent 30% of the total VPNs market in the near future (once a stationary growing of these two services is reached).

Still, in both service models of the 3rd generation, the customer is required to have strong network competences, since the SLAs are based on technical parameters, such as bandwidth, Classes of Service, Quality of Service, etc. It is the customer responsibility of defining the mapping policies between the applications and the Classes of Service, even when the effective mapping (marking) is done by a service provider’s network equipment. Also the customer has to understand the benefits and drawbacks of the Layer 2 and Layer 3 solutions in order to select what solution better fits his specific needs. When both service models are available, a customer would prefer to have a hybrid solution; for example, adapting the technology per type of site (small/large branches connection, data centers interconnection, etc.).

Both 3rd generation IP and Ethernet VPNs enable basic connectivity between multiple customer sites that is used as the foundation for value-added services (VASs) to provide additional value proposition with premium services. Such VASs are in particular important when IP VPNs are becoming a commodity that results in pricing erosion. For instance, this could be end-to-end quality assurance (with several Classes of Service). Another example is OBS “Enterprise Application Management” offering, which is a value-added service for IP VPN. In fact this OBS service is a clear move towards the 4th VPN generation, yet not a case of Business-Aware VPN that is introduced in the White Paper.

The 4th Generation: Business Aware VPNs

The driver for the 4th VPNs generation is the customer desire to completely outsource VPNs to service providers. Indeed, as indicated in recent market analysis, CIOs prefer to focus on their business processes, their applications, the level of criticality of the applications and the Quality of Experience expected by the end-users (individuals and communities of end-users) rather than deal with technical issues of network functionality. They don’t want to be involved in technology choices (e.g. level 2 or level 3), nor to be familiar with technical concepts like bandwidth and classes of services.

In Service Level Agreements (SLAs) of the 4th generation VPNs, the customer will define its requirements in a non-technical language that he understands: the sites he has, the applications he runs, the expected Quality

of Experience (instead of Quality of Service), the relative criticality of the various applications, the required level of security, etc.; of course, along with the classical metrics of reliability and availability.

A main advantage is that the VPN service definition is agnostic with respect to the type of “enabling” connectivity. This will be a huge benefit for service providers that introduce the 4th generation. Instead of being confronted to the complexity of marketing Layer 2 and Layer 3 VPNs and presenting techno-economical comparisons, the service provider will focus his marketing efforts and differentiators on how close his service is to the real customer needs, as understood from a business process point of view.

The service provider can decide whether it is preferable to deploy layer 2 or layer 3 VPN solutions per client, per network segment and even per site. Indeed, such a solution can combine Layer 2 Ethernet and Layer 3 IP connectivity (see Figure 3).




Figure 3: Network technology agnostic VPNs

The key concept of the 4th generation is that the service provider identifies the application that has generated a flow (a voice call, a file download, a Citrix activity, etc.) and uses this information for handling the entire flow end-to-end (the sequence of packets of the flow), according to the type and level of criticality of the application, as defined in the SLA.

In addition to the controlled application aware connectivity, some added-value functionality is provided: visibility (end-to-end quality monitoring, resource usage monitoring), online configuration, security, expenses control (alerts when new resources are required), applications acceleration, etc.

We will call these 4th generation VPNs “Business-Aware VPNs” (BA-VPNs). The following two major features differentiate this new generation:

1. Service requirements are defined in terms of quality of experience per application and performance of the business processes (not technical parameters of required connectivity)
2. The service is provided transparently to “enabling” connectivity and network type (“connectivity-agnostic”), so it might be delivered with IP and/or Ethernet connectivity of multiple customer sites.

Understanding the BA-VPNs Market Opportunity

The customer demands for this type of VPN solutions is not new. Historically, when an enterprise was facing Quality of Service problems, the first “natural move” was to negotiate the increase of the access bandwidth of the sites facing the quality degradation. This solution is expensive, and doesn’t always solve the problem. Indeed, the congestion may be in other parts of the network, in the servers, etc. Therefore, in the late 90th, several start-ups introduced new enterprise monitoring appliances. Such equipment, owned by the customer allowed monitoring the usage of the access (WAN) links and optionally provided basic traffic control over available access bandwidth.

This allows the customer to validate whether the problem is in the access or not. In the cases the problem is indeed in the access link, a natural solution is to increase the bandwidth, but this is not always necessary. Indeed, the impact of the degradation of the quality of experience on the performances of the business processes depends on the application, and usually non-critical applications’ bandwidth requirements increase much faster than critical applications’ bandwidth requirements. To handle this issue, the monitoring appliances were enhanced with application-aware flow control. This functionality enables to identify the applications and to control the flows in case of access link congestion in order to protect the critical applications.

As noted in the previous sections, in order to better fit the requirements of enterprises’ business processes, the 3rd generation of VPNs was introduced, allowing, for example, services with no restriction on the VPN topology, specifically enabling direct traffic exchanges between sites (usually based on the so-called “hose model”). The new enabled business processes are more distributed and the consequence is that controlling the traffic access link by access link is suboptimal and a global control of the whole traffic of the VPN is required. A very few companies in the world are today proposing solutions for this requirement.

In any case, today the companies see the application / traffic awareness solution as appliances to be deployed in their sites, usually under their control (even when selected and bought by the service provider).

In the new generation VPNs, this functionality is outsourced and sold as an added value service, a key enabler for advanced business processes and a new source of revenues for service providers.

Overall Benefits of BA-VPNs

The most important benefits of BA-VPNs for customers are as follows:

- Network transparency from the CIO’s and individual end-user’s point of view. The CIO can outsource ALL possible networking functionality to a service provider (of course, this type of service allows also for a limited outsourcing adapted to each CIO strategy). The end-user obtains the expected Quality of Experience on critical applications under any circumstances.
- As a consequence, the business process of the company can be enhanced, but more important, new business process, with an increased productivity and generating new revenues, become possible. Indeed, the dynamic protection of critical applications, in a

seamless approach from the end-user point of view, and without the requirement of any specific expertise for the IT team, is a main enabler for deploying new business processes.

- The obtained global visibility allows the enterprise for better expenses control and planning (predictable services).

The most important benefits of BA-VPNs for service providers are as follows:

- The service provider sells the BA-VPN as an integrated service; it doesn't need to market different "technically defined" offers (IP VPN, L2 VPN), requiring technical insides and controversial discussions. This is becoming more and more important as the trend of many CIOs is towards the outsourcing of networking facilities, to focus on their core activities, more closely related with the enterprise business processes.
- Differentiating new VPN services with high value proposition (from basic connectivity to enabling business processes and collaboration) that leads to increased revenue and preventing customer churn.
- The new VPN offering is highly adaptable to any specific needs of every customer.
- Optional additional VASs (application management, security, etc.) are facilitated.

Functional Requirements of BA-VPNs

To supply BA-VPNs, the service provider has to identify the flows generated by the various sites of a VPN instance, then recognize which application and at which site has generated the flow, and after that make a decision regarding forwarding (which destination, which path) and quality control (scheduling among the various flows based on the type and criticality level of the application that has generated them).

For this purpose, the service provider may have to analyze the customer traffic. The analyses can be done in different ways; the most popular today is the so called Deep Packet Inspection (DPI). In DPI, at least for the first packets of a flow, the content of the traffic is analyzed up to the application layer in order to identify the application. DPI can be complemented with a statistical analysis of the temporal structure of the flows.

Traffic awareness may therefore require analyzing the traffic at several layers, including the application layer, as well as certain traffic semantics. This does not impose any restriction on the traffic forwarding; usually traffic is forwarded at Layer 2 or Layer 3. In other words, a packet will be analyzed to decide to which flow it belongs to, or if the packet initiates a new flow. From this analysis, a decision will be taken on functionality like scheduling (the order at which packets are transmitted) and queue management (which packets are dropped in case of congestion). Still all the packets are forwarded based on Layer 2 or Layer 3 addresses, as in the case of VPNs without application awareness. More precisely, the forwarding will be based on MAC addresses forwarding tables or on IP addresses routing tables, while the application awareness enables flow identification, and mapping the traffic over the required network Class of Service or specific network connection.

Implementing BA-VPNs

This section describes a possible approach to implement a BA-VPN, which is a simplified overview of the solution being developed by RAD and Ipanema.



Based on the facilities of the 3rd generation VPNs, like IP VPNs or Ethernet VPNs, and of new technologies (like IP Telephony), the direct traffic exchanged between multiple sites of a company is increasing (let us remind here that we call direct traffic the traffic between two sites that is not forwarded through a third site). As an example, in IP telephony the signaling for session (call) set-up and the media may follow different paths: typically the signaling will be exchanged through a server located in the Head quarters or datacenter, whereas the media (voice) will be exchanged directly between the terminals that may be located at different branch offices, therefore resulting in a meshed traffic. In such a distributed environment, the traffic aware control has to be done in a distributed way and requires specific functionality at customer premises. This functionality is under full control of the service provider in the BA-VPN service model.

Moreover, a service provider will deploy at the customer premises a device that will enable smart demarcation (fault management, quality of service control, etc) at Layer 2 (Ethernet) or Layer 3 (IP). It seems natural, for CAPEX and OPEX reduction, to integrate the whole functionality (traffic awareness and smart demarcation) in the same device. These customer-located devices will collaborate in the VPN architecture with centralized servers that will allow for a policy-based configuration of the customer-located devices, therefore further reducing OPEX. The routing and switching functionality is provided by the existing Layer 2 and Layer 3 VPN infrastructure, no additional investment is required here, previous investment is leverage for the provision of the BA-VPNs.

Conclusion

The present document introduces the concept of BA-VPN, its benefits, a short overview of its requirements and a possible implementation framework. BA-VPNs are defined here as a new service model instead of an added value to existing service models, proposing an integrated new value proposition and cost reduction approach. They leverage existing investments in Layer 2 and layer 3 VPNs, no additional investment is required for connectivity.

The White Paper presents a very first draft of the new BA-VPN model. Its main target is to facilitate further discussion on the concept, its business model and implementation in specific carrier environment.