



Multicast Protocols

The Case for Virtual Routing Support
November 2006



INTRODUCTION

As convergence becomes more of a fact than a theory, the complexities that were being addressed (i.e.: separate physical and logical networks) continue to exist to some degree or another.

Voice traffic is still not able to tolerate more than a fifty millisecond delay between spoken word and reception or subscribers look for another provider.

Data traffic based on the IP protocol will tend to consume as much bandwidth as is available to it without regard to its neighbors on the link.

Multimedia traffic requires not only timely (and on-time) delivery, but larger and larger amounts of bandwidth as quality improves in response to consumer demand.

Carriers are deploying IP Next Generation Networks to deliver triple play services (i.e.: Voice, Video and Data [VVD]), including video-based services like IPTV. Carrier success in delivering the 'triple-play' of voice, video, and data is not only dependent upon the proper choice of service and content partners but also on the right network infrastructure. This network infrastructure must be capable of evolving as business and consumer needs change, as new services and applications are introduced into the marketplace, and as bandwidth needs grow. New consumers, providers and services must be able to be introduced without disrupting existing services and services delivery. The dilemma, then, is how to maintain service level commitments, preserve service quality, and please consumers without placing overwhelming demands on equipment, personnel and budgets. One strong answer is separating consumers, providers and services along the lines of explicit domains.

VIRTUAL ROUTING

Virtual routing, the ability to provide the appearance of multiple individual machines on a single physical platform, is a vital tool in the workshop of service (and content) providers in meeting the demands of NGN markets. VR allows for the partitioning of a single platform (although, admittedly, one with the necessary horsepower) into a number of dedicated, isolated virtual domains addressing one, or at most a few, specific needs of the installed base.

The convergence of networks has been attributed, in some degree, to the desire to reduce expenses associated with meeting market demands. Capital expense is reduced by consolidating multiple physical links (and thus ports) onto a single link/port combination. Operational expense is reduced by reducing the number of network elements that require administration, resulting in productivity gains for NOC staff and installation professionals.

The consolidation of multiple network elements, then, onto a single platform results in even greater efficiencies for providers. The problem is how to guarantee this additional savings by reducing or eliminating any additional expense that might result from contention between transports, services or content?

The answer is to partition the network element using VR into multiple, logical domains; each domain exists in isolation from the others by virtue of platform and OS replication on a domain basis, and; each domain, as a separate logical entity, is incapable of affecting concurrent services on the same network element due to its isolation from critical platform and infrastructure components. Services competing for processor, memory and network resources have minimal impact on each other with the efficient distribution of resources among them on either physical or other boundaries (e.g.: low- and high-water slices of time). In addition, new domains may be added in real time as new services or consumers come on-line, ensuring that revenue generation is not affected.

SERVICES SEPARATION

What is necessary, then, is a coherent and flexible set of policies governing the distribution of resources on a per-element basis.

How many links, of what kind, will be given to each virtual machine?

How many processor cycles, or even processors, will each VR have access to during normal operations?

How much memory is required for each VR to provide sufficient storage for tables and buffering to ensure on-time delivery of data to each subscriber with minimal disruption arising from interference or capacity mismatch in the links making up the network path?

How will surviving platform resources be divided when a component fails (as is sure to happen)?

With this set of policies in place, the network element is divisible into multiple logical entities tuned to the security and provisioning requirements of a specific content provider, delivery of a particular revenue-generating service, or perhaps a number of services to a particular subscriber.

By Content Provider

Content providers differ by their content (mind-numbing revelation, here) which will have very different requirements for network resources on a path basis. From origin to end all the traffic originating from a particular content provider will have the same set of characteristics. Those characteristics in an NGN will have the potential for disruption of the other traffic on the network. Audio content is streaming in nature, and has little tolerance for delay or packet loss between source and destination. Competing traffic on the network cannot impose either on a flow or flows in transit over the infrastructure.

Video content can be separated into two different types:

Standard Programming generally has numerous subscribers, presents content at specific times, and is more global in nature; standard programming is ideal for transport by multicast protocols, as this will significantly reduce contention for resources in the core, and final distribution can be done as close to the leaf nodes as possible;

Video on Demand [VoD] has few subscribers, presents content when it is requested, and is very local in nature; VoD can be transported by multicast protocol, but is more suitable to unicast transport with the danger of large amounts of duplication in the core if the subscription set is large enough, and distribution must be done at the source, as requests are unpredictable.

Application content has none of the bias toward low-latency, guaranteed delivery that is present in audio or video traffic, tends to be comprised of a few large packets, or a lot of small packets, yet may share some other characteristics with them. Subscribers are, on the whole, known and limited in number. Requests for services occur at predictable, regularly scheduled times, and may be suited for multicast distribution (e.g.: inventory updates to retail stores from national or regional distribution centers).

The ability to provide a separate VR to a content provider allows for the ability to tune that VR to the specifics of the expected traffic. It also allows for the isolation of that provider's network element from other elements, eliminating the possibility of the VR becoming collateral damage in the event of external catastrophic failure. Higher levels of SLA can be provided, at additionally higher cost, maximizing the revenue stream from each source.

By Service

Services may also be segregated by VR, taking maximum effect from the efficiencies of scale provided by generalizing service-specific components across large numbers of subscribers. Critical

services can be protected from interruption or failure as a result of instability in another network service.

Voice traffic can coexist on an NGN with video (or some other high-demand) service, yet be protected from loss of bandwidth or ports in the event that a recovery takes place in the video network requiring failover to backup paths. By virtue of the VR the voice traffic transits links that, while physically attached to the same platform, are invisible to the virtual network elements that exist on the network element.

Video traffic continues to generate revenue while high-speed Internet access and application services are disrupted by virtue of the fact that they have been relegated to network resources that are not shared, or are shared in a minimal way with the video service. Viewer satisfaction is not compromised, and forms the base for the layering of new video services as they are introduced.

By Subscriber

Some network demands may be best addressed by providing a VR on a subscriber basis. CLECs can be supported by an ILEC through a VR, guaranteeing SLA while providing integrity and security to subscribers from malicious access originating in external networks.

Municipalities having differing contractual relationships can be segregated by VR, allowing favorable status to be preserved. Local and regional regulation can be addressed without proliferation of network elements at the network edge.

Even smaller concentrations of subscribers can be served by VR: a multi-tenant unit can negotiate preferential treatment at the content provider or service level; a corporation can negotiate preferential treatment at the services or SLA level, and; even individual subscribers can be protected from outside interference, or prevented from interfering from the outside, if the price is right and the market will bear it.

CONCLUSION

As networks, the services they carry, and their providers move further into the NGN, one technology will stand out from others in the delivery of data. Virtual routing, and particularly virtual routing of multicast protocols, will provide the saving in capital and operating expense required to efficiently meet the emerging demands.

As multimedia, and video in particular, becomes a more critical component of the total offerings from Service Providers, only Virtual Routing will provide the insulation and flexibility necessary to maximize this revenue stream. Protection from disruption, favorable allocation of network resources and a secure subscription base will retain existing subscribers, and encourage the uptake of new subscribers.

Through all of this, Virtual Routing will protect existing, legacy services (e.g.: Voice or Data) contributing to the retention of subscriber base and enhancing the standing of the Service Provider. IP Infusion provides a proven, real-world hardened Virtual Routing implementation that supports Layer 2 and Layer 3 protocols, and is extensible beyond unicast to multicast protocols, guaranteeing a firm foundation for today, and NGNs to come.



© Copyright 2006 IP Infusion Inc. All Rights Reserved.

ZebOS and IP Infusion are registered trademarks and the ipinfusion logo is a trademark of IP Infusion Inc. All other brands or product names are trademarks or registered trademarks of their respective holders. All specifications within this document are subject to change without notice. Contact Sales for current feature availability.