APPLICATION-AWARE NETWORKS

Evolving Carrier Business Models with Application Aware Technology

Table of Contents

Executive Summary 3
Introduction 3
Imperatives for Service Providers 4
Evolving “As a Service” Business Solutions 4
“Open Garden” Partnerships 5
Reducing the Risk Associated with Service Definition 5
Satisfying Dynamic Service Needs 6
Visualizing and Assuring Network Performance 6
How Juniper Networks Enables the Application Intelligent Network 7
Control Functions 7
Traffic Processing Systems 7
Juniper’s Commitment to Open Development 8
Conclusion 9
About Juniper 9

Table of Figures

Figure 1: Juniper Networks open software development process 8

Executive Summary

Based on network bandwidth, traditional carrier offerings for enterprise networks are disconnected from the enterprise CIO's typical method of assessing network value based on application performance. While dimensions such as geographical reach, service quality, and available throughput have provided sufficient scope for carriers to differentiate themselves, the carrier's service offering has focused solely on the provision of connectivity. In some cases, coarse service quality definitions are applied to network services in an attempt to enhance their value. These definitions are made after a manual assessment of the individual application mix of each enterprise customer. Because the impact of each new application requires serious consideration, applications are grouped within different classes of service. An operational burden is associated with the manual changes required to support the dynamic range of applications found within the enterprise.

An Application-Aware Network combines application awareness with dynamic network policy, enhancing traditional managed service offerings and simplifying the act of balancing the needs of applications against available network resources. For carriers, this ability to synchronize network resources with end customer requirements ensures a more efficient use of network capacity and infrastructure, providing better return on investment. Through greater visibility into network use, enterprise customers are empowered to directly customize their own service definitions—in a highly granular fashion—as quickly as their shifting business needs demand it.

This paper is directed to service provider administrators and executives. Through this paper, you will learn about the growing market need for application-aware network infrastructure, how you can leverage this technology to streamline your business operations' application-based service-level agreements (micro-SLAs) and increase the value of your managed services portfolio. In addition, this paper explains how deep inspection (DI) technology, combined with dynamic policy, delivers a cost-effective and flexible toolkit for enabling new business models and revenue streams for carriers.

Introduction

Located strategically between enterprise customers and applications, service providers are uniquely positioned to address the three main priorities of enterprise customers.

- Lowering total cost of ownership (TCO) of the enterprise network
- Improving network responsiveness to evolving and often unpredictable needs of the enterprise's dynamic application mix
- Measuring conformance to application performance and availability SLAs.

While service providers must come to grips with these priorities, care must be taken to avoid the ever increasing number of potential service offerings that create new obstacles and impose on service delivery costs.

Rather than identify each and every service combination, service providers are looking toward more flexible service packaging to allow enterprises to select the range of services that are most relevant to their organizations. Flexible service packaging requires new techniques to manage the complex relationship between the needs of various applications and the resources that are available within the network. The key to ensuring enterprise acceptance and the widespread adoption of these service packaging techniques is to provide measurement and assurance mechanisms.

An Application-Aware Network can therefore be defined as a network that understands and supports the individual needs of applications and their users. The use of dynamic policy enforcement enables the Application-Aware Network to rapidly adapt to the flexible needs of the high-performance enterprise most cost effectively. Per-application assurance and reporting tools are provided to validate that the application-based micro-SLAs are met at all points across the carrier's domain.

Imperatives for Service Providers

The drive for service providers to increase both the range and value of service offerings—moving away from providing simple connectivity and moving toward a large range of business services that can be delivered cost effectively—is part of a broader evolution of carrier business models. The Juniper Networks® vision is to allow service providers to integrate new capabilities such as application awareness into existing Juniper Networks technologies, using open standards-based interfaces that enable carriers to support future business models without costly forklift upgrades.

The Application-Aware Network also helps service providers address their own most pressing needs.

- **To increase revenue from services:** The current growth in enterprise spending is driven by services, whether on or off the network. Service providers must move "up the OSI stack" to generate revenue by focusing on delivering more than just Layer 3 connectivity. To do this, the network must become smarter and more aware of the needs of individual applications. The alternative is to offer a "dumb wire" over which third parties profit by delivering services to customers, with carriers reduced to a simple utility status.
- **To lower the operating cost of service delivery (OpEx):** Reducing the cost of transporting network traffic adds bottom line value and increases the cash available for investment in innovation and continued network expansion. Addressing the cost part of the equation is just as important as generating additional revenue. As the bulk of carrier cost is represented in OpEx, attention must be given to the contribution of people and processes in the total cost equation. Automation and simplification of processes through self service portals can offer significant savings.
- **To expand the ecosystem:** Even the largest service providers understand that they cannot do it alone. They must work closely with application and content providers, strategic partners, and even other carriers to develop innovations that will differentiate their service offerings. For business services, the emergence of "as a service" solutions represents a clear opportunity for service providers to expand service offerings without significant investment in new service delivery platforms.

Evolving "As a Service" Business Solutions

In the current economic climate, "as a service" business solutions are of heightened importance to the enterprise, as they draw upon emerging cloud computing capabilities to overcome the entry barriers of established alternatives. By implementing an "as a service" solution, enterprises can reduce their up-front CapEx and scale OpEx in line with the adoption rate of the particular application. When using an "as a service" solution, a concession is made by the enterprise to place additional importance on the security, reliability, and responsiveness of the underlying network services.

Some typical business solutions and their relationship with an Application-Aware Network include:

- **Software as a service:** As the hosted application is emulating the functionality of a local application, the network must identify and prioritize business critical applications to ensure that the responsiveness of the application is as close as possible to that of a locally hosted application. Additionally, the network should provide visibility into the performance of the application.
- **Security as a service:** In this model, the network must be able to identify applications or even specific message flows within an application that should be redirected for examination by hosted security infrastructure. Care should be taken when diverting traffic to such a service, as unnecessary redirection of traffic may lead to unacceptable application performance—particularly for real-time services—however; lax redirection may leave security violations undetected.
- **Infrastructure as a service:** Often associated exclusively with cloud computing platforms, infrastructure as a service can represent both network and IT services. An example of a network-based infrastructure as a service (NaaS) includes bandwidth on demand, where additional bandwidth capacity can be provided as needed. Additionally, connectivity-based services can be deployed on demand—particularly when combining resources that may be spread across large geographical distances—to assure end-to-end network performance. The Application-Aware Network must be aware of changes to the current resource mix to ensure that applications can make full use of all available resources.

"Open Garden" Partnerships

Increasingly, even the largest carriers understand that they cannot do it alone. They must work closely with application and content providers, strategic partners, and even other carriers to develop innovations that will differentiate their service offerings. For service providers looking to introduce one or more of the "as a service" business solutions, the open garden model has the potential to deliver a high value, high quality customer experience.

The open garden model makes a critical assumption: No single organization is likely to have all of the resources needed to deliver the next generation of high value services. More specifically, the model posits that:

- Any provider can be an owner/proprietor of a service
- Any provider can also contribute key components of a service, without necessarily being the owner/proprietor
- All contributors to the service should be compensated for their offering.

An open garden model lowers the barrier for multiple stakeholders to offer and combine resources in a way that the services of the future will demand. Operational costs can be managed because existing network and IT resources can be reused. Through adoption of the open garden model, service providers can take advantage of the wide range of existing "as a service" product offerings to quickly launch new service models. These "white label" services can be branded with or without reference to the external partners, enabling service providers to present a uniform brand to network users.

Reducing the Risk Associated with Service Definition

Traditional service definition—or more correctly product definition—occurs well in advance of network or device configuration. The consequences of the introduction of a new service are identified and classified as network, IT, or business process impacts. Each impact is examined and an integration plan is developed as needed. This necessary process adds greatly to the time-to-market and cost of the delivered product.

To mitigate the risks associated with the introduction of a new service, a business case must conservatively consider the cost to define and deliver the service against the potential revenue opportunity. The high cost associated with service definition ensures that only services with wide appeal will be approved for development. This decreases the likelihood that a carrier will deliver upon an innovative but unproven market opportunity; as such, service offerings today target mass-market opportunities.

Flexible service definitions allow the carrier to reduce the product definition burden by directly productizing the resources of the network for versatile inclusion in an assortment of products. The Application-Aware Network leverages the power of software to automate the process of deriving network policy changes based on the selected service mix. Resources are selected and a vetted configuration can be automatically pushed into the network immediately following the definition of a new service. These requests can be made using one of the following methods.

- **User driven:** In this method, a user who is typically an enterprise administrator defines the resources required across the managed network via an enterprise portal. Templates or predefined service definitions can assist in the management of a large-scale deployment in a "one click" activation model. However, granularity should exist that allows for the detailed manipulation of a services associated with a single site. This lightweight approach enables the user to define potential quality-of-service (QoS) policies without requiring any specific knowledge of the underlying network technology.
- **Application driven:** In this method, the network is made directly aware of each application's resource requirements. In simplistic deployments, the resource requirements can be based on a per application static definition provided by either the operator or the user. However, for the maximum benefit, the network should be able to sense the application's demands in real time, particularly in cases where the application may vary its resource needs. Call admission control (CAC) provides such a mechanism, enabling applications to request resources with feedback provided. Call admission control allows the network to warn applications of resource unavailability—which can have a number of valid causes—decreasing user frustration and support calls.

By allowing the user or user applications to define the service mix, the carrier has removed the lengthy process of determining which services to offer. Each service can now be offered as an application-based SLA (micro-SLA) that is bound to the characteristics of the offered network assets. Without the high cost of service introduction, the service

provider carries little risk in offering resources in a variety of flavors, regardless of the size of the eventual service uptake. Niche services and new service innovations can be prototyped directly on the network without requiring exhaustive evaluation of market opportunity. Should a new service capture the attention of the market, it can be released quickly to maximize "first mover" advantages. The customer is rewarded with faster access to new innovative services that are matched with their own application priorities.

Satisfying Dynamic Service Needs

Service fulfillment can loosely be defined as the set of required actions that make a network service available to the application it supports. With the needs of each application previously defined in the product definition stage of the service life cycle, the fulfillment stage focuses on the propagation and enforcement of network policy change requests, ensuring that each application receives the appropriate network resources.

Within the network forwarding plane or infrastructure layer, two tools are used to enable the fulfillment of these policy change requests: application identification and application prioritization.

- **Application identification** of an individual flow within the mishmash of enterprise traffic can be problematic. Applications may use customized (nonstandard) ports or may share a control protocol with other applications. For example, Session Initiation Protocol (SIP) can be used to control the delivery of multiple applications, including voice and video services. For the Application-Aware Network, DI technology is a crucial technique required to discern an application's identity through examination of its message flows. Deep inspection has the ability to thoroughly investigate an IP packet by analyzing data associated with the higher OSI layers—which routers typically do not consider—to distinguish one application from another. True application awareness ensures that differentiated handling of traffic is applied correctly against each application.
- **Application prioritization** already exists in today's managed business services. The classical approach to prioritization is to allow the customer (or managed customer premises equipment) to mark traffic into one of several classes of service, according to a predefined policy. Each class of service is then given weighted access to network resources with priority given to the highest class of service. Dynamic policy is the application of adjustable, service-specific policies on the network that ensure appropriate application prioritization. With the introduction of the enterprise portal in the service definition stage, the Application-Aware Network can dynamically support changes in the prioritization policy and even in the total available network bandwidth. Policy enables the carrier to mix high value services that typically have stringent network requirements with lower value services that are more tolerant of network characteristics, such as delay or jitter. With dynamic policy imposing traffic hierarchies based on application type, the carrier can increase the total utilization of the network while preserving a high level of user satisfaction.

Visualizing and Assuring Network Performance

With the infinite potential of flexible service packaging, both service providers and customers will require per-application visibility of network performance. Traditional SLA metrics are insufficient as proof that the offered micro-SLAs have been met. For the Application-Aware Network, this service assurance represents the most critical aspect of the service life cycle. Per-application assurance and reporting tools are required to validate that the application-based micro-SLAs are met at all points across the carrier's domain.

Conventional methods to provide SLAs include:

- **IP-based probes**—emulating applications to measure packet loss, round trip delay, jitter.
- **Separate application intelligent appliances**—deployed to monitor the quality of specific applications.

IP probes represent a simple performance reporting tool that can be deployed to provide coarse information on packet loss, latency, and jitter. It is worth noting that these results do not directly relate to any specific application's performance or the impact of the recorded variables on the application performance. Another caveat concerning IP probes is the need for the performance monitoring system to interoperate with the heterogeneous range of devices found in a typical service provider's ecosystem. Despite these conditions, IP probes can still provide a simple and scalable method of application performance reporting. Standards-based approaches such as the IETF's Two-Way Active Measurement Protocol (TWAMP) can be deployed to simplify impacts on these performance management systems.

Application intelligent appliances provide more accurate quality measurements than IP probes, but tend to be more costly due to the need to purchase and deploy dedicated hardware across multiple sites. As DI technology has already been deployed to scrutinize each packet, DI technology within the Application-Aware Network can be assigned with the additional responsibility of generating per-application reporting. The combination of subscriber-id with the DI technology assigned application-id enables the carrier to generate customer reporting without specific inline reporting hardware.

For enterprise customers, reporting information can be used to confirm that applications have received the anticipated service levels, validating the network's performance. Reports can also be used to identify application usage trends or even rogue application usage within the enterprise. With the ability to view packets at a granular level, DI can detect applications or threats that might otherwise go unnoticed. DI can detect protocols that have previously been able to evade simple signature and port-based detection mechanisms. For example, traffic hungry applications such as P2P can implement a variety of techniques to avoid identification such as port masking (using a common port) or encryption; DI can help mitigate these techniques. Per-application reports allow the enterprise to maximize the use of application prioritization to protect business critical applications while avoiding the "just buy more" approach to congestion management.

For carriers, the increased visibility through application awareness enables network planning tasks, including traffic engineering, demand forecasting, and network investment decisions. DI technology enables service providers to mine application-based statistics in an effort to hone their existing services or to identify new tailor-made service opportunities.

How Juniper Networks Enables the Application Intelligent Network

Juniper Networks offers a large number of technology options to build the Application-Aware Network. An Application-Aware Network solution is comprised of components that fall under two broad categories:

- Control functions
- Traffic processing systems.

Control Functions

Juniper Networks SRC Series Session and Resource Control Modules enable service providers to deliver differentiated products and services through dynamic allocation of network resources to users or applications as required.

The SRC Series connects the service layer with the network layer of service provider networks by providing a feedback loop between applications, users, and the network. Its open interfaces integrate with any network and any service offering, regardless of the demand location. The SRC Series allows service providers to generate additional revenue on their existing network infrastructure by adding dynamically-activated services.

Unlike competitive solutions that use authentication, authorization, and accounting (AAA) for policy management or solutions that deploy static policy enforcement, the SRC Series delivers granular dynamic policy enforcement at a per-service level. This enables service providers to deliver revenue generating services on top of existing sessions. Because the SRC Series readily interfaces with existing subscriber management data bases, available network resources can be mapped to subscriber and service profiles. The result is differentiated services that are based on dynamic allocation of network resources with the ability to provide service-specific accounting.

Additionally, the reporting system generates various usage reports that provide visibility into the network operation. This is a vital component of the solution, as it provides feedback on the effectiveness of the policies configured. Juniper Networks Application Usage Manager (AUM) can perform this function.

Traffic Processing Systems

Juniper Networks traffic processing systems used within the Application-Aware Network deliver two key components.

DI Engine: leverages Juniper Networks leadership in security and policy management to provide stateful, subscriber-aware application monitoring, enabling advanced application-specific service levels and QoS treatment. Either standalone Juniper Networks IDP Series Intrusion Detection and Protection Appliances, or integrated Juniper Networks MX Series 3D Universal Edge Routers and Juniper Networks SRX Series Services Gateways, or both, can perform this function.

Enforcement System: enforces the action of the policy based on the classification performed by the DI engine. MX Series routers can perform this function.

Juniper also offers several Juniper Networks Junos® operating system-based integrated applications that greatly simplify the network edge and improve operational efficiency by reducing the number of network elements and operating systems within the network.

- **StreamScope eRM**: multilayer monitoring and analyzing of video streams as they cross the IP network to ensure service quality and facilitate rapid fault detection and isolation
- **Telchemy ePM**: multilayer monitoring and analyzing of voice and IP services as they cross the IP network to ensure service quality and facilitate rapid fault detection and isolation.

Juniper's Commitment to Open Development

The Application-Aware Network fits Juniper Networks corporate commitment to a top-to-bottom open development stack, as show in Figure 1.

- Juniper Networks Partner Solution Development Platform (PSDP) helps providers and others create an enhanced high-performance network service delivery infrastructure. The PSDP is also a key innovation enabler that is driving convergence, increasing operational efficiency, and reducing capital investment. The application awareness capabilities are exposed to applications built using the PSDP, which can reduce the overall time and cost to deliver new innovations. Using the PSDP, Juniper supports a wide variety of router integrated services, including video monitoring, VoIP, and IP service monitoring, Dynamic Application Awareness (policy driven DI technology), and intrusion prevention system (IPS), among others. The ability to physically and functionally integrate these services with the routing infrastructure eliminates the cost and complexity of qualifying, deploying, maintaining, and sparing standalone appliances. Furthermore, because all of these services are licensed Junos OS elements, administrators can realize increased operational benefits. By implementing a single operating system with a single release train (ease in management), a common set of administration, configuration, and management tools are utilized.
- Juniper Networks Open IP Service Creation Program (OSCP) further extends openness of the Session and Resource Control Policy Engine (SRC-PE), which provides the intelligent link between routing infrastructure and applications using standards-based application to network APIs. The OSCP enables the customization of services based on network capabilities. This ensures that the network and application work together for each subscriber to create a positive user experience.
- TM Forum's IPsphere framework transforms the way that next-generation network services can be created and delivered by multiple stakeholders. Using Web services principles, the Juniper Networks-developed TM Forum IPsphere framework defines a sophisticated business developer layer that automates offer, purchase, and provisioning of service components among multiple stakeholders.

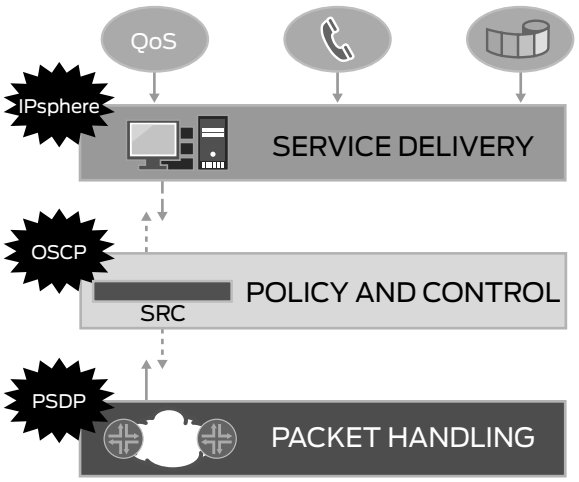

Figure 1: Juniper Networks open software development process

Conclusion

An Application-Aware Network combines application awareness with dynamic network policy, enhancing traditional managed service offerings and simplifying the act of balancing the needs of applications against available network resources. For carriers, this ability to synchronize network resources with end customer requirements ensures a more efficient use of network capacity and infrastructure, providing better return on investment. Through greater visibility into network use, enterprise customers are empowered to directly customize their own service definitions—in a highly granular and flexible fashion—as quickly as their shifting business needs demand it.

For service providers looking to take advantage of the benefits of an Application-Aware Network, Juniper Networks offers a rich set of tools extending from the infrastructure that is the foundation of all services, through policy and control, and service delivery layers. Developer programs, including the OSCP and the PSDP, demonstrate Juniper Networks commitment to the philosophy of a top-to-bottom open development stack. This openness gives carriers new paths to innovation, which, when combined with open standards-based interfaces, enable new capabilities such as those required for Application-Aware Networks to be easily activated across Juniper Networks infrastructure. As a result, this strategy limits the risk and operational costs associated with deploying services demanded by business. Furthermore, tight physical and functional integration of these new capabilities within the router and routed infrastructure permit carrier deployments without adding operational cost and complexity.

About Juniper

Juniper Networks, Inc. is the leader in high-performance networking. Juniper offers a high-performance network infrastructure that creates a responsive and trusted environment for accelerating the deployment of services and applications over a single network. This fuels high-performance businesses. Additional information can be found at **www.juniper.net**.

Corporate and Sales Headquarters
Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER (888.586.4737)
or 408.745.2000
Fax: 408.745.2100
www.juniper.net

APAC Headquarters
Juniper Networks (Hong Kong)
26/F, Cityplaza One
1111 King's Road
Taikoo Shing, Hong Kong
Phone: 852.2332.3636
Fax: 852.2574.7803

EMEA Headquarters
Juniper Networks Ireland
Airside Business Park
Swords, County Dublin, Ireland
Phone: 35.31.8903.600
EMEA Sales: 00800.4586.4737
Fax: 35.31.8903.601

To purchase Juniper Networks solutions, please contact your Juniper Networks representative at 1-866-298-6428 or authorized reseller.

2000328-003-EN Aug 2010 Printed on recycled paper