WHITE PAPER

The Public IP Network

Ericsson IP Infrastructure

October 1999

Ericsso n bal

From The Internet to the Public IP Network

The growth of the Internet in the past five years, fueled by the popularity of the World Wide Web and increased use of intra- and inter-company e-mail, has been unprecedented, and has shaken the foundations of the data communications & telecommunications industries. Both industries are aggressively moving towards a common future, based on IP protocol and a packet-switched infrastructure model - a vision of "IP Convergence".

Today's Internet is an ad-hoc overlay built above existing telecommunications and wide-area data communications networks. Entrepreneurial Internet Service Providers (ISPs) have for the most part built it around existing carrier networks, with minimal participation from carriers. The Internet today scores points for its ubiquity - it is accessible from virtually anywhere on the planet via modem, and over existing high-speed telecommunications services in all major markets. But it falls far short of meeting the reliability and performance of traditional wide-area data communications services. The vision of "IP Convergence", along with many bold predictions on IP telephony, e-commerce, and the transformation of the workplace, could not be fulfilled by today's Internet.

New Public IP Networks are already being built to supplant and replace not just today's Internet, but also major elements of the telecommunications network. The Public IP Network is not being built simply as an overlay by independent ISPs; it's an integrated part of the carrier network, built and delivered by innovative, entrepreneurial carrier-service provider hybrids. The Public IP Network will maintain the ubiquity of today's Internet, augmented with the reliability of a telecommunications service and with performance enhancements that enable the most aggressive applications to operate seamlessly. Today's Internet has grown in a short time to a multi-billion dollar market, has changed the shape of the communications industry and has become a social phenomenon. But today's Internet is only the beginning.

Converged IP Services

To some, the term "IP Convergence" describes the transition from today's telecommunications infrastructure technologies (ATM/Frame Relay Switching, SONET/SDH TDM) to new, IP-centric solutions. To others, it describes a transformation of today's multiprotocol service model to a bundled service model based on IP. The former concept - "IP Infrastructure Convergence" is technologically fascinating and will certainly have a major impact on carrier backbones over a 5-10 year period. But the latter concept "IP Service Convergence" is happening far more quickly and will have a much more visible near-term impact.

The concept of IP Service Convergence is a simple one. Today, the typical business contracts with a variety of different providers of diverse communications services. A medium-sized corporation could have different providers for basic telephony (POTS, FAX, etc.), videoconferencing service, private wide-area networking, remote LAN access and Internet access. Each of these services potentially requires its own WAN connection and its own CPE; each service also likely demands specialized expertise from the subscriber. What IP Service Convergence promises is a simplified model, where all the services above (along with new ones not yet imagined!) are delivered via one connection (optionally redundant) to an IP network. Via one wire from one service provider, subscribers will be able to access a full menu of communications services.

Today's Internet can be used to prototype new services, but is not equipped to deliver converged IP services on a large scale. The routers used to forward IP traffic within the Internet today, and the IP protocol suite running as a distributed operating system within these routers, is weak in two key areas: performance and topology. The New Public IP Network will rely on innovations in both areas to enable a bundled service model.

Performance

Performance issues in the Internet today are easy to observe -- anyone who has used a browser to surf the Web has firsthand experience with them. Statistical multiplexing of bandwidth is an inherent characteristic of packet-switched networks, and contention for use of the network creates variations in latency ("jitter") and non-zero data loss. IP routers today make no effort to improve upon these issues for any traffic types; all service is unreserved, best effort and connectionless. The converged service model will require more intelligent handling of different traffic classes within the Public IP Network, insuring that certain applications (e.g. voice) get prioritized handling over other traffic types.

Topology

Topology is a more subtle issue than performance in today's Internet. Today's routers maintain a single topology database (the "routing table") for the global Internet; from any specific router to any point in the Internet, there is one unique path, dynamically maintained through exchange of topology information between routers. The protocols used to exchange this information today make no accommodation of "private" traffic (e.g. a corporation's own WAN backbone), are unable to re-route traffic around congestion points and do not easily allow topologies to be constrained by commercial factors. The Public IP Network will require enhancements to topology management allowing IP-based Virtual Private Networks (VPNs), topologies based on Class of Service (CoS) and easier implementation of commercial constraints (e.g., preferring one backbone alternative over another due to lower transport costs).

With enhanced Quality of Service/Class of Service (QoS/CoS) capability and more flexible topologies supporting IP-VPNs, the Public IP Network will be able to deliver the full range of telecommunications and data communications services available today, with "public" Internet access as a bundled feature. New architectures, products and standards are being developed that will be used to build the Public IP Network and make IP Service Convergence a near-term reality.

Building the Public IP Network

Like today's Internet (and today's telecommunications networks), the Public IP Network will not be one single network, but a mesh of parallel networks interconnected at major "peering points". Each of these parallel networks will be an IP network on its own, and may in turn be broken into different elements owned and operated by different carriers. The structure of each parallel IP network can be broken into two major elements.

Backbone Networks

Large carriers and service providers will operate backbone networks nationally and internationally. These will always run over large fiber plants, typically with a DWDM layer enabling terabits per second of aggregate bandwidth in the network. These networks are always built as a mesh of interconnected IP routing nodes richly interconnected by point-to-point links, although there are numerous approaches to how the IP network is layered over the fiber optic backbone. These approaches can be loosely grouped as follows:

IP over ATM Multiservice Backbones. These networks use ATM Switching to multiplex IP traffic with other traffic types across the backbone. Routers interconnect via point-to-point virtual circuits over the ATM backbones; in some cases, ATM switches are active routing nodes as well, using MPLS technologies (more on this below). The advantage of this architecture is its ability to support existing backbone traffic (non-IP data, non-data) alongside IP, enabling easy migration to the New Internet. However, it carries a price both in the loss of throughput to ATM overhead and in the cost of ATM network management, and not all network architects believe this price is offset by ATM's benefits.

IP over SONET/SDH Backbones. These networks eliminate the ATM layer, implementing point-to-point links between IP routers directly over SONET/SDH rings (which in turn run over DWDM). If non-IP traffic is carried at all, it is carried over separate point-to-point connections in the same SONET/SDH structure. Like the IP over ATM approach, this can be implemented today, based on proven standards.

IP over DWDM Backbones. These networks exist today only in theory; the idea is to replace the SONET/SDH layer with a new, lightweight physical layer mapping IP traffic directly to DWDM fibers. The argument for this approach is logical; much of the structure of SONET/SDH is optimized for circuit, not packet switching, and a simpler approach optimized for IP packets will result in better price-performance. IP/DWDM nodes could use IP over SONET as an interoperable interface to IP over DWDM networks and optimize performance for IP traffic.

The vision of IP Infrastructure Convergence calls for rapid migration to IP over DWDM backbones. In reality, all three types will exist; migration to a common approach in the backbone may take many years.

Aggregation Networks

Both large and small carriers and ISPs will operate aggregation networks within service areas (as large as a country or even a few states, or as small as an industrial park). Logically, an IP Aggregation Network will look like a funnel: thousands or even tens of thousands of subscriber connections will be transported via carrier switching & multiplexing networks into an IP Aggregation Point, where powerful Aggregation Routers map the subscriber traffic streams to backbone connections. Functionality in these new IP Aggregation Networks can be viewed in three domains:

Subscriber Access. The connection between the subscriber and the IP aggregation point will logically be a point-to-point IP connection, and carrier networks will use diverse Layer 1/Layer 2 multiplexing and switching technologies to deliver thousands of these connections within a service area. Alternatives will range from high-speed leased lines and ATM/Frame Relay PVCs to IP/PPP/ATM connections over xDSL networks to wireless and cable modem networks. IP Aggregation Routers will need to offer tens of thousands of virtual IP interfaces on a variety of physical port types to insure easy integration with these diverse networks.

Subscriber Traffic Processing. At the IP aggregation point, many thousands of subscriber streams are being terminated and aggregated. The IP Aggregation Router must be able to quickly classify received packets according to pre-defined policies, per subscriber and per application. This must be more than simple classification into a finite pool of priorities - each traffic class may require mapping to a different VPN, and may require unique shaping and prioritization.

Backbone Integration. A primary function of IP Aggregation Routers is to route the aggregated traffic onto IP backbone networks. This demands comprehensive support for routing protocols (OSPF, IS-IS and BGP-4) as well as the ability to map traffic to supported backbone CoS levels and traffic-engineered topologies. Two new standards for IP networks nearing completion by the Internet Engineering Task Force (IETF) will play a key role here: DiffServ and MPLS.

The new Differentiated Services, or DiffServ, standard is used to allow IP traffic to be marked for preferred handling by the network. DiffServ redefines a byte in the existing IP header (the Type of Service, or ToS byte) to include a 6-bit "DS" field indicating the service requirements for the packet. DS-capable nodes will examine this field on each packet and condition forwarding operations according to its value. Of the 64 possible DS values, the IETF intends to define up to 32 as "global" service classes and to leave the other 32 open to network specific definition. Nodes can also rewrite DS values in transit. The DS value will allow certain packets to be prioritized ahead of others at each network node, reducing jitter and increasing "goodput" (throughput of payload packets) for specific traffic streams (albeit at the expense of less important traffic).

The Multiprotocol Label Switching (MPLS) initiative is much broader than DiffServ, and has evolved into a family of standards within the IETF. The basic concept of MPLS is to prepend the IP packet with an extra label, allowing intermediate nodes in an IP network to perform simple label processing to determine the packet's egress path instead of the more complex lookup normally used to find the destination. Since the forwarding within an MPLS "cloud" is based only on the label, it supports IP-based Virtual Private Networks (VPNs) readily. And because the process that manages the labels in an MPLS network is decoupled from the basic topology processes of the network, MPLS can be used to augment the basic topology with new paths - a capability called "traffic engineering". By engineering extra "Label Switched Paths" (LSPs) for certain traffic classes, and by using MPLS LSPs to optimize use of a complex mesh, network operators can improve goodput and jitter across the board.

MPLS and DiffServ are both enhancements to basic IP networks, but they don't make any presumptions about the underlying protocol sandwich. Standards have been defined to allow MPLS to run directly over ATM or Frame Relay backbones, as well as over any network carrying IP traffic in PPP frames. And DiffServ operates strictly within the IP packet; it can be used in any IP network. Infrastructure convergence isn't necessary to take advantage of these new standards - all that's required is that IP routing nodes within the infrastructure be upgraded to support them.

Major telecommunications equipment providers are rolling out a new generation of solutions designed for these Public IP Networks. Ericsson, the global leader in mobile communications infrastructure, is delivering a powerful set of products for building Public IP Networks, including:

Backbone Network Solutions

Ericsson's portfolio spans the variety of backbone alternatives planned for Public IP Networks. The AXI 520 IP Core Router can be used to build IP core networks at up to OC-48/STM-16 (2.5 Gbit/second) speeds over SONET/SDH connections, and fully implements new MPLS and DiffServ standards. The AXD 301 IP/ATM Switch supports both ATM cell switching and MPLS label switching, and scales to over 100 Gbits/second of switching capacity. And Ericsson's DWDM solutions can be used to build metropolitan-area backbones for IP as well as other telecommunications technologies.

Aggregation Network Solutions

Ericsson offers a broad set of solutions for new IP access and aggregation networks. Optimized access nodes serve narrowband access; xDSL access and new wireless IP access systems. And Ericsson addresses the market for high-speed fixed access and Point-of-Presence aggregation with its new AXI 540 Edge Aggregation Router.

The AXI 540 Edge Aggregation Router

Service providers and carriers building the Public IP Network have very specific requirements for their IP Aggregation Routers. These are:

- Can the Product Route?
- Will the Product Scale?
- Is the Product Reliable?
- Will the Product Deliver Advanced IP Services?

The Ericsson AXI 540 Edge Aggregation Router and its IPaction routing software have been optimized to address these requirements.

Can the Product Route?

Ericsson's IPaction Routing Software supports a complete range of routing protocols: RIP, OSPF, IS-IS, IBGP/EBGP, DVMRP, PIM-D and PIM-S. The IPaction software suite also implements a comprehensive range of advanced BGP features, including Route Mapping, Communities, Route Flap Dampening and Route Reflectors, as well as a broad set of other value-add features such as BOOT-P, DHCP Relay, Proxy ARP and NTP. Major ISPs worldwide have tested the IPaction software suite, and have demonstrated full interoperability with existing router backbones.

Will the Product Scale?

The architecture of the AXI 540 Router supports scaling in every important dimension. It scales throughput to over 20 Gigabits per second or over 20 Million packets per second, aggregated from as many as 40,000 virtual IP interfaces per system. It will support as many as 400 IBGP/EBGP peers and can accept over 6 million routes from these peers. The active routing table maintained at each port can hold as many as 400,000 network prefixes or 64,000 S-G pairs for multicast routing. The overall system supports up to 50,000 classification filters at each port, mapping traffic to as many as 120,000 independent queues across the systems switch fabric. The inherent scalability of the AXI 540 architecture addresses every key requirement for long-term growth.

Is the Product Reliable?

Ericsson's AXI 540 Routers deliver fully resilient routing at the hardware, software, system and network levels. The AXI 540 platform uses redundant power supplies, redundant fan trays, redundant 20 Gigabit-per-second switching fabrics and redundant Route Processors to insure maximum availability at the hardware levels. Ericsson's IPaction software partitions all major capabilities into separate software tasks, running in protected memory over a multitasking operating system, thereby insuring that a fault in one task can't affect operation of others. And the overall design of the AXI 540 uses SONET/SDH alarms and multipath routing features to quickly identify and re-route around any faults in the network.

Will the Product Deliver Advanced IP Services?

Delivering routing that works is a matter of software development and testing; delivering scalability and reliability requires a solid architecture and implementation. But to truly enable IP Service Convergence in the New Public IP Network, a product must incorporate visionary breakthroughs in technology. The AXI 540 Edge Aggregation Router is built around such breakthroughs.

The AXI 540 design uses a custom silicon processor on every physical port to examine each received packet at wire speed, classify it according to flexible rules and determine its next-hop router using patented search techniques in the full routing table. Packets matched to different traffic classes can be shaped according to QoS policies and encapsulated for backbone VPNs based on the MPLS protocols. The design also handles transformation of transmit packets to different formats and on-the-fly rewrite of DiffServ values for mapping to backbone CoS values.

By using its powerful classification and shaping capabilities, the AXI 540 Router can map traffic to a range of service levels based on the subscriber's user profile and the type of application. By using DiffServ and MPLS, the same traffic can be mapped into backbone CoS reservations and traffic-engineered paths. And with wire-speed silicon routing for all traffic types, the AXI 540 Router insures that every application from every subscriber sees the performance it needs. This powerful capability set, together with the scalability and reliability of the system and the power of IPaction routing software, make the AXI 540 Router the clear technology leader for new IP aggregation applications.

Conclusion

Today's Internet is a phenomenon. Its emergence in the last four years has shaken the foundations of industries, created and eradicated fortunes and catalyzed social and economic change. But today's Internet is only the beginning of a new era of IP-based communications - the era of the Public IP Network.

The Public IP Network will transform the communications industry and will enable convergence of a variety of services to an IP-centric model. It will require both powerful "core routers" for IP backbones and a new class of IP Aggregation Routers for subscriber consolidation and service delivery.

Ericsson, the leader in mobile communications worldwide, is delivering a comprehensive set of solutions for new Public IP networks, from the core to the edge. The AXI 540 Edge Aggregation router is a key member of this powerful family of next-generation networking solutions.

Ericsson Inc.
Datacom Networks
77 South Bedford Street
Burlington, MA 01803
www.ericsson.com/datacom