

Executive Summary

Network bandwidth usage is growing at annual rates of 60 percent or more because of the rapid adoption of video and cloud services. Control plane traffic is increasing even more quickly as the Internet moves from serving fixed locations and fairly static information sources to one characterized by personalized, media-rich applications (app), and mobile services. Service providers that are already concerned with data plane scalability must also add control plane scalability to their list of network planning issues.

To address the new variables in their networks, operators need solutions that are scalable in multiple dimensions (data, control and service). The Ericsson SSR 8000 family of products is a flexible platform for building the next generation of converged IP networks and services. It provides the scalability required to meet future data and control plane traffic requirements.

ACG Research conducted an analysis of the sources of growth in control plane traffic: increases in end-user connections, device types, apps, app policy and control requirements, and mobility management requirements. This study projects data and control plane traffic growth at a typical packet core node and analyzes the scalability of the SSR 8020 IP service delivery platform and of a competing service router when configured to serve this traffic growth. The study finds that the SSR 8020 has lower total cost of ownership (TCO) by 66 percent and higher scalability by two to more than three times that of the competitor's service router.

Key Takeaways

The Ericsson Smart Services Router (SSR 8000) provides the control plane scalability required to meet users' demands for personalized, socially-inclusive, media-rich applications and mobile devices. Compared to a competing service router, the SSR 8020 has the following:

- 66% lower cumulative five-year TCO
- 66% lower CapEx
- 64% lower OpEx
- 2–3+ times scaling advantage when subjected to rapid increases in IP sessions and total bandwidth requirements
- The cost advantages of the SSR are due primarily to better software and hardware architecture

Introduction

As network usage moves to cloud, personal and mobile services control plane traffic is exploding. Router vendors and network architects who are already struggling to meet the bandwidth requirements of video traffic must now add control plane scalability to the network design equation.

The rapid increase in control plane traffic is driven by the move from an Internet that served fixed locations and fairly static information sources to one where users demand personalized, socially-inclusive, media-rich apps and mobile devices. Growth in fixed and mobile broadband subscribers, the number and type of network devices and applications are all fundamental drivers of increasing control plane traffic. They combine to produce a multiplier effect that makes control plane traffic grow faster than data plane traffic.

The Ericsson SSR 8000 family provides the data and control plane scalability needed to deliver the services users are demanding. The SSR 8000 provides a smooth transition from the Ericsson SmartEdge Router by using the same transport technologies, policy control, and management interface employed by SmartEdge.

The following sections examine the drivers and sources of control plane traffic growth, provide a projection of expected control plane traffic growth rates, describe the architecture of the SSR 8000, and analyze the scalability benefits of the SSR 8000.

Growth in Control Plane Traffic

The demand model for data plane traffic growth is relatively simple as compared to the demand model for control plane traffic¹. Data plane traffic growth is driven by increases in the number of end-user connections and by the rate of change in average bandwidth usage per connection. Control plane traffic, in contrast, has a very complex growth model. Just as with data plane traffic growth, it is driven by the number of end-user connections. However, unlike data plane traffic growth control plane growth is driven by the number and nature of each transaction, call or data flow handled by the network. The number and types of network uses are determined by:

- Number of end-user connections
- Number of different end-user device types
- Number of apps
- App policy and control requirements; Deep Packet Inspection is one example
- Mobility management requirements

¹ In this report traffic is measured at the peak usage period because network elements such as routers must be sized to meet peak demand. Data plane traffic is measured as the bandwidth (Gbps) of the user data (payload). Control plane traffic is measured in two dimensions: 1) Number of IP sessions and 2) signaling transactions per second.

The following sections discuss how the control plane traffic is increasing in each of these five dimensions. Consequently, a growth multiplier effect is driving control plane traffic growth at a much higher rate than that of data plane traffic growth. Data plane traffic growth is expected to increase in excess of 60 percent² per year. It is therefore reasonable to expect that control plane traffic will expand at annual rates well in excess of 100 percent.

End-User Connection Growth

End-user connection growth is driven by the continued adoption of broadband service. Wireline services lead the first wave of broadband development. Although wireline broadband development is still robust wireless broadband is leading the second wave of broadband development. In the U.S., for example, wireline broadband grew at 10 percent during the last year; wireless broadband grew at 58 percent.

There is much room for broadband connection growth. For example, the Federal Communications Commission defines minimum current performance standards for broadband connections as an advertised downstream speed of 3 Mbps and an upstream speed of 768 Kbps. Thirty-five percent of U.S. households meet this standard. The growth potential is therefore 65 percent of U.S. households. On a global basis U.S. broadband development is judged to be near the middle of all industrialized countries.

U.S. wireless market statistics also show that most broadband connection growth has yet to come. Currently, there are 328 million wireless connections in the U.S., which is 104 percent of the U.S. population. Twenty-nine percent of these connections are smart phones or PDAs; five percent are wireless-enabled laptops, notebooks, tablets or wireless broadband modems. This implies that 66 percent of U.S. mobile wireless connections have yet to move to wireless broadband.

The potential for control plane traffic growth is even greater than that for broadband connections. As a rule of thumb a smart phone uses three times the control plane traffic of a mobile feature phone. Therefore, if there are two potential smart phones³ for every existing

² ACG Research projection

³ Sixty-six percent additional market potential implies that approximately two-thirds of the market potential remains or two potential connections for every existing connection.

M2M Connection Growth

Machine to Machine (M2M) growth is projected to explode. One projection is that one billion connections will be in service by 2015. Types of devices include:

- Smart grid transponders
- Utility meters
- Security sensors
- Location based systems
- Inventory control transponders
- Medical sensors
- Smart cards
- Traffic control systems
- Home appliance transponders
- Environmental sensors

M2M applications are different from personal communications in almost all respects:

- Service payloads are much smaller than signaling and overhead
- Nearly all traffic is upstream; downloads are rare
- Data loads are light; smart electric meters use less than 1 MB per month
- Devices operate autonomously, should be maintenance free and some battery life requirements are 20 years. The strategy to preserve battery charge of a connected device has been to sever its connection to its corresponding server in idle mode and re-connected when invoked. This strategy applies further pressure on control plane scaling requirements of the network and its components.
- Many devices have long inactive/dormant periods; alarms and power outage monitors may never be activated
- Periodic status reporting generates many signaling events but no data payloads

one, there is a potential six times rate of increase in the existing control plane traffic.

Device Proliferation

The large number of different device types also is a fundamental driver of control plane traffic in that messages must be sent and acknowledged before data flows that meet the specific requirements of each device type can be established. More than 630 different handsets and devices are manufactured by 32+ companies for the U.S. market, and in the last year vendors brought an additional 120 new smart phone models to the market according to the CTIA⁴.

Device types include:

- Smart Phones
- Tablets
- PCs
- Laptops and Notebooks
- Mobile Hot Spots
- Gaming Devices
- PDAs
- M2M
- Wireless Modems

Application Growth

The number of network connections and the increasing diversity of device types provide the foundation for control plane traffic growth. The rapid growth and acceptance of network apps provides further leverage for control plane traffic to grow more rapidly than data plane traffic. Figure 1 shows a global projection of expected app downloads.

⁴ Other estimates on the number of different devices are much higher. For example, AT&T says it has certified more than 1,000 devices for use on its network.

M2M control plane and data traffic characteristics, consequently, are dramatically different from those of personal communications. Control plane traffic dominates data traffic.

M2M apps are chatty. For example, keep alive packets and status checks are frequently made to ensure that monitoring devices are in working order. Also, since most devices have long dormant periods a new communications session must be established each time a device uploads data. Location-based services, in particular, generate a great deal of control plane traffic. For example, taxi location systems report each taxi's location every 10 to 15 seconds, but each transaction payload is a few hundreds of bytes.

Some M2M apps have the potential to create signaling storms. For example, a major power outage could trigger millions of smart grid devices to send outage reports.

The business case for development of millions of M2M applications is simple and compelling. M2M devices are used to monitor and control field services remotely and at very low cost. Low-cost M2M devices and communications networks reduce expensive resources such as taxis, delivery vans, snow plows, ambulances and drivers, service technicians, EMTs, and power company linemen. Such a compelling business case is a driver for the rapid and large-scale deployment of M2M devices. This in turn drives rapid and large-scale growth in control plane traffic.

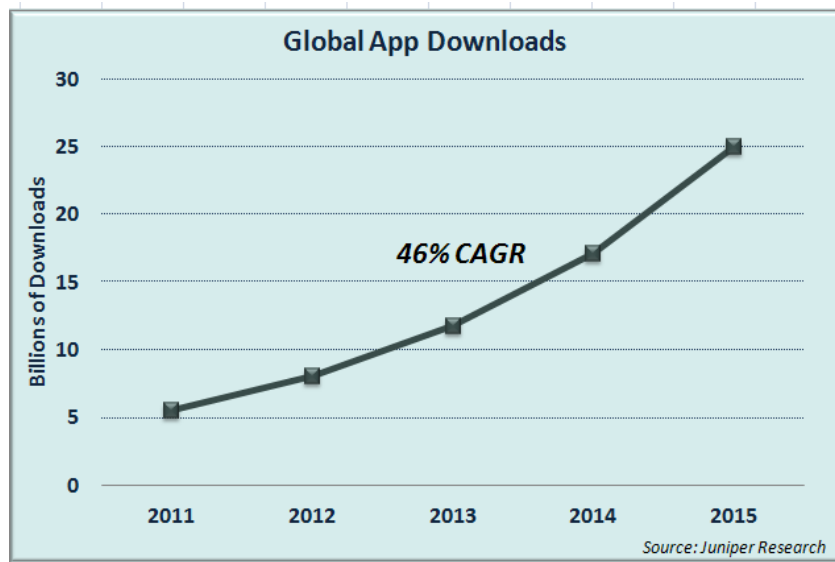


Figure 1 – Global App Downloads

The projection has a compound annual growth rate (CAGR) of 46 percent. Control plane traffic, however, will grow faster than the rate of growth in app downloads. See Figure 2.




Figure 2 – Control Plane Traffic

Control plane traffic is growing at 131 percent CAGR compared to projected app growth of 46 percent CAGR. Figure 3 shows why this is so.

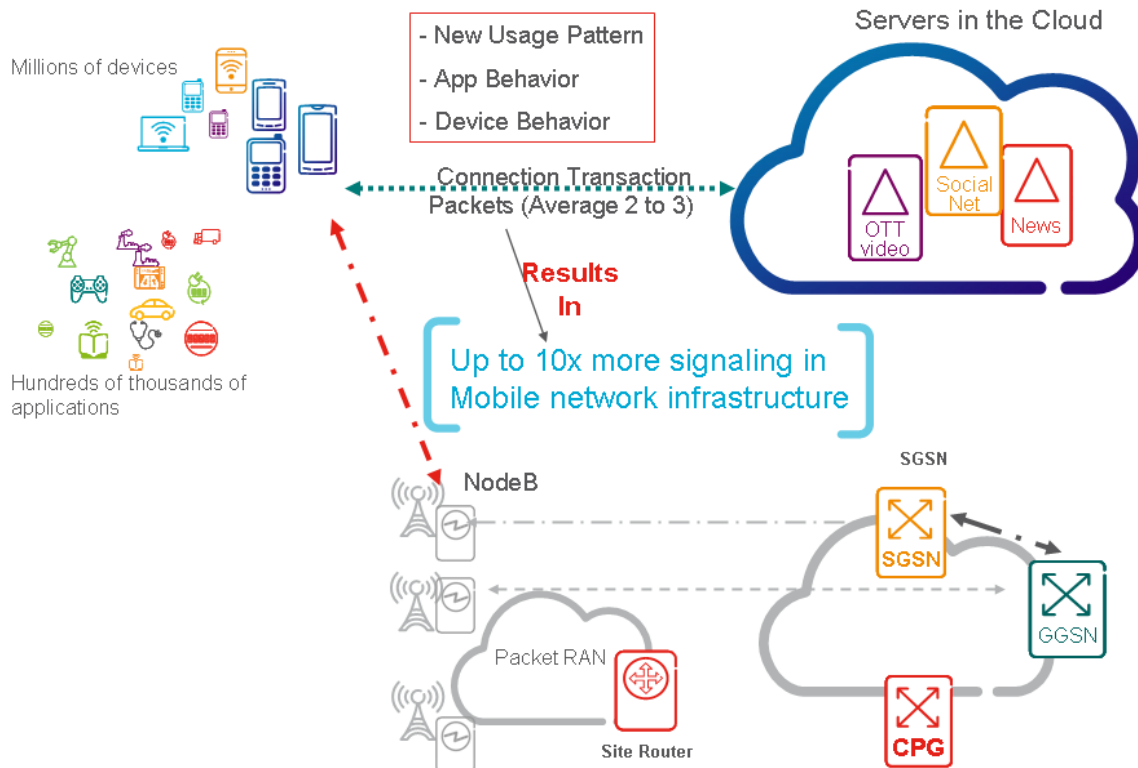


Figure 3 – Drivers of Control Plane Traffic

Signaling events are triggered by new usage patterns created by apps. For example, watching a YouTube video has a different usage pattern than simply viewing a web page. Watching a YouTube video begins by viewing web pages to locate the desired video. This signaling is the same as for viewing any other web page. However, additional state change signals are required when the video link is selected and when it is subsequently exited.

New application behaviors also increase control plane requirements. For example, the FindMe application allows friends to find each other on a digitally-displayed map. The app has substantial control plane requirements, cell I.D.s, GPS information and network information that must be exchanged. Most importantly, updates are traded each minute so that the friends can observe each other's movements.

Device behavior also generates additional control plane traffic. Specifically, many apps require an always-on behavior mode. Conversely, device manufacturers are working to minimize device uptime to conserve battery power. Consequently, devices periodically go into an idle or standby state and then to an active state so that battery power is conserved even though the app has the appearance of being always on. This creates large quantities of control plane traffic relative to data traffic.

Figure 4 illustrates the control plane traffic flows for an LTE network.

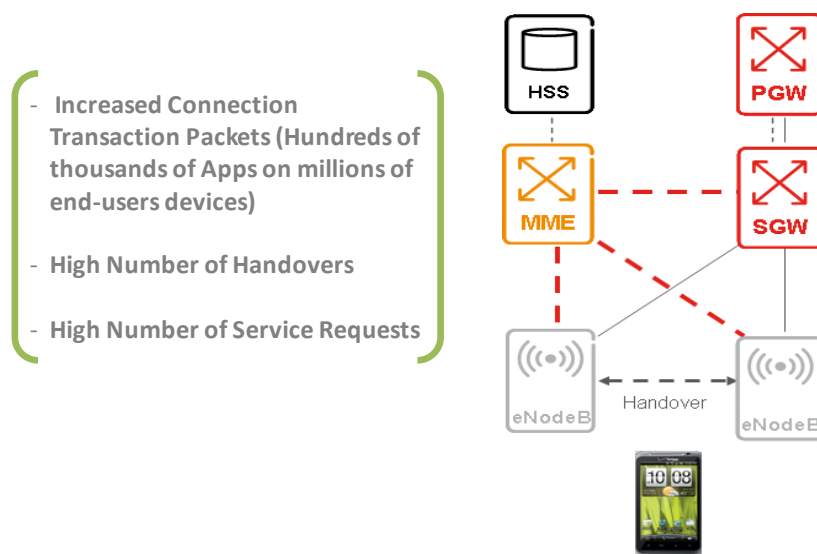


Figure 4 – LTE Smart Phone Impact on EPC Signaling

In order to maximize cell data throughput high cell handover rates must be maintained. At the same time, device inactivity timers should be adjusted to the lowest possible settings to prolong device battery life. The combination of high handover rates and short device activity times maximizes the evolved packet core (EPC) control plane traffic.

The net effect of rapid growth in mobile apps is to accelerate growth in control plane traffic. The rapid adoption of smart phones, apps that require regular updates, evolved pricing and segmentation models and the increased control over network resources are all driving control plane traffic increases.

The remainder of the paper analyzes the ability of the SSR 8000 to economically accommodate high control plane traffic growth (scaling) by comparing the TCO of the SSR 8020 versus that of a competing service router when both routers are subjected to high rates of control plane traffic growth.

Total Cost of Ownership Analysis of SSR 8020 Control Plane Scaling

The control plane scaling advantages of the SSR 8000 are analyzed by comparing the TCO of the SSR 8020 with that of a leading competitor where the routers are deployed as service routers in the EPC. The comparison is made by projecting control and data plane traffic requirements at a packet core node over five years, configuring the SSR 8020 and the competing service router to meet the projecting traffic requirements, and then computing capital expense (CapEx) and operations expense (OpEx) associated with the equipment configuration of each router.

The Role of the Service Router in Processing Control Plane Traffic

Figure 5 illustrates the many interfaces between the service router and other elements of the EPC.

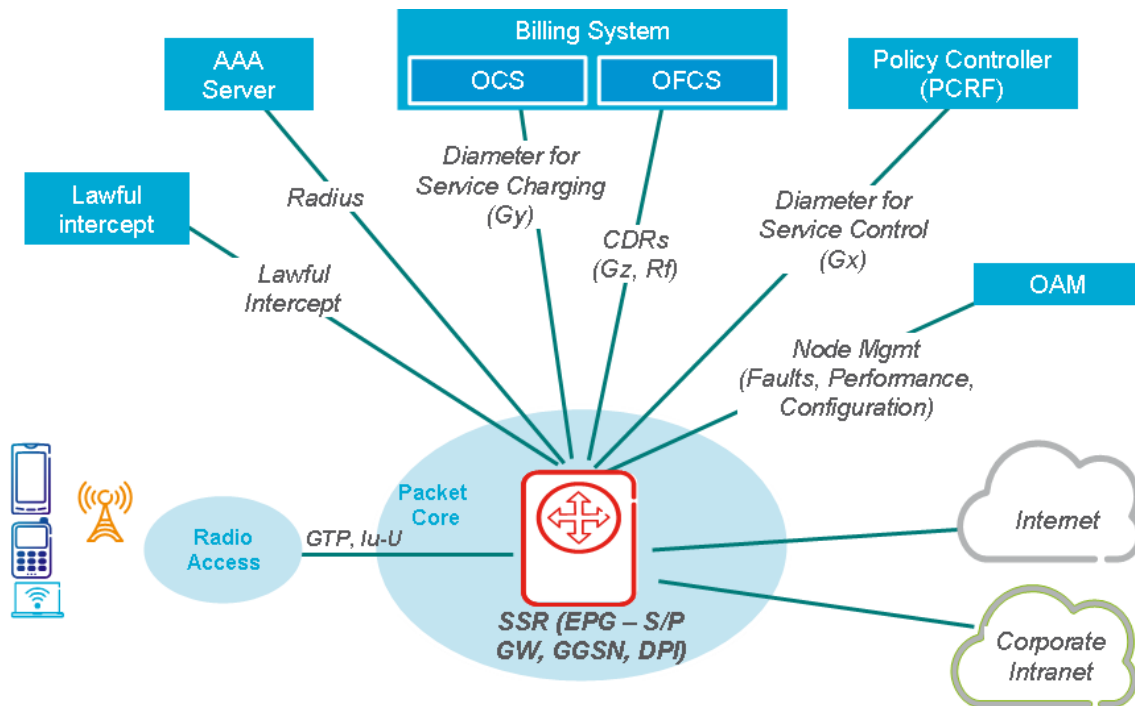


Figure 5 – EPG on Service Router with Its Interfaces

The service router provides Ericsson’s Evolved Packet Gateway (EPG) function, which includes S-GW, P-GW, GGSN and DPI for the EPC and is central to most control plane traffic in the EPC. Ericsson’s EPG provides interfaces to the radio access network (RAN), public Internet and corporate or private intranet. It also exchanges significant volumes of control plane traffic used for lawful intercept, AAA, billing system, policy control and operations, administration and maintenance (OAM) functions.

The service router must be highly scalable in the data plane and control plane. In the next section projections are made for data plane and control plane traffic servicing a packet core node and then used to calculate the TCO of the SSR 8020 versus a competing service router.

Control Plane Traffic Projections

Control plane traffic, however, has two dimensions:

1. IP sessions
2. Signaling transactions

IP sessions have a one-to-one relationship with Packet Data Network (PDN). A single user or device can have one or more IP sessions active at the same time. IP sessions consume memory on the services processor blades of the router. Signaling transactions are measured in transactions per second. They can potentially consume a high level of processor resources.

Figure 6 shows the projections used to model control plane IP session and transaction capacity requirements.




Figure 6 – Traffic & IP Sessions

The number of IP sessions corresponds to the number of active users served by the node. IP growth (CAGR 25 percent) is driven by the migration of wireless subscribers to smart phones and the adoption of M2M solutions. Signaling transactions are subject to a multiplier effect applied to the number of active users and is growing at 131 percent CAGR.

Figure 7 provides an overview of the SSR 8020.




Figure 7 – SSR 8020

The SSR 8000 family of products is a flexible IP service delivery platform for building a next generation of converged IP networks and services with a common management platform. It enables high-end business services based upon MPLS or Carrier Ethernet, the mobile packet core, subscriber management, residential triple-play services and over-the-top services such as DPI, content caching and security.

The SSR 8000 features include slot independency so that signaling-heavy or throughput-heavy environments can be supported without compromising asset utilization. Control plane scalability is enhanced by high-throughput smart services cards (SSCs) that provide services such as Deep Packet Inspection (DPI), Border Gateway Function (BGF), Content Delivery Network (CDN), and Evolved Packet Gateway (EPG with S/P GW, GGSN and DPI). The use of DPI is of particular importance in managing and optimizing bandwidth traffic in a mobile network infrastructure. Clearly, the invocation of DPI can reduce throughput of the network. The SSCs are supported by storage cluster technology for scaling and caching efficiency to ensure high throughput whether or not DPI is invoked.

The SSCs use advanced multithreading architecture with two processors per card. They are designed to offload process-intensive services from the routing processors. Two or more SSCs per service can perform load sharing based on application. The chassis can be populated with as many SSCs as needed.

Competing Service Router

This study compares the SSR 8020 to a competing services router. The competing router has a large global base of installations and is the preferred choice in many large service providers' networks.

The competing services router supports MPLS, Ethernet and legacy interfaces. It enables converged network infrastructure for next-generation service delivery. It is a multiservice edge router designed to deliver high performance and high availability with service-aware operations, administration, management and provisioning.

The router can be deployed as a Broadband Network Gateway (BNG), multiservice edge for Carrier Ethernet and IP VPN business services, and as a mobile gateway for wireless services. In this analysis it is configured with 10 GE ports to connect to the RAN and mobile gateway servers and 100 GE ports to connect to the Internet and corporate intranet.

TCO Analysis Results

The SSR 8020 has 66 percent lower five-year cumulative TCO than the competing router; CapEx is 66 percent and OpEx is 64 percent lower. The primary sources of the cost advantage of the SSR 8020 over the competing router are:

1. Its services cards have more capacity to process IP sessions and support more traffic throughput.
2. Its service cards run S/PGW, GGSN and DPI on the same card and can load share across multiple cards.
3. The SSR 8020 has more router slots than the competitor. This reduces the number of chassis and service cards required by the SSR 8020 to meet the traffic and processing requirements as compared to the competing router.

Figure 8 compares the TCO of the SSR 8020 with the competing service router.




Figure 8 – TCO Comparison of SSR 8020 vs. Competing Router

The SSR 8020 enjoys a cost advantage in all OpEx categories except testing and certification, and training costs (See Figure 9). The cost advantage of the SSR in the network management system (NMS), network care, network upgrades and patches, and capacity management categories are due to the SSR 8020s use of fewer service cards and router chassis. This reduces these costs in two ways. First, total man-hours are decreased because there is less work to do when the number of network elements is reduced. Second, fewer chassis and cards cut down the number of network interfaces and thus the complexity of conducting the operations activities.

Environmental expenses (cooling, power and floor space) are lower for the SSR 8020 because it requires fewer chassis to meet the network control plane traffic load. Power and cooling costs increase with the number of chassis in service as there are overhead costs associated with each chassis. The SSR 8020 has lower power and cooling costs because it uses many fewer chassis than the competing service router.

Service Contract and Engineer, Furnish and Install (EFI) expenses are directly associated with CapEx. With its lower CapEx, the SSR 8020 also has correspondingly lower Service Contract and EF&I expenses.

Figure 9 compares operations expense over five years for the SSR 8020 versus the competing router.

The figures compare increases in scale of IP sessions and signaling transactions versus the corresponding annual increase in the TCO of each router. The figures illustrate a two to over three times scaling advantage for the SSR 8020 compared to the competing router. Note that the slope of the SSR 8020 scaling effect is flatter than that of the competing router. Flat slope indicates stronger scaling effects, especially when compared to the exponential curvature of IP sessions and signaling transactions during the five-year study period.




Figure 9 – Five-Year Operations Expense Comparison

Figure 10 and Figure 11 compare the scaling effects of the two service routers.




Figure 10 – Scaling with Increase in IP Sessions

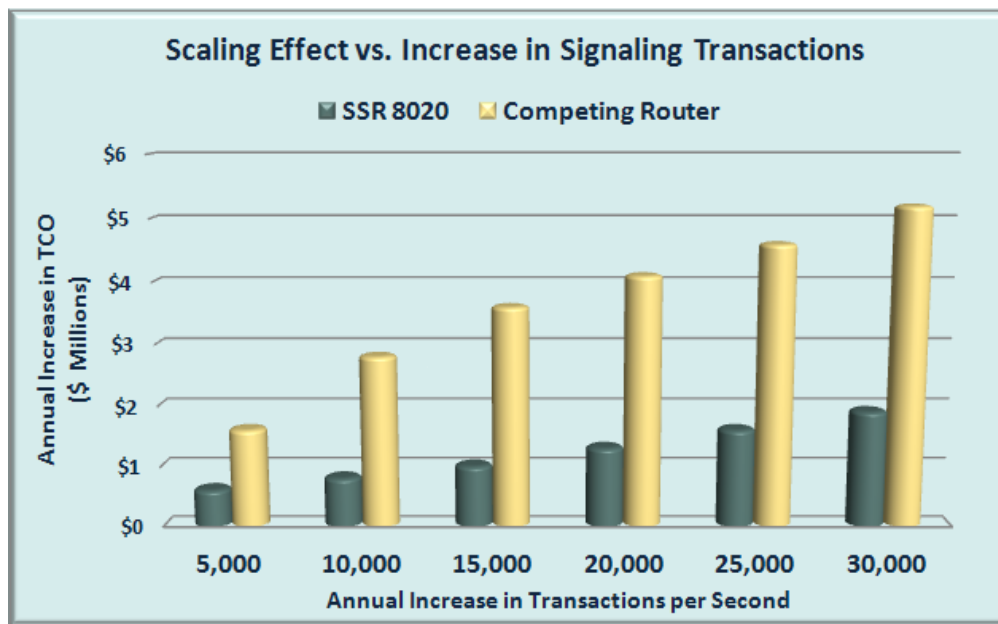


Figure 11 – Scaling with Increase in Signaling Transactions

Conclusion

Data plane traffic growth is being driven by the rapid adoption of video services and cloud services; control plane traffic growth is being driven by migration from fixed and fairly static information sources to personalized, socially-inclusive, and mobile information sources. Rapid growth in the number and types of network uses is determined by:

- Number of end-user connections
- Number of different end-user device types
- Number of apps
- App policy and control requirements
- Mobility management requirements

These drivers of control plane traffic are merging to produce a multiplier effect that makes control plane traffic grow faster than data plane traffic.

ACG Research analyzed the control plane scalability of the SSR 8020 compared to that of a competing services router and determined that the Ericsson Smart Service Router provides the data and control plane scalability needed to meet end-users' service expectations.

The analysis shows that the SSR 8020 has 66 percent lower five-year cumulative TCO than the competing router; CapEx is 66 percent and OpEx is 64 percent lower. The primary sources of the cost advantage of the SSR 8020 over the competing router are: 1) its services cards have more capacity to process IP sessions and support more traffic throughput; 2) its service cards run S/PGW, GGSN and DPI on the same card and can load share across multiple cards; and 3) the SSR 8020 has more router slots than the competitor. Analysis of the scaling effect of the SSR 8020 as compared to the competing service router when subjected to rapid increases in IP session and total bandwidth requirements indicates that the SSR 8020 has a two to over three times scaling advantage than the competing services router.

ACG Research

ACG focuses on providing market analysis and consulting to help service providers, enterprises, and vendors monetize their existing infrastructures and increase operational efficiency and profitability. Through ROI and TCO analysis, product and service message testing, and business model review, reports and forecasts, ACG gives you strategic and tactical advice, services and products, and timely answers so that you can better understand market dynamics and grow your telecom operations more efficiently and profitably. Copyright © 2012 ACG Research.