Designing a Scalable High Capacity Super Softswitch

Amir Atai, Ph. D.

September 22, 2003

***Abstract**: Carriers have already begun deployment of the first generation softswitches for applications such as Voice over IP, Class 5, Packet Tandem, and Internet Offload. Traffic patterns are shifting from traditional access to new technologies. Some carriers are experiencing small negative growth in traditional access lines, while experiencing positive growth in the broadband access, long distance, and wireless. As these carriers enjoy the economic savings offered by the convergence of PSTN and IP networks, their first generation softswitches must carry rapidly changing and expanding amount of traffic originated and terminated from a variety of network interfaces. This change and increase in traffic sources, coupled with the fact that new IP/PSTN signaling protocols require the exchange of numerous signaling messages to establish calls, demand that the second generation of softswitches process a record number of packets and signaling messages with minimal delay. This paper will discuss system architectures that can be utilized to achieve record call volumes while maintaining scalable and flexible network architecture.*

1.0 Introduction

Scalability is an essential requirement for softswitches if they are to be a viable and economical replacement technology for existing circuit switches. Traditional circuit switches were designed to serve well-defined communities of interest. Depending on whether the switch was intended to be deployed in a rural area, a typical residential area in a city, or a highly populated metropolitan area, switch suppliers offered different switch types designed specifically to have a number of lines, trunks, and the call processing capacity to economically serve these communities. After decades of experience, taking into account the statistical nature of traffic generated by residential and business users and standardized service performance requirements, several typical sizes for circuit switches were determined. For the US market switches that economically process about 300 thousand calls per hour and are able to further grow their capacity are the most desirable for typical deployments. However, for rural communities, a much smaller system with only thousands of lines and the ability to process about 50 thousand calls per hour is sufficient. For local tandems or long distance tandems applications, where traffic from many Class 5 switches are concentrated, switches capable of processing one million calls per hour or more are economical. These varying traffic profiles and associated economics produced a segmented switch market place with different circuit switch types specifically designed to carry the anticipated traffic and growth for a particular network deployment.

Having a high capacity scalable softswitch is highly desirable because it avoids the additional cost of deploying a second system when the first switch runs out of capacity. There are considerable costs in addition to the hardware/software cost of the second system. The administrative costs of installing and provisioning the second switch, code splitting, and necessary trunking to interconnect the new switch to the rest of the network are all additional costs that service providers will incur when their switch runs out of capacity. A switch that is truly scalable and can start from a few thousand lines and grow to several hundred thousand lines while processing millions of calls, will provide tremendous flexibility to service providers.

Flexibility in configuring the capacity of a multi-service softswitch is an important requirement. Service providers have started to deploy new technologies and services such as VoIP and SIP-based services via Feature Servers and Media Servers. However, there is not a long history of user behavior and traffic measurements to use for accurately sizing the softswitch for these new applications. An ideal softswitch provides a single platform where the switch resources are optimally shared among various applications according to their bandwidth requirements and the associated QoS. The ability of a service provider to add processing units to the platform quickly and with ease in order to increase the capacity of the system will help remedy the uncertainty in the user behavior and various bandwidth requirements of applications.

At a high level, softswitch architecture can be viewed as unbundling or decomposition of functions performed by existing circuit switches. Call processing function, which is typically performed in dedicated internal processors in circuit switches, now is performed by the Media Gateway Controller (MGC). Line and Trunk modules of the circuit switch where a large number of interfaces can be terminated are replaced by the functions performed by the Media Gateway (MG), and finally the network signaling necessary to establish calls from/to other switches in the PSTN is performed by the Signaling Gateway (SG)[1]. For a more complete explanation of softswitch architecture refer to [4].

[1] This is somewhat of an oversimplification. Softswitches are designed to have additional capabilities such as converting TDM voice into packetized voice and the ability to interface with IP/ATM networks and IP-based application servers.

This decomposition of functions provides great flexibility and allows efficient inter-working between PSTN and IP/ATM networks. However, like any distributed architecture the sub systems in a softswitch must frequently communicate with each other in order to work as a cohesive single system. For example, a typical call in today's PSTN uses about 5 SS7 ISUP messages (average signaling message length of 26 bytes) for call setup and tear down resulting in about 130 total bytes, while a similar call implemented in the distributed softswitch architecture may involve 18 or more signaling messages (average signaling message length of about 250 bytes) resulting in 4500 bytes of total signaling traffic. Therefore, a scalable distributed softswitch must be able to efficiently process numerous signaling and bearer packets while meeting the rigid real-time performance requirements for voice services.

From a network or service viewpoint, MGC, MG, and SG together can be viewed as comparable to a circuit switch (particularly for Class 5 replacement applications or tandem replacement). Thus, the first instinctive thought would be to apply the same capacity and performance requirements developed for circuit switches to softswitches. This notion is certainly valid when the service performance is viewed from the end user's perspective. Certainly, softswitches must provide dial tone as fast as circuit switches, particularly in applications where the eventual service being provided is POTS replacement. Call processing and call setup timing, and reliability requirements previously established in standards and LSSGR for circuit switches would naturally be also applicable to softswitches. One could pose a high level requirement that in general softswitches must provide the same or better performance than circuit switches. After all, that is what is expected of new technology that is designed to eventually replace the existing one.

This paper reviews a number of system architectures that may be used to implement a high capacity scalable softswitch (refer to Section 3). Section 4 summarizes key performance requirements from PSTN that should be applicable to softswitch applications for Class 5 and tandem replacement. In order to be able to make apple-to-apple comparisons regarding switch capacity, a benchmark call mix is defined in Section 5, which exercises most of the softswitch capabilities in a realistic network deployment scenario. Subsequent materials in section 5 define the softswitch capacity in more detail and provide a generic methodology for measuring or estimating the call capacity. Section 6 contains a review of the capacity requirements for current and future applications. Conclusions and key points are summarized in Section 7.

2.0 Scalability: Definition

At a high level, softswitch scalability is defined as the ability to grow or ramp up capacity of the platform in inexpensive and incremental steps, as well as the ability to expand the capacity of the system beyond the present requirements. While there could be other considerations, the main objective is to achieve a direct linear relationship between the amount of resources (processors, memory, interface cards, etc) added and the resulting gains in the performance without degrading availability/reliability and other system requirements.

More specifically, softswitch scalability can be classified into two distinct categories: 1) Vertical Scalability, 2) Horizontal Scalability.

In the vertical dimension, scalability is within a single node. The vertical scalability is defined as the ability of the system to accept more lines or trunks terminations, and call processing CPUs as the demand grows, while maintaining all subsystems as part of the original single node architecture. In an ideal system with infinite Vertical Scalability the processing capacity should increase linearly with the addition of processing units.

The second dimension of softswitch scalability is the ability to distribute MGs, SGs, and MGCs in a way that optimizes a particular network deployment. For example, if a carrier has multiple physical sites, the MGs can be located close to each customer site in order to reduce the transport costs, and then a centralized MGC can be used to process the calls among various locations. In other cases a carrier, may prefer to deploy MGs and MGCs in each location for additional survivability while using a pair of SGs for signaling. The ability of the softswitch to be scalable and flexible in this manner (distribution of MGs, MGCs, and SGs across a geographical area) will be called Horizontal Scalability.

A softswitch that offers both Vertical and Horizontal Scalability will provide the maximum flexibility and value to service providers. A service provider can both grow the capacity of the node as successful deployment and conversion to applications such as Voice over IP occur and at the same time implement the most optimum network deployment architecture that best meets customers' survivability and network operations needs.

3.0 Softswitch System Architectures (Implementation)

When deciding on scalable system architecture to implement a softswitch, careful considerations must be given to availability requirements, type of applications (Class 5, tandem, offload), size of deployment scenario, the types of signaling interfaces supported, and network management and operations. It is also necessary to take into account the type and volume of processing resources that the major components of softswitch architecture require (summarized in Table 1).

Components that perform functions for call processing and signaling require very high availability (1:1 redundancy) because a failure in this equipment will impact a large number of users. In the 1:1 redundancy scheme there is a one-to-one relationship between the working unit and the hot standby. Other components may use a 1:N redundancy scheme where N can be selected by the service provider based on the overall service reliability objectives and economic considerations. In the 1:N scheme there is one stand-by unit for every N active units.

In order to implement a high throughput and scalable system, software components for various functions needed for MGC, SG, and MG must be distributed to the processors and other system resources in an optimum manner.

Table 1 – Softswitch Elements Processing Characteristics

	MGC	*MG*	*SG*	*FS*	*MS*
	Call Processing	I/O & DSP Intensive	SS7 & IP Signaling	Large sub. & Feature Databases	Streaming Traffic
Tasks	Resource management and routing	Real-time Transport	Routing & translations	Mostly IP Traffic & transactions based	High IP Traffic & Real-time
Scalability	CPU Scalability	Port Scalability	SS7 Port & Point code Scalability	CPU & DB Scalability	No. of Sessions scalability
Availability	1:1	1:N	1:1	1:1	1:N

Though it is possible to implement each component as a separate system or implement all of the elements shown in Table-1 in one physical system, there are both architectural and business related reasons for taking a different and more flexible approach. Somewhere between these two extreme alternatives, there is an option of building a flexible system architecture in which components can be either integrated or distributed based on service providers' needs and the particular applications requirements.

Feature and Media Servers provide a platform where value added services and features can be offered by service providers. Third-party companies that specialize in developing software for creating new services such as IP Centrex, IVRs, video streaming, etc are well positioned for building the best in class equipment in that domain. Outsourcing FS and MS from the core softswitch platform enables the platform to dedicate its resources to call processing. However, well-known and standardized PSTN services such as CLASS and other essential PSTN voice services (e.g., 911), where there is close interaction between call control and the features, can be implemented in the MGC platform for superior performance and high throughput due to less signaling communications overhead. This flexibility of providing an open interface to FS/MS and integrating a package of standardized services with MGC both optimizes performance and follows a proven business model.

Distributed parallel computing hardware implementations can be classified either based on the particular resources that are shared among multiple processors or based on the way the processors communicate with each other. Examples of the former are shared memory systems, shared disk systems, or shared nothing systems. These systems can also be classified into "loosely coupled" or "tightly coupled" systems based on the way the nodes/processors communicate with each other.

There are several hardware architectures that can be used:

- Monolithic
- Loosely Coupled
- Symmetric Multi-Processors (SMP)
- Multi-processor Cluster Computing Model

Monolithic models tend to require very complex software, provide limited scalability, and result in higher cost. They will not be discussed in this paper.

3.1 Loosely Coupled Architecture

A Loosely Coupled system is comprised of a set of independent processors in constant communications. This model typically consists of a set of lower level processors called Front End Processors (FEP) performing lower level functions (e.g. Layer 2 signaling) while there are another set of processors, called Back End Processors

(BEP), performing higher level functions. There is usually a dual bus or high speed LAN interconnecting the FEPs with BEPs. Typically, PSTN and IP termination and signaling are implemented on the FEPs while applications and call processing are implemented in the BEPs. The FEPs are responsible for load sharing among BEPs. Databases can be either distributed in the BEPs, FEPs, or alternatively in a separate server. The advantages of the Loosely Coupled architecture are that they provide incremental growth and that the failures are local; if one processor fails, others stay up.

Softswitch call processing is an ideal application for Loosely Coupled architectures because at least half of the computing requirements are signaling protocols, which don't require a common shared data.

The application software architecture must take into account the Loosely Coupled architecture to partition the tasks among the multiple CPUs thereby taking advantage of distributed processing. This is at the heart of having a scalable solution.

3.2 Symmetric Multi-Processors Architecture

A Symmetric Multi-Processor system contains multiple CPUs sharing memory and devices. In the SMP architecture, the same OS runs in all processors. Each processor has its own cache and memory management unit. However, processors must cooperatively share the same memory and devices. Parallelism can be achieved in the SMP architecture by having each CPU run a different process. Memory bandwidth and lock contentions are the key factors for determining the overall throughput. However, most memory usage can be processed by the per CPU cache.

Motivations for using the SMP architecture are scalability and price/performance. In this architecture, adding a CPU to a computer is very cost effective[3]. Only a fraction of the overall cost of the computer is due to the CPU (e.g., 10%-15%). The remaining costs are allocated to motherboard, memory, disk, power supplies, etc. For example, while adding a second CPU in the SMP system can increase the performance by 1.9x, the corresponding increase in the price will only be about 1.1x.

Thus, massive multiprocessor architectures are a huge win if the memory and lock retention issues can be managed for a particular application.

3.3 Multi-Processor Cluster Model Architecture

A Multi-Processor Cluster system is comprised of a group of seemingly independent computers working together in a cooperative manner to provide SMP-like transparency. A cluster consists of a number of peer nodes where one or multiple nodes can run call processing applications or signaling terminations with traffic correctly being forwarded to the right nodes and processes. The challenges in the cluster models are that nodes must agree on the states of resources and communicate with each other (requires a communications protocol). For certain tasks, it requires a processor to act as the central coordinator or master. However, clever implementations can keep this overhead to a minimum (e.g. less than 20%).

The motivations for using this architecture are high availability and capacity scalability. This architecture overcomes SMP memory contention issues. The system survives when a particular node or link fails.

For carrier class implementation it is necessary to consider a multi-processor Cluster model. Multi processors will allow scaling up of call processing capacity. The state of the art of computing technology provides for multiple powerful processors (e.g. 4 processors at about 1GHz) on a single board. However, pure hardware speed will not translate into high call carrying capacity unless multi-process software is designed to take advantage of the multiple processors. It is possible to either purchase an off the shelf multi-processor server or design a platform where the individual processor boards are integrated with the input and output modules into a uniform platform. The advantage of integrating the multi-processor boards into a cohesive platform is that true carrier class redundancy can be built in the entire softswitch platform while providing a uniform operations view for the entire system. This approach will take advantage of the famous "Moore's Law"[2] in the processing power while allowing the supplier to build a very flexible configuration and high availability system.

[2] Moore's law is based on the observation of Gordon Moore (co-fonder of Intel) made in 1965. Moore observed that the density of transistors per square inch on the integrated circuits is doubling every year. He predicted that this trend would continue for the unforeseeable future. However, this trend has slowed down a bit to doubling of the number of transistor to every 18 months. This is the current definition of Moore's law. Most experts believe this trend will continue in the next two decades.

3.4 Plexus System Architecture

A hybrid system can be very advantageous, because it can unite the strengths of several architectures while overcoming the respective limitations. Telica's Plexus softswitch architecture is based on a Loosely Coupled distributed computing architecture utilizing Cluster Computing to achieve scalable call processing.

The compute intensive processes (connection control, call state, signaling protocol) are distributed to the cluster of Computing Modules (CM) while certain common processes (routing, subscriber database, and configuration management) are centralized. By minimizing the number of tasks that require centralized processing and fully distributing the larger remaining tasks to a cluster of Compute Modules, this hybrid architecture can achieve near linear scalability (up to the processing capacity of centralized functions).

Flexible software architecture allows configuring the same platform in various network element configurations. For example, it is possible to configure the platform such that MGC, MG, and SG are all integrated within the same platform. Or it is possible to have MGC and SG in one platform and multiple MGs each in a separate platform. And finally, it is possible to have each element MGC, MG, and SG in a fully distributed architecture where each element is in a separate platform.

This flexibility is achieved by implementing a cluster of Computing Modules where software processes for various functions such as call connection and control, SS7 level 3, TCAP, MEGACO, BICC, SIP are mapped to multiple processors. Various grouping of these processes and associated hardware creates a particular node (e.g. MGC). This grouping is done via provisioning the system during the installation phase.

A multi-threaded[3] approach has been utilized to allow for parallel execution of several processes within a processor. A very high throughput and fault tolerant IP interconnect network (15 Gbs) is provided within the Plexus platform. This allows multiple processors to communicate with each other with minimal delay and without the need for an external router/switch. Thus the arduous task of provisioning multiple IP addresses associated with an external router/switch is eliminated.

[3] A thread is an independent sequence of software instructions that can be executed in a processor.

Plexus distributed computing platform utilizes a cluster of Compute Modules. Each Compute Module contains several microprocessors at 1 GHz or higher. Each microprocessor has a dedicated memory system (1-2 GB) and cache. It is possible to configure a single Plexus chassis with 16 Compute Modules. To ensure 1:1 availability for the critical call processing function, there is one hot standby Compute Module for every active Compute Module. In case an active Compute Module fails, the hot standby Compute Module would resume call processing. Each Compute Module can process about 1 million calls per hour, depending on the call mix.

Figure 1 shows how the Plexus MGC call processing can scale by installing additional processors.

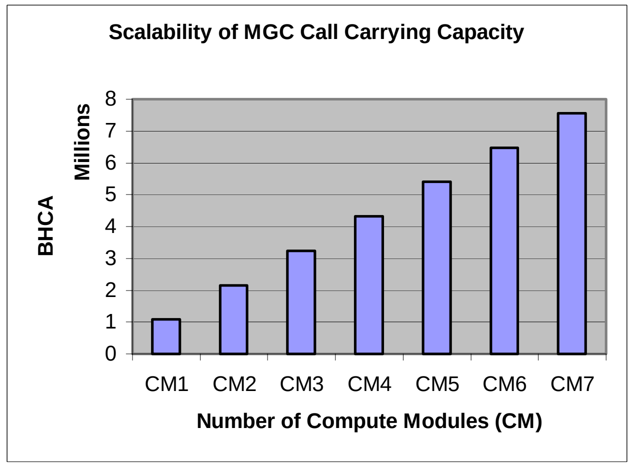

Figure 1 – MGC Scalability Through Installing Multiple Call Processors (redundancy 1:1)

Both TDM and IP ports can be added to the Plexus MG platform by installing additional Input Output Modules (IOM) cards. Each IOM card can contain 8 DS3 terminations resulting in 5,376 DS0 ports per IOM card. A single Plexus MG chassis can be configured with 15 IOM cards to have about 80,000 TDM ports (refer to Figure 2). Similarly, IP ports can be added by installing Voice Processing Server cards in an incremental fashion. Depending on the type of voice codecs used (e.g., G.711, G.726, and G.729) a range of IP-based voice channels can be supported. Access interface cards are typically configured for N:1 redundancy (e.g., N=8). Plexus interface cards have flexible configurations. A TDM port can act as either a line port or a trunk port. Similarly, an IP port can be used either for line side access or trunk side access.

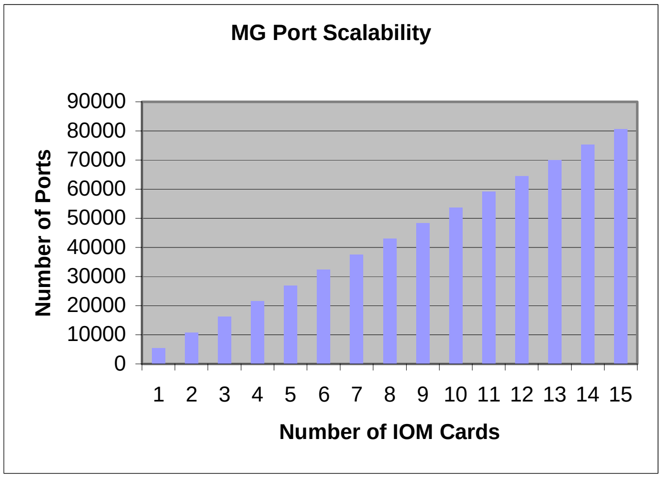

Figure 2 – MG TDM Port Scalability Through Installing Additional IOM Cards

4.0 Softswitch Voice Performance Requirements

There is a direct relationship between scalability and performance. The performance requirements are the constraints or objectives that a scalable system must meet when resources are added to increase the system capacity. As argued in the introduction, a softswitch platform should meet or exceed the performance/capacity of the existing circuit switches. Capacity of the switch is typically measured by the number of calls that it can process during a busy hour while meeting the service performance requirements. There are many service performance requirements developed for voice switches.

Voice service performance standards are described in terms of blocking and delay requirements that apply to the entire switching system regardless of the hardware or software implementation. In this section, we summarize the key requirements[1] that should be applicable to softswitches deployed in a Class 5 or Tandem replacement applications.

- Dial tone delay (measured during ABSBH)
  1. Average Dial Tone Delay < 0.6 s
  2. Probability (Dial Tone Delay > 3 s) < 1.5%
- Probability of Cut-off Calls < 0.000125
- Probability of Ineffective Attempts <0.003
- Cross-switch Call Setup Delay < 400 ms

A valid bid for a call can be blocked inside the switch due to the failure to establish a "path" between a line or trunk terminations when both lines or trunks are idle/available. The blocking requirements (also called matching loss) for the switch are provided for four types of call terminations:

- Line-to-trunk < 1%
- Trunk-to-line < 2%
- Line-to-line < 2%
- Trunk-to-trunk < 0.5 (for tandem)

5.0 Methodology for Measuring the Capacity of a Scalable Softswitch

5.1 Defining Softswitch Capacity

Capacity of a complex system such as a softswitch with multiple processors, a variety of input output interfaces, and signaling protocols is best understood[2] if it is divided[4] into three distinct categories: 1) port capacity, 2) call processing capacity, 3) traffic usage capacity.

MG Network Port capacity is defined as the maximum number of line side terminations plus trunk side terminations that a MG can accommodate. Typically, in central offices individual lines or trunks are aggregated and converted into DS1 (24 DS0) or higher rates (DS3 or STM1) digital format before terminating on the switch line/trunk interface cards. Providing DS3 or higher interfaces results in a very compact switch that occupies very little space. For softswitches that have both TDM and IP line side terminations, the total line port capacity is the sum of the TDM ports and IP ports. Voice over IP calls require additional DSP and processing power in order to perform voice packetization and transport the packetized voice using a real-time protocol such as RTP over IP. Therefore, a softswitch may have different capacity limits for IP type line/trunk termination than TDM type line/trunk terminations.

Thus, the MG Network Port Capacity is defined as:

MG Line Ports= TDM line ports + IP line ports
MG Trunk Ports=TDM trunk ports + IP trunk ports
MG Network Port Capacity = Line ports + Trunk ports

[4] This section follows a methodology similar to LSSGR GR-517 for circuit switching systems. This methodology is generic and does not depend on a particular implementation. However, it is necessary to augment this methodology for new IP/ATM interfaces.

The MGC call processing capacity is defined as the maximum number of originating plus incoming calls that a MGC can process during the busy hour while meeting the dial tone delay requirements.

Traffic Usage Capacity or Active Call Capacity is defined as the amount of traffic generated by customer usage that the switching network (internal to the softswitch, also called switch fabric) could support while meeting the service performance requirements. The customer usage is the sum of the originating and terminating usage. This traffic usage is measured and expressed in units of CCS (Centi Call Seconds or 100 call seconds)

SG capacity is defined as the maximum number of signaling messages per second that the SG can process while meeting the SS7 delay requirements. The SG should be scalable and provide enough capacity to terminate multiple link sets for survivability.

5.2 Benchmark Call Mix

The traffic impact on softswitch capacity depends on the types of the calls and the services associated with the calls. Traffic can enter a softswitch via multiple interfaces and traverse different paths within the switch fabric, consume different amounts of CPU, require the searching of various routing databases, and possibly require feature level processing, etc.

The actual call volume offered to a softswitch will depend on the geographical area, call and service mix, and time of the day. It is useful to construct a benchmark call mix that would represent a realistic traffic mix that a softswitch much handle in a typical deployment in LEC environment. Though it is possible to define several benchmark call mixes to represent various traffic profiles for different communities of interest, in this paper, we suggest a simple call mix that should represent the aggregate average traffic.

The call flow information is represented by the H-chart pattern (Figure 3) where the arrows represent the internal call flows within a softswitch[2].

- Originating Calls 63%
  1. Outgoing Calls 41%
  2. Intraswitch Calls 22%

- Incoming Calls 37%
- Terminating Calls = Incoming Calls + Intraswitch Calls – Line Busy

Note: Originating plus Incoming calls=100%

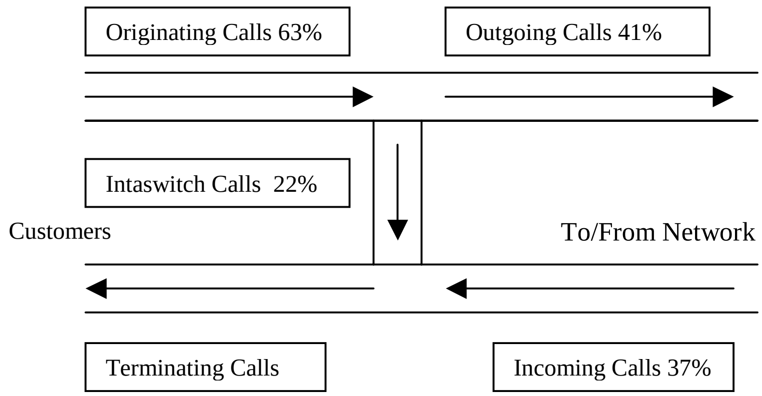

Figure 3 – Call Flow Distribution

It is also necessary to define the distribution of calls for TDM vs. IP. The number of IP calls can be represented as a percentage of total calls. The same call Flow distribution can by used for the total traffic while the percent of IP calls is varied from 0 to the maximum number allowed by IP lines and trunk ports.

Finally, all subscribers' lines should be provisioned with CLASS type features such as Call Waiting, Caller ID, Calling Name Delivery to ensure feature processing is invoked as part of call processing.

5.3 Processor Capacity Estimation

Most MGC architectures are likely to use a set of distributed processors 1 through n for call processing. In a fully distributed architecture, it is possible that each processor may perform certain aspects of call processing. It is also possible to have a hierarchy of processors where certain low level functions are performed by the second tier processors and certain higher level functions such as routing and call control are performed in the first tier processors. In implementations where the calls from 1 through n processors of the second tier feed into a single processor in tier 1, then the sum of the capacity of the processors in tier 2 should be compared with the capacity of the single processor in tier 1. In such distributed implementations, the overall MGC capacity is determined by the processor with the lowest capacity. This processor with the lowest call carrying capacity is named the "limiting processor".

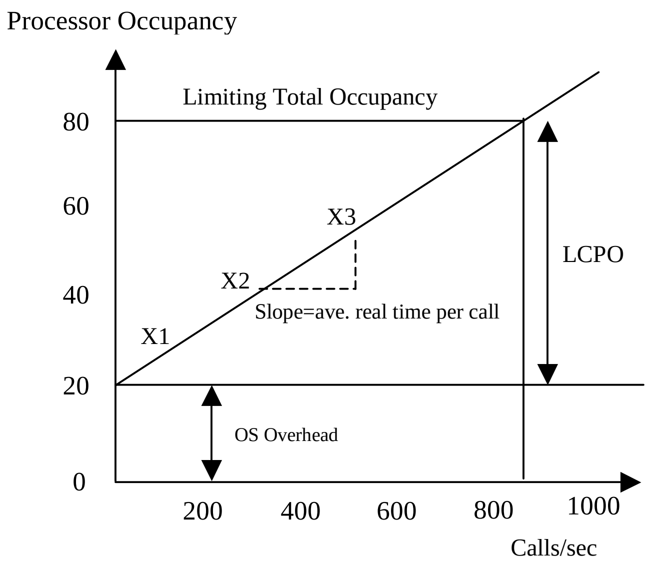

Figure 4 – Estimation of Processor Call Capacity

Call capacity is traditionally measured in thousands of calls per hour. However, processor occupancy is expressed as a dimensionless ratio. The amount of processor time that a call consumes is typically measured in milli seconds. Thus 100% occupancy of the processor means there are 3,600,000 msec available. However only a portion of that occupancy is available for actual call processing. The Operating System and other essential software (e.g. maintenance) will occupy a certain percentage of the processor occupancy. This amount is called the "overhead" (refer to Figure 4). Also it is impractical to run the processor to 100% occupancy. At levels close to 100% delays violate the service performance standards. The upper limit for which a processor can process calls and still meet the service standards is called the "Total Limiting Occupancy". Thus, the occupancy levels between the "overhead" and the "Total Limiting Occupancy" are what is available for call processing. This amount is referred[2] to as "Limiting Call Processing Occupancy" or LCPO. The most reliable method for determining the LCPO is by test measurement.

The second factor that determines the processor call capacity is the average real time per call. Each call type will require a different amount of processing power depending on call types and features. The average real time also depends on the frequency at which various call types are invoked, as seen by the processor. This obviously depends on the call mix. Since each deployment will have its own call mix, the average real time per cal will vary from office to office. In order to make some apple-to-apple comparison, it is necessary to use the benchmark call mix, as was discussed in the previous section. To calculate the average real time per call one has to compute the weighted average of all call types that the switch will be processing. Although this weighted average computation is theoretically possible, it is difficult to perform in practice. The weights (frequency) for each call and the amount of processing power for each call type are not always known in advance.

An alternative method to analytical computation is to measure the average real time per call by running the call mix benchmark at various calls per hour and then measuring the limiting processor occupancy. Once these measurements are obtained, it is then possible to derive the slope of the occupancy vs. calls per hour curve. This slope corresponds closely to the average real time per call (refer to Figure 4).

Once LCPO and average real time per call is determined, then the processor call carrying capacity in calls per hour can be obtained [2] by:

$$\text{Processor Call Capacity}\,(calls/hour) = \frac{3{,}600{,}000\ \ (ms/hour) \times (\frac{LCPO}{100})}{\text{Average Real Time per Call}\ \ (ms/call)}$$

5.4 Line Capacity Estimation

Line Concentration Ratio (LCR) is the ratio of the number of access lines to the number of available line ports on the switch. In circuit switches often a concentration stage is used before the switching stage. In Plexus platform for TDM ports, LCR=1, meaning every access line is directly connected to a line port.

For the GR-303 interfaces, service providers typically provision more subscribers than there are available DS0 channels (between the RDT and COT) to take economic advantage of the statistical nature of the access traffic.

When IP technology is used for the line side access via a protocol such as MGCP, it is possible to provision more subscribers behind the IAD and other residential or access gateways than there is IP voice channel capacity available on the MG and MGC. Since not all users are likely to use their SIP phones at once, some level of concentration ratio may be used. An admission control policy can be implemented in the softswitch to control the incoming traffic when the offered load exceeds the system capacity. This capability requires implementing a resource manager that works with the traffic management software. For example, consider a MG that can support up to x number of IP voice channels. A resource counter can keep track of the total number of IP voice channels being used, and when x+1 call requests service, the service can be denied as part of the admission control scheme.

Therefore, the number of lines that can generate traffic is obtained by:

$$Line\ Capacity = Line\ Ports \times LCR$$

5.5 Active Call Capacity Estimation

Softswitch Active Call Capacity is the maximum carried load (measured in CCS) that can occupy the softswitch resources in one hour while meeting the performance requirements. For a given traffic environment, processor call capacity is independent[5] of call holding time and is relatively constant and varies mostly based on the call mix within a range. However, the softswitch Active Call Capacity decreases as the average call holding time increases (inverse relationship). Therefore, in order to compute the Call Capacity based on Active Call CCS Capacity, the following formula can be used.

$$Call\ Capacity = \frac{Active\ \ Call\ \ CCS\ Capacity \times 100}{Average\ Call\ \ Holding\ Time\ (\sec/\, call)}$$

Though one can use an average call holding time of 3 minutes as an aggregate average. It is important to note that various applications have very different call holding time characteristics. For example, for call center applications the average call holding time may be on the order of 45 seconds while for the internet offload applications the call holding time could be in the order of 20 minutes or longer.

5.6 Call Capacity Estimation

Overall call capacity estimation is the most complex to calculate because it depends on the Line Capacity, Active Call Capacity, and Processor Capacity. Depending on how the softswitch is engineered one of the above categories may become the limiting factor in the overall call capacity of the softswitch.

Call Capacity = min {Processor Call Capacity, Line Capacity × Calls/per Line}

Calls per line statistics mostly depend on whether the line is a residential or business line, and the time period during the day (morning, afternoon, evening). For a residential line a calling rate of about 1-2 and for business lines a calling rate of about 3-5 is typical. The mix of residential to business lines will determine the overall average calling rate for a particular switch. For a switch with a mix of 50% residential and 50% business lines, the weighted average is about 2.5 ABS[6] calls per line or 3 HD[7] calls per line. The average call holding time per call is about 3 minutes.

Example: Consider a Plexus configuration with 500,000 line ports and 7 CM (28 processors) for call processing.

Call Capacity= min {7.56 M calls/hour , 0.5M lines × 2.5 calls/line}

Call Capacity=min {7.56M calls/hour , 1.25M calls/hour}

Call Capacity=1.25 M calls/hour

In the above example, the softswitch had plenty of call processing capacity but not enough ports to use all of the call processing capacity. So the determining factor was the line port capacity. Of course, the number of processors installed can be chosen to match the amount of traffic generated by users.

Or inversely, we can compute the number of ports needed to match the processing capacity by (7.56M / 2.5) or 3.024 Million lines.

6.0 Capacity Requirements for Current and Future Applications

Current capacity requirements for several applications, presently served by circuit switches, are listed in Table 2. A typical range is provided for the number of ports available in typical circuit switches, and the corresponding Busy Hour Call Attempt (BHCA). For internet offload a 20-minute call holding time is assumed.

---

5 Examining tasks purely from the call processors point of view, it takes the same amount of call processing resources to set up and tear down a short call that lasts just 10 seconds versus a long call that lasts 20 minutes. However, a long holding time call occupies the bearer paths within a switch for a longer period than a short call.

6 Average Busy Season (ABS) is defined as the three months with the highest average traffic in the busy hour.

7 High Day (HD) refers to the one day, annually recurring, that has the highest traffic during the busy hour.

Thus, a RAS modem port can support 3 calls per hour on each port, for a 720 port RAS this results in about 2K calls per hour. For other applications a 3-minute call holding time is assumed.

We note that a super softswitch with call carrying capacity of 7 million calls per hour or greater can potentially replace the traffic of not just one Class 5 or Class 4 circuit switch but several (e.g., 3-5). This type of high capacity softswitch node offers a new network rearrangement opportunity for service providers to further consolidate their central offices to achieve operational efficiencies.

In deploying a high capacity node particular attention must be paid to redundancy issues. Both geographical and local equipment redundancy can be utilized to achieve high availability.

Future applications are likely to require more interactions with FS and MS for more customized services. Other future directions may also include inter-working with popular instant messaging software to locate users, forward calls, or in general offer Unified Messaging.

Table 2–Current Applications Capacity Requirements

Applications	*Port Capacity (range)*	*BHCA (range)*
Internet Offload	720 ports per RAS	~2K per RAS
Class 5	10K-140K lines 10K-100K trunks	50K-2M
Packet Tandem	20K-120K trunks	400K-2.4M
MSC Gateway	20K-130K trunks	400K-2.6M

7.0 Conclusions

While some suppliers have decided to use general purpose servers to build a softswitch, a few suppliers have built an optimized integrated softswitch platform utilizing the latest technology in the multi-processor and distributed computing.

Loosely Coupled platforms are expected to have better price/performance than the Tightly Coupled architectures. Furthermore, more research and development is directed towards the Loosely Coupled systems rather than Tightly Coupled as is evident by the progress made in the cluster computing and blade technology vs. N-way processor market.

The advantages of general purpose servers are that they offer a quick time-to-market avenue for suppliers that either don't have the necessary expertise in building the platform or for business reasons prefer to allocate their resources for the development of the application software. Since general purpose servers are built for multiple market segments, they are not fully optimized for a particular application such as softswitches.

Both general purpose servers and the integrated platforms can take advantage of Moore's law. However, integrated platform can offer uniform operations across call processing and input/output modules with built in redundancy in all component levels.

Loosely Coupled system architecture with Cluster Computing can effectively be used to implement a scalable softswitch platform. Once the software is partitioned to scale, additional call capacity can be realized by incrementally installing new call processors. Similarly port capacity can be increased by installing additional interface cards. In order to fully realize the above benefits, the switch supplier must design and implement the application software in such a way that it is partitioned into appropriate modules and executed in parallel in multiple processors.

Flexibility in configuring the softswitch components (MGC, MG, SG) for various nodal configuration (fully distributed, hybrid, or integrated) will allow the service provider to choose the optimum network deployment based on customer needs, economics, and operations issues.

The capacity methodology described in this paper can be used to determine the overall capacity of a softswitch. Three aspects of capacity (call processing, port capacity, and active call capacity) should be determined based on a benchmark call mix and compared against each other in order to derive the overall call throughput. A scalable softswitch should be able to increase its capacity in the Horizontal and Vertical dimensions, and strike a balance between these three aspects.

In summary, a well-integrated Loosely Coupled system with Cluster Computing offers multiple advantages over general purpose servers or tightly coupled architectures. These advantages include higher call capacity, built in redundancy in all levels, smaller footprint, and easier operations and management.

References

[1] "LATA Switching System Generic Requirements (LSSGR), Service Standards", GR-511-CORE, Telcordia Technologies.

[2] "LATA Switching System Generic Requirements (LSSGR), Traffic Capacity and Environment", GR-517-CORE, Telcordia Technologies.

[3] "UCLA Presentation on Distributed Computing", March 2003.

[4] "Reference Architecture", International Softswitch Consortium, June 2002

Acronyms

ABS	Average Busy Season
ATM	Asynchronous Transfer Mode
BEP	Back-End Processor
BICC	Bearer Independent Call Control
CLASS	Custom Local Area Signaling Services
COT	Central Office Terminal
CPU	Central Processing Unit
DSP	Digital Signaling Processing
FEP	Front-End Processor
FS	Feature Server
GR	Generic Requirements
HD	High Day
IAD	Integrated Access Device
IP	Internet Protocol
IVR	Interactive Voice Response
MEGACO	Media Gateway Control
MG	Media Gateway
MGC	Media Gateway Controller
MS	Media Server
MSC	Mobile Switching Center
POTS	Plain Old Telephone Service
PSTN	Public Switched Telephone Network
RDT	Remote Digital Terminal
SIP	Session Initiation Protocol
SG	Signaling Gateway
SMP	Symmetric Multi-Processors
SS7	Signaling System 7
TCAP	Transaction Capabilities Appl. Protocol
TDM	Time Division Multiplex
VoIP	Voice over Internet Protocol
QoS	Quality of Service

Dr. Amir Atai is Telica's Lead Scientist performing research related to the convergence of broadband and the PSTN. Prior to Telica, he was the Executive Director & Chief Scientist of Network Performance & O,A,M Practice at Telcordia where he consulted on network planning, performance, and operations for new and emerging technologies to major carriers/suppliers. Dr. Atai has been at the forefront of developing generic requirements for SS7, AIN, Internet Offload architecture, and voice/data integration. He served as ITESF chairman from 1999-2001 and has published numerous white papers/articles on the impact of Internet traffic on the PSTN. Dr. Atai received both a BSEE and MSEE from the University of Florida and a Ph. D. in EE from Rutgers University.

About Telica, Inc. - Telica provides softswitch-based voice switching solutions for today's evolving telecommunications networks. The company's Plexus™ 9000 is a carrier-grade system that includes a media gateway, media gateway controller and signaling gateway to provide Class 4/5 switching for the PSTN as well as ATM and IP networks -- all in a distributed softswitch model or an integrated open solution. Breakthrough levels of performance, reliability and scalability assure the highest quality of service while easily accommodating subscriber growth. Telica, founded in 1998 and currently deployed in 50 service provider networks, is headquartered in Marlboro, MA. For more information on Telica, visit http://www.telica.com