



SPEAKEASY
NETWORK
ARCHITECTURE

January 2009

Nationwide
800.556.5829

Seattle Office
206.728.9770

Partner Program
888.886.6116

1201 Western Ave
Seattle, WA 98101

www.speakeasy.net

TABLE OF CONTENTS

1.0 Overview	4
2.0 National Network Architecture.....	4
2.1 Introduction	4
2.2 The Basics of High Performance Network Design	4
2.3 Customer Application Benefits	4
2.4 Network Topology Overview.....	5
2.5 National Data Infrastructure.....	6
2.6 Speakeasy Network Routing.....	6
2.6.1 Fiber Path as the Performance Limiter	6
2.6.2 MPLS Overview.....	6
2.6.3 MPLS Implementation Specifics.....	6
2.7 Peering and IP Transport Policy	7
2.7.1 Autonomous System Presentation	7
2.7.2 Network Traffic Classes	7
2.8 Quality of Service	7
2.8.1 IP Prioritization Techniques	8
2.8.2 Speakeasy's Prioritization Techniques	8
3.0 Regional Network Architecture	8
3.1 Introduction	8
3.1.1 Secure and Reliable Physical Facilities	8
3.2 History of Broadband: 1st and 2nd Generation Networks.....	9
3.2.1 Overview	9
3.2.2 Network Designs	9
3.2.3 Regional Networks	9
3.3 Redundant Architecture	10
3.4 Capacity Management.....	11
3.4.1 Introduction.....	11
3.4.2 Circuit Capacity Management	11
3.4.3 IP and ATM Transport Network Protection	12
4.0 Access Technology.....	12
4.1 Introduction	12
4.3 Access Technology Comparison	13
4.3.1 Description of DSL Access.....	13
4.3.2 T1 Service Overview	13
4.3.3 Channelized T1 Circuits	13
4.3.4 T1 Hybrid Model	14
4.3.5 Speakeasy T1 Bonding.....	14
4.3.6 Speakeasy Last Mile Description	15

4.4 Customer Routing and IP Services.....	15
5.0 Speakeasy Network Services	16
5.1 Business Voice	16
5.2 Private WAN	17
6.0 Inside Speakeasy	17
6.1 Network Engineering and Network Support.....	17

TABLE OF FIGURES

Figure 1: Speakeasy Network Topology	5
Figure 2: Speakeasy’s Top Markets.....	10
Figure 3: Logical POP Architecture.....	11
Figure 4: Speakeasy Logical Last Mile	15

1.0 OVERVIEW

Speakeasy is a Seattle-headquartered, national Internet access and broadband solutions provider serving both residential and business customers. It provides service to all major markets in the United States as well as several hundred tier 2 and tier 3 markets. Speakeasy's national data network delivers reliable, secure, and high performance data and application services throughout the continental United States.

This document is an overview of the Speakeasy network, ranging from a discussion of the national data transport network, the regional POP architecture and the last mile transport network. This document helps customers understand the benefits of the Speakeasy network and how it enables them to meet their needs. At a minimum, it presents the information needed to use the Speakeasy network effectively in meeting business needs. If there are additional questions, please contact a Speakeasy representative and they will provide additional information.

2.0 NATIONAL NETWORK ARCHITECTURE

2.1 INTRODUCTION

Speakeasy maintains a national data transport network that connects all Speakeasy POPs and provides transport for customer data and Speakeasy services. The network is a private, layer-2 network that does not run over the Internet. Speakeasy has a sizable data transport network with current capacity well over 30 gigabits. This offers many advantages including a high performance, exceptionally stable network with low jitter and packet loss. These attributes are very useful for accessing Internet-based services and enabling office-to-office applications.

2.2 THE BASICS OF HIGH PERFORMANCE NETWORK DESIGN

A core component for operating a high performance transport network is the elimination of potential bottlenecks throughout the entire network. It requires that the network has certain characteristics and that there is careful attention to any expansion of the network. The network is only as strong as its weakest link.

It is critical that the "intelligence" of the network, the core routing component, is not a constraint. It is especially so when required to interact with the Internet, with network peers, and with a complicated internal network. Core routing must be setup in a resilient fashion and be able to handle network interfaces at line rate with complex filtering and network policy.

Large data transport bandwidth availability for network access is another major requirement for a high-performing network. Proper capacity management is important for this. With poor capacity management, highly utilized links with a high resource demand can lead large portions of the network to suffer from degraded performance. Speakeasy's solution to this problem is to maintain a highly meshed network with a significant amount of capacity.

Finally, in order to optimize network application support:

- »» The entire network must be able to act as a harmonized whole.
- »» The network edge interacts with the core.
- »» Network services operate transparently across the entire network.
- »» Resiliency of all points throughout the network maintains application performance appropriately in an outage or degraded performance event.

2.3 CUSTOMER APPLICATION BENEFITS

Speakeasy's national network offers many benefits to customers and their application traffic. Because the data packets route across a precise physical path and Speakeasy controls the loading and routing of the network, the network is stable, fast, and

secure. Business-to-business data applications, such as VPN traffic, do not have to traverse the Internet and the traffic will not suffer from many of the security risks, performance anomalies, and service interruptions that occur from accessing services located on the public Internet.

2.4 NETWORK TOPOLOGY OVERVIEW

Speakeasy's national network is fully redundant and consists of multiple physical fiber paths that form a mesh of connectivity around the country. Speakeasy maintains a cutting-edge, 2nd generation, high-performance broadband network. It was recently upgraded with world-class network hardware. This enables Speakeasy to offer the same performance and service features throughout the country.

By controlling the network completely, Speakeasy ensures it can offer innovative services when the technology is available for future deployment. It also ensures Speakeasy continues to offer the best service to customers, both now and in the future.

Figure 1 illustrates the network topology. The network consists of a pair of network "rings" that link all the Speakeasy POPs together along with a series of "express" routes that link Speakeasy's major customer markets together. All of this result in improved performance and network load scaling.

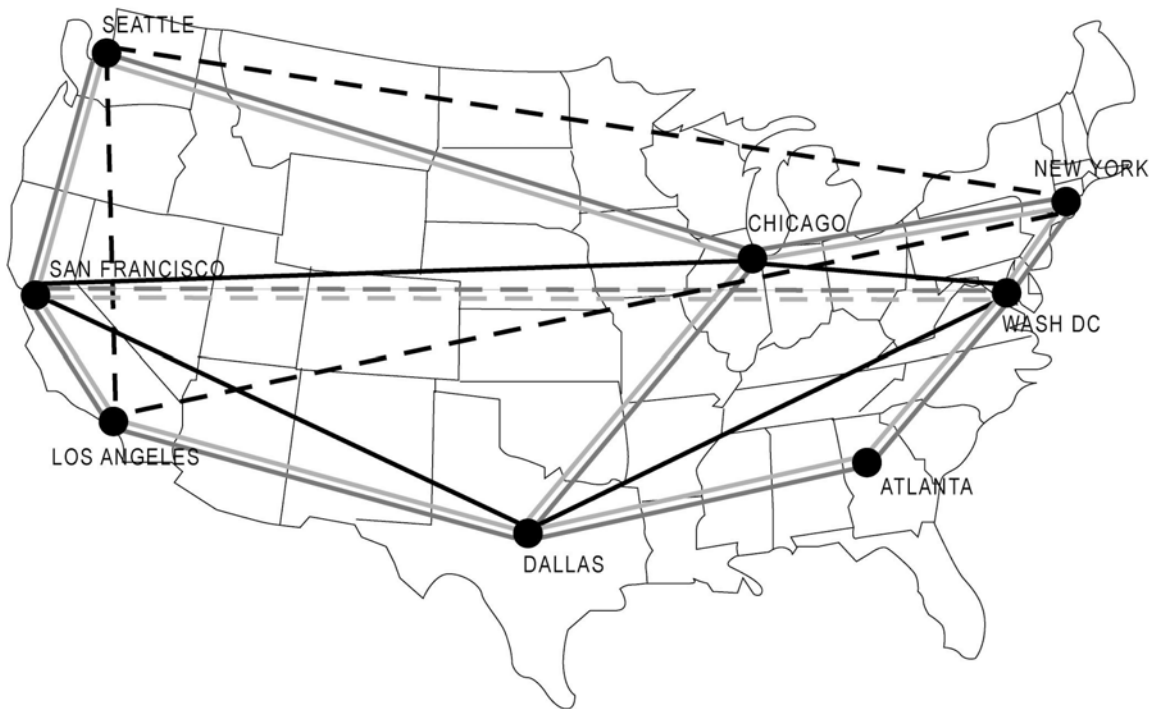


Figure 1: Speakeasy Network Topology

The light gray and dark gray paths show the main primary and redundant "ring network". The black paths show supplemental links for major markets. Speakeasy, also maintains an overlay of "express routes" that cover the redundant ring topology to provide optimal performance between markets.

2.5 NATIONAL DATA INFRASTRUCTURE

The Speakeasy network is a dedicated layer 2 network, which operates across a series of redundant fiber pathways around the country. The network design allows traffic to follow the shortest possible physical fiber pathway across the country. Additional pathways allow for failover with minimal performance impact.

Simply upgrading and running parallel links across the network increases Speakeasy network capacity. Speakeasy deliberately operates the national network at a fraction of its total capacity. Additionally, the network can handle multiple simultaneous circuit failures and still maintain the ability to pass data successfully with performance at acceptable levels.

Since there is heavy management of capacity on the network and the bandwidth availability for the customer traffic greatly exceeds demand, the network is able to demonstrate performance stability even during abnormal network events such as high bandwidth utilization and network attacks. All of this allows business application traffic and Internet access to operate with reliability.

2.6 SPEAKEASY NETWORK ROUTING

As mentioned previously, Speakeasy maintains consistent network infrastructure at all POP locations throughout the country enabling reliable network operations that performs for all application needs.

2.6.1 Fiber Path as the Performance Limiter

The physical length of the fiber network itself is Speakeasy's major determinant of network performance. Due to the design and capacity of the network, Speakeasy suffers from no capacity or router overload latency.

2.6.2 MPLS Overview

Speakeasy utilizes a series of techniques in order to optimize the network for performance. One of the first design components is the Multi-Protocol Label Switched (MPLS) topology, which offers several advantages over a traditional network design. Simply having MPLS does not guarantee improved network performance. MPLS essentially offers an additional toolset beyond standard routing protocols in order to determine an improved network path. MPLS enables Speakeasy to abstract the physical topology from the logical topology that is essential for Speakeasy to add optimized routes and to handle circuit failures in a fashion that delivers a high degree of flexibility and performance.

MPLS enables the creation of additional logic that determines the network pathways used by traffic traversing the network. These are called label switched paths (LSPs). On the Speakeasy network, LSPs are setup to handle a metric not normally covered by a standard routing protocol. That metric is distance. By default, all traffic chooses the path that is the least amount of distance. In the event of a failure, all traffic moves to the path of the physically closest route.

2.6.3 MPLS Implementation Specifics

Speakeasy maintains a partial MPLS mesh, whereupon the network is broken into several components. They are direct network "ring" paths, direct market connections (or "Express" Paths), and various tertiary paths that are used for redundancy in the face of significant network outages. The LSPs support fast reroute that allows the logical topology to update itself very quickly in the event of an outage along the traffic path.

Speakeasy utilizes an interior routing protocol for handling inter-network device reach-ability and runs the border gateway protocol (BGP) to pass internal routes around the network for the customer and service network. It maintains all internal routing decisions in as simple a fashion as possible. Clearly delineated mechanisms for handling different network reachability elements ensures a clear troubleshooting path and fast network convergence times in the event of network device failures. In the event of an outage, this technology handles network failures in a transparent fashion that will not impact network services.

2.7 PEERING AND IP TRANSPORT POLICY

2.7.1 Autonomous System Presentation

Speakeasy maintains a single autonomous system AS and presents it to all neighbors and peers across all of POPs and facilities. All portable customer address space is aggregated into this space and announced equally with the Speakeasy owned space.

2.7.2 Network Traffic Classes

The network design handles three major classes of traffic and optimizes them individually for their specific application. Broadly speaking, these major classes of traffic are—Internet traffic, peering traffic, and customer & Speakeasy content traffic.

Speakeasy follows a network data prioritization scheme that ensures the different classes of traffic are protected against network attacks. The result is service reliability. For more information on this topic, refer to Section 2.8, Quality of Service.

Internet Traffic

In regards to Internet traffic, Speakeasy maintains various upstream peering providers at as many POPs as possible. This gives Speakeasy rich access to provider networks and provider peers, thus enabling Internet traffic to pass to the best possible path on a per POP basis.

Speakeasy immediately passes all Internet-directed traffic directly to the Internet at each POP location. This ensures customer traffic moves directly to the Internet as directly as possible.

Peering Traffic

Speakeasy handles peering traffic differently. Because peering is inherently a direct attachment to the content, there is no benefit in going across the Internet to access this content. In this case, Speakeasy will route the traffic across the Speakeasy national network, thus, ensuring that peering traffic directly attaches and is as reliable as possible.

Customer-to-Customer Traffic

Finally, Speakeasy ensures that all customer-to-customer traffic and Speakeasy services-to-customer traffic is kept on the secure private network. This ensures that all customer office-to-office communication stays off the Internet and that all access to Speakeasy services stays on the network as well. This provides the best possible performance, access, and reliability available for running applications across the country.

To see a real time map of network performance, please visit www.speakeasy.net/network, and check out the regional looking glasses that are available at <http://lg.speakeasy.net>.

2.8 QUALITY OF SERVICE

On the Internet and in other types of networks, Quality of Service (QoS) is the idea that transmission rates, error rates, and other attributes can be improved by measuring them. QoS is of particular importance for transmission of high-bandwidth video and multimedia information. The ability of a network to provide better service to selected traffic over a variety of technologies, including frame relay, ATM, and other IP routed networks, is done by raising or lowering the priority of a given traffic flow over another traffic flow.

QoS identification and marking techniques are ways to prioritize traffic. Because Speakeasy controls the entire network, it can control in a granular fashion the QoS of the traffic that passes over the network when necessary. One example of prioritization of traffic in Speakeasy's network in this manner is Speakeasy's prioritization of Voice over IP (VoIP) traffic over other kinds of traffic, on a customer-by-customer fashion. This traffic prioritization ensures delivery of the critical real time protocol (RTP) stream, which is the active component of the VoIP phone call, or there is degradation of the call quality.

By using network intelligence to support mission critical applications, Speakeasy is able to guarantee quality and top-notch service in a guaranteed fashion that many other providers simply do not have the ability to provide.

2.8.1 IP Prioritization Techniques

Network QoS requires the use of a series of technical tools, especially in a degraded network outage event. The elemental enemy of any critical service is a lack of available bandwidth to support the application. A variety of potential constraints across the WAN and LAN networks can cause a bandwidth issue.

The most common bandwidth constraint point is the customer circuit. This last mile pipe is usually the greatest constraint point. Speakeasy specifically engineers solutions to help with this constraint with offers such as business voice service. It is important to maintain plenty of bandwidth to the office. This will prevent the network from being a constraint.

Some of the other common prioritization methods fall into a generic bucket that can be simply called "queuing techniques". Techniques that follow this method would be things such as maintaining a diff-serv discipline, packet TOS bit marking, and other IP header modification techniques. Essentially, the essence of queuing is reserving available bandwidth for packets that meet the classification requirements for utilizing those resources. By carefully designing the size of the queue resource "bucket", it is possible to ensure that certain types of network traffic are buffered and put "ahead of the queue". In extreme cases, the network drops certain kinds of traffic so that the protected network traffic can be sent.

2.8.2 Speakeasy's Prioritization Techniques

Speakeasy maintains a generic prioritization policy that offers tangible benefits for its customers. In order to protect customer traffic, Speakeasy classifies different types of traffic to have different priorities on the network. It is important to note that Speakeasy maintains a tremendous amount of capacity, and it would require a significant series of outages throughout several parts of the country to trigger use of these routing policies.

Speakeasy's highest prioritization is for Speakeasy's own router-to-router updates. It ensures that the equipment is always able to update itself and, thus, the network is able to protect and communicate with itself in an emergency. This is not a lot of data and there is plenty of bandwidth left for important customer services. The second highest prioritization is voice traffic – this protects against latency and loss of sensitive voice traffic. Speakeasy then prioritizes traffic that is "on-net"; this ensures that customer-to-customer, and customer-to-Speakeasy services are highly available.

Speakeasy's QoS implementation is highly flexible and capable of additional modification and customization for new services.

3.0 REGIONAL NETWORK ARCHITECTURE

3.1 INTRODUCTION

Speakeasy maintains regional points-of-presence (POPs) and network access points (NAPs) throughout the country. These regional POPs are completely stand-alone components designed to offer optimal access and regional redundancy to ensure that any problem on a small part of the overall network will not affect the rest of the network.

Speakeasy is operating a second generation broadband access network and a modern, redundant architecture that is consistent throughout the country. By replacing all legacy equipment, Speakeasy is running a consistent and modular network architecture duplicated at all network facilities. Many other providers operate multiple networks with multiple technologies. This results in an inconsistent customer experience with different performance and service attributes.

3.1.1 Secure and Reliable Physical Facilities

All network equipment is located in world-class data centers and all equipment and circuits are locally and remotely monitored 24 hours/7 days. Network equipment in each POP is setup in a redundant fashion, with local equipment sparing available at all facilities to reduce MTTR ("Mean Time to Repair") in the event of hardware failures. Each Speakeasy network facility has controlled access and a high degree of physical security. Maintained redundant power to each of the sets of active Speakeasy equipment types along with generator and battery backup systems ensure operation through extended outages.

3.2 HISTORY OF BROADBAND: 1ST AND 2ND GENERATION NETWORKS

3.2.1 Overview

Broadband networks have evolved quite significantly in the last decade. Initially, general-purpose routers with advanced code upgrades enabled them to support a reasonable amount of broadband subscribers for basic connectivity. As networks evolved, the requirements on the equipment to scale and perform increased tremendously. A significant amount of evolution and progress in Broadband networks and equipment has resulted.

From a hardware architecture perspective, 2nd generation broadband network equipment has some elemental characteristics. They include the usage of extremely specialized and dedicated hardware to replace interrupt driven, centralized processor/memory software driven architectures. In order to improve performance and scaling, the central route-processing model is supplemented with the distributed routing on the line cards, significantly enlarged interface modules, and dedicated hardware application specific integrated circuits (ASICs). This handles many of the functions once barely supported on previous generation hardware.

3.2.2 Network Designs

Having scalable broadband subscriber management hardware is only part of the battle. Once there is a solution for the performance and scaling issues of the network hardware, the performance of the physical network once again becomes the limiting factor to optimizing performance.

As operating physical facilities and maintaining a complex network is expensive, it was a common design component of early broadband network designs to minimize POP locations and backhaul traffic across the country in order to lower the operational costs per user. This significantly reduces network performance to certain destinations, especially local markets, and generally increases the significance of equipment failures on the network as a whole.

Improved performance results by installing regional facilities around the country. This allows customers to get on the Internet as directly as possible and allows a high degree of resiliency. For Speakeasy, potential failures of different regions do not affect other regional networks because each one is completely autonomous.

3.2.3 Regional Networks

Speakeasy's network architecture has eight POPs as shown in Figure 2. Each POP is a complete, standalone facility with multiple, redundant points of access to the Internet, the Speakeasy network, and to local Speakeasy applications and services.

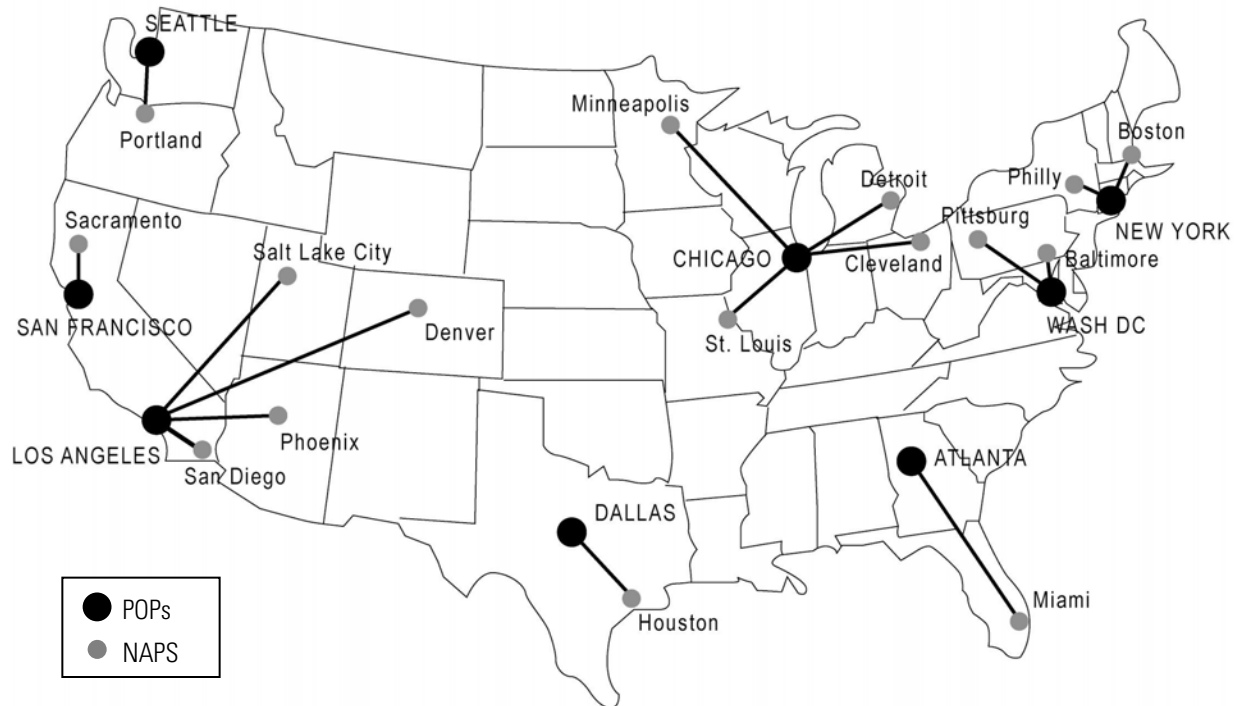


Figure 2: Speakeasy's Top Markets

In order to maintain optimal performance for customers, Speakeasy also uses secondary network access points (NAPs) to extend the customer footprint while still maintaining optimal regional performance. NAPs are point-to-point aggregation circuits that run from the Speakeasy POP equipment to the regional facility that the NAP is servicing. Speakeasy also uses remote locations including: Albuquerque, Austin, Indianapolis, Dayton, Columbus, Milwaukee, Raleigh, Richmond, Birmingham, Nashville, New Orleans, Kansas City, Grand Rapids, Tucson, Charlotte, Hartford, Norfolk, Louisville, Memphis, Santa Barbara, Jacksonville, Tampa, San Antonio, Greensboro, Las Vegas, Orlando. Figure 2 illustrates the Speakeasy POPs and NAPs.

3.3 REDUNDANT ARCHITECTURE

Three primary logical components comprise each Speakeasy POP—the core routing logic, the local services, and the customer management system.

The Internet access core is a redundant routing function that handles the logical routing across the Speakeasy network for customer-to-customer traffic and the Internet for access to hosts that are not on the Speakeasy network.

Another important part of the POP is the local services component. This consists of local content services and critical infrastructure such as authentication and domain name server (DNS) services that enable the POP to continue functioning autonomously despite a potential failure of other components in the network.

The last component is the customer management system. The functional network element manages the customer routing, circuit aggregation, and individual customer network services. Figure 3 illustrates the logical POP architecture.

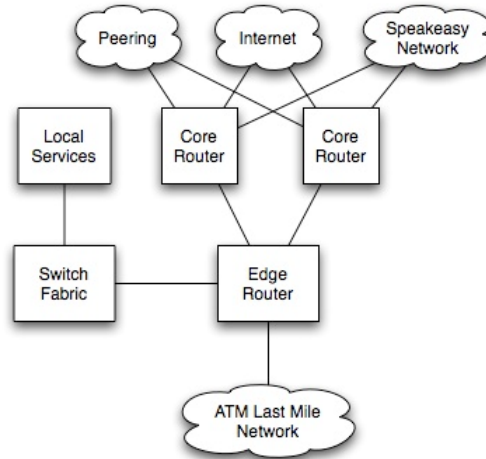


Figure 3: Logical POP Architecture

3.4 CAPACITY MANAGEMENT

3.4.1 Introduction

Speakeasy can offer premium services because no single one link is at capacity and there is proper maintenance of all components of the POP. Speakeasy maintains redundant cross-connects, as well as very stringent capacity management guidelines. This provides all the bandwidth and equipment horsepower required to ensure the customer receives optimal performance. Due to the standardization on very high capacity circuits, and the Speakeasy second-generation broadband platform, Speakeasy is able to ensure high network performance, guaranteed capacity, and advanced services managed on a circuit-by-circuit basis. It is common for companies to “overload” circuits in order to reduce operational costs. Speakeasy works hard to ensure that this never happens and that all the capacity needed is available.

3.4.2 Circuit Capacity Management

It is common when using dialup services and commodity broadband services to focus on ratios like oversubscription and other metrics that describe network capacity. Speakeasy looks at network capacity from a different angle, one that more adequately defines the network capacity needs for different regions and customer types.

Speakeasy does not have an “over-subscription ratio”. Speakeasy currently maintains all circuits below capacity and has a process that triggers an augment when additional capacity is required. By carefully forecasting the growth in different regions, it is able to add circuits ahead of demand in order to guarantee that appropriate capacity is always available to the customer.

Another Speakeasy differentiator is that it manages capacity by peak utilization instead of average utilization. Additionally, if a circuit reaches capacity even for a short period each day and normally remains below capacity, it will trigger the necessary augment. Diminished performance for even a short period each day is a problem for applications.

This focus on ensuring appropriate capacity, allows Speakeasy to offer advanced services, such as voice over IP (VoIP), in a reliable fashion.

3.4.3 IP and ATM Transport Network Protection

In addition to Speakeasy’s robust capacity management program, Speakeasy maintains additional protection of business class customers on the network. Speakeasy does not operate the customer access network at capacity. In the event high traffic flows hit the network and make it past the layered defenses, several mechanisms ensure protection of prioritized traffic.

There are agreements with Speakeasy vendors that business circuits have a higher priority than normal residential customers do. This ATM traffic shaping profile enables the customer to receive prioritized bandwidth in the event of a major ATM transport network outage. Combined with network QoS techniques, this helps ensure the delivery of business traffic.

4.0 ACCESS TECHNOLOGY

4.1 INTRODUCTION

In order to offer world-class connectivity, and to be able to understand why Speakeasy service has real advantage to the customer, it is required that a thorough understanding of the market place and of the techniques used in offering these services.

Speakeasy takes advantage of multiple last mile technologies and vendors in order to offer exceptional service throughout the country and each one of Speakeasy’s partners are required to deliver exceptional network quality. By leveraging multiple last-mile provider partners, it allows Speakeasy to be selective about the providers and be able to offer the best price and best quality vendors to the customer.

Speakeasy’s T1s deliver the guaranteed performance of a point-to-point T1 with the resiliency of a redundant network. This offering is far superior to point-to-point T1s and traditional frame relay T1s. The following section describes the differences and benefits.

Table 1: Product Offering Comparison

Type	Speed	Technology	SLA
Bonded T1s	3.0 M ↑ , 3.0M ↓	Multilink PPP over ATM PVC's, with Multilink PPP over Frame Relay Last Mile Handoff over an engineered circuit	99.99% uptime 4 hrs MTTR 100% bandwidth
T1s	384 k ↑ , 384 k ↓ 768 k ↑ , 768 k ↓ 1.5 M ↑ , 1.5 M ↓	ATM PVC Transport with Frame Relay Last Mile Handoff over an engineered circuit	99.99% uptime 4 hrs MTTR 100% bandwidth
SDSL	384 k ↑ , 384 k ↓ 767 k ↑ , 767 k ↓ 1.1 M ↑ , 1.1 M ↓ 1.5 M ↑ , 1.5 M ↓	2B1Q Based Line Encoding over POTS copper wiring.	99.99% uptime 24 hrs MTTR
ADSL	256 k ↑ , 1.5 M ↓ 768 k ↑ , 1.5 M ↓ 768 k ↑ , 6.0 M ↓	G. Lite, CAP Based Line Encoding, Frequency Division over POTS copper wiring.	36 hrs MTTR
IDSL	144 k ↑ , 144 k ↓	Dual ISDN B channels-64Kbps +D control channel-16Kbps	36 hrs MTTR

Business service is not just about bandwidth and speed. It is about quality and assurances that there is quick resolution if something goes wrong. The T1 services have a higher uptime guarantee and faster mean time to repair (MTTR) on the product. If

business data is mission critical, purchasing a service with these additional guarantees will ensure greater reliability and improved restoration times.

For more information about Speakeasy service offers, please go to Speakeasy's online qualification page. Refer to Speakeasy's Terms of Service for further information about the SLA and service components of all the service offerings.

4.3 ACCESS TECHNOLOGY COMPARISON

4.3.1 Description of DSL Access

Digital subscriber line (DSL) service brings high bandwidth to the customer over copper telephone lines. It is a rate adaptive technology and cost competitive but all forms of xDSL are distance sensitive. Symmetric digital subscriber line (SDSL) service has the same data rate in both upstream and downstream directions. Asymmetric digital subscriber line (ADSL) is symmetric because most of its bandwidth is devoted to the downstream direction. Yet another flavor of DSL is ISDN DSL. IDSL, which is a variety of traditional ISDN, has the B channel and the control channel available that gives it a slight bandwidth boost from traditional dual channel ISDN service.

4.3.2 T1 Service Overview

One of Speakeasy's primary offerings is high quality T1 service. As there are many methods of offering business-grade T1 service, the following paragraphs describe the main mechanisms involved and the differences between them.

Traditional Frame Relay Overview

Traditional frame relay service offered dedicated data access to the office for many decades. It was developed and used before there was a public shared data network known as the Internet. Legacy private network access commonly used frame relay as compared to general Internet access and server hosting. Almost all large carriers that historically operated frame relay service have migrated to ATM-based networks due to difficulties in managing a large frame relay network.

Frame relay is a fast-packet technology that means that the protocol does not attempt to correct errors. When there is error detection in a frame, it is thrown away. The end-points are responsible for detecting and retransmitting any dropped frames. This is very similar to ATM and other forms of layer 2 protocols.

Traditional frame relay runs on a single permanent virtual circuit (PVC). While this technology is reliable, it has several challenges that cause other forms of data service to be more popular in the modern business environment. The primary reason that frame relay is a declining solution is due to the limited bandwidth availability and the very high cost of "guaranteeing" data rates on frame relay. Frame relay generally has a committed data rate (CDR) and a burst data rate. While the circuit is capable of "bursting" to a high speed, it is often not able to sustain this, especially during peak hours. These issues with frame relay often result in the product not being useful for Internet savvy businesses. Another important issue with frame relay is simply the cost. Modern data communication methods generally have a significant cost/performance advantage over frame relay with a greater flexibility in equipment choices for terminating the connectivity.

4.3.3 Channelized T1 Circuits

Another method of delivering T1s is to offer it over point-to-point DS3s. This technology has some advantages over frame relay since it offers, "dedicated bandwidth" to the customer over the last mile. Since the customer receives all the bandwidth "dedicated" on the circuit, it may seem like this is best of all possible worlds. It is critical that the upstream network equipment, which terminates the channelized T1, is managed appropriately to ensure that bandwidth continues to be available. Although the dedicated bandwidth pitch seems attractive, the potential bottleneck simply shifts from the T1 circuit to the network aggregation equipment. The aggregation piece rapidly becomes part of a shared infrastructure. Appropriate management by the service provider is necessary in order for the customer to receive the required bandwidth.

Another challenge a service provider faces with the utilization of point-to-point circuits is the difficulty in making the channelized point-to-point circuits redundant. There can be several failure points that can be difficult to mitigate in the event the circuit fails. If any piece of the transport infrastructure fails, the dedicated circuit will fail. Additionally, there are high initial operational costs

required to provide the service because of the need to dedicate equipment to the customer and the limited ability to share and efficiently utilize the equipment at higher capacities.

4.3.4 T1 Hybrid Model

Speakeasy offers a different option for customers. For T1 services, Speakeasy runs a hybrid model utilizing frame relay for the last mile to the central office (CO), and then shifts to an ATM transport model for the extended connectivity back to the Speakeasy network.

By enabling individual customer's traffic prioritization schemes across the ATM network and ensuring all transport circuits do not operate at capacity, Speakeasy is able to offer the best of both worlds. It can offer the dedicated bandwidth of a point-to-point solution with advance failover features which will allow the circuit to continue to be available in the event of a failure in the transport network. Additionally, ATM networks are quite scalable and are well understood and standardized throughout the industry.

Speakeasy's network technology strategies help create the cornerstone of quality, high availability, and guaranteed performance that differentiate Speakeasy from many other service providers. This allows Speakeasy to offer high quality products cost efficiently to customers.

4.3.5 Speakeasy T1 Bonding

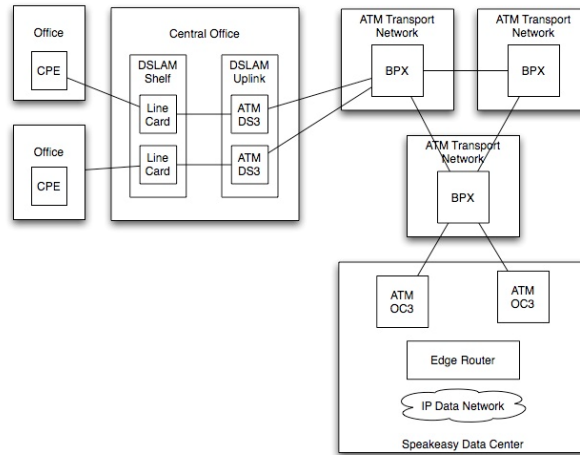
In the event that greater bandwidth is required for T1 services, Speakeasy offers a bonded T1 service. This service has the same characteristics and SLAN as the normal T1s. This service is essentially two T1s bonded together utilizing a multilink point-to-point protocol (PPP) bundle. There are some solid advantages to this product, namely that it offers significant additional bandwidth and greater failover resiliency than a standard T1 product.

Unlike load balancing or other similar solutions, the bonded T1 supports up to a full 3-megabit flow across the circuit. This allows an application with high bandwidth needs to utilize both circuits simultaneously.

The T1 bonding solution allows a true 3-megabit pipe on both the download and the upload. It appears to be a single interface from the perspective of the network routing. Bonded T1s will continue to operate in the event of a circuit failure and it will take only a few minutes for the MLPPP bundle to negotiate. Even when a circuit fails in the bundle, it will continue to pass traffic over the remaining circuit. This circuit failure protection is not instantaneous, and in the event of an outage, the Edge customer premise equipment (CPE) and the Speakeasy edge router will need to handshake and renegotiate the bundle. However, it will give greater protection and a faster recovery than if dependent on a standalone T1.

From a protocol perspective, the edge CPE will synchronize with the Speakeasy edge router, and they create a logical circuit bundle of the T1 circuits. The PPP encapsulated data will then operate over the multiple circuits bundle and will automatically split the loading across both T1s. PPP encapsulated packets, once received by the edge CPE or the Speakeasy edge router, are assembled into a single data stream and de-capsulated into IP datagrams ready for transmittal across the LAN and the Internet.

Figure 4 shows a diagram of the Speakeasy logical last mile.



Grossly Simplified and Theoretical Last Mile Topology

Figure 4: Speakeasy Logical Last Mile

4.3.6 Speakeasy Last Mile Description

In order to help explain Figure 4, the following is an overview of the last mile network elements.

The CPE is located in the customer office and connects the office network to the last mile network.

The Digital Subscriber Line Access Multiplexer (DSLAM) is the device located in the local central office or remote terminal that converts the CPE physical layer signaling and transports the network data encapsulated within the ATM cell or vice versa.

The ATM transport network is a collection of large ATM circuits, generally DS3 or above, that transports many DSL circuits. The ATM transport network normally comprises of multiple ATM switches that handle the logical switching. It protects the network transport circuits within the ATM network in the event of a failure. This network is also capable of prioritizing and queuing different kinds of traffic on a PVC-by-PVC basis. In the event of a circuit failure, the ATM switches reroute the customer's PVC down another DS3 that is available.

The ATM network handoff is generally an OC3 or DS3 that hands the subscriber traffic from the ATM network to the edge router. The customer edge router handles tens of thousands of customer connections over dozens of circuits. It is redundant equipment that is designed to have very high availability.

The edge router performs a multitude of functions. The most important of which is reassembling the IP data from the ATM PVC and then handing it off to the core network for transport to other customers on the router, the Internet, the Speakeasy network, or local services.

The edge router then talks to the Core network, which maintains the "intelligence" to forward and direct traffic to the best path and provider.

4.4 CUSTOMER ROUTING AND IP SERVICES

It is important for customers to understand the Speakeasy IP address policy when getting a Speakeasy circuit. Speakeasy will also help with IP address assignment and routing.

Speakeasy works with the American Registry for Internet Numbers (ARIN). This group allocates Internet numbers and maintains reverse DNS names. Speakeasy handles all the negotiations, address record management, and handles the fees and justifications required to get address space with this group. If the customer's organization maintains a portable address block assigned from ARIN, Speakeasy is happy to announce this address space.

As IP addresses are a limited resource, and have a tangible cost associated with their allocation, Speakeasy has a formal policy around the assignment of address space. For a bridged business circuit, customers provisioned using a maximum of 10 IPs. To have more than 10 IPs, the customer must have a routed connection. Up to 32-routed IPs are available to a business customer free of charge. After that, routed IPs are available in blocks of 8, 16, 32, 64, etc. for a charge. Customers must use all the IPs assigned to them.

If a customer needs more than 32 static IP addresses, it is required that they work with the Speakeasy account team to complete an address justification form used to grant additional address space. This information is required for records to help justify the address allocation with ARIN.

Speakeasy aggregates all the IP addresses regionally and announce these addresses globally to the Internet at large, including portable addresses that a customer moves to Speakeasy.

5.0 SPEAKEASY NETWORK SERVICES

5.1 BUSINESS VOICE

Speakeasy's business voice service provides an integrated voice and data solution that delivers voice and technology over a broadband connection with a company-wide web portal. VoIP is a method by which voice is digitized and transmitted in digital packets rather than using traditional circuit-committed protocols of the public switched telephone network (PSTN).

Speakeasy's Voice Quality (VQ) technology prioritizes voice packets over data. Because Speakeasy has a private, dedicated IP network, it provides dependability and performance. This allows Speakeasy to maintain voice QoS guarantee from the desktop to the PSTN in a manner that is not possible for most providers.

Speakeasy's VQ technology solves the problems of receiving reliable voice service over the Internet. By prioritizing voice traffic at all the constraint points of data networks, Speakeasy is able to provide a high quality and reliable service. Speakeasy eliminates the traditional bottlenecks of the customer WAN outbound/inbound data traffic, network router prioritization, and transport across the Internet.

Speakeasy addresses these issues with a variety of techniques that are appropriate for each problem. The customer WAN bottleneck is protected with the use of a specialized router that is capable of ensuring that voice traffic receives priority leaving the LAN router. This is the largest and most common bottleneck. Speakeasy then tracks all the voice traffic, and using a combination of queuing and data packet prioritization ensures that the voice traffic is giving priority on both the edge router and the core network. Finally, the Speakeasy Network directly connects to all the voice network equipment so that there is no voice traffic passed over the Internet, nor any constrained network points. This final piece ensures a high quality voice service.

It is a critical to realize the benefits of voice traffic carried over the Speakeasy broadband network and passed off to directly peered partner networks without touching the public Internet. Keeping voice traffic secure is critically important. Running all voice traffic over a private network ensures voice calls are safe from individuals listening as calls pass over the Internet.

The Speakeasy website has further information on Speakeasy's Business Voice Service.

5.2 PRIVATE WAN

Speakeasy's Private WAN (PWAN) product is available throughout the country and offers a solution for private network access needs. Speakeasy Private WAN is more secure and simpler than a VPN and has none of the traditional hardware and maintenance costs typically associated with the product.

The topology of this solution generally takes a main branch office and attaches multiple secondary offices to the main branch. By routing all the data traffic over a private network, the traditional risks associated with putting traffic on the public Internet are not there.

Speakeasy sets up the customer edge router to act as a "network bridge" which becomes a virtual Ethernet switch for multiple DSL circuits located anywhere around the country. By registering the MAC addresses of devices at each office location and only forwarding traffic between the locations that have these MAC addresses, this solution conserves bandwidth across the network. This solution also does not pass non-directed broadcast traffic that further improves scaling. This creates a very powerful and simple solution—a national WAN network with the complexity of an Ethernet switch.

6.0 INSIDE SPEAKEASY

6.1 NETWORK ENGINEERING AND NETWORK SUPPORT

One of the most important components of a quality service is the support groups that service the connectivity and are there for support if there is a problem with the service. Speakeasy provides support through dedicated business manager (DBM) service that comes with all business class T1 circuits. The DBM model assigns an account manager to every customer. That account manager serves as a representative for all customer issues, technical or otherwise.

Backing up the DBM group is the Network Operations Group that supervises, monitors, and maintains the Speakeasy network. All business customers are equipped with a Proactive Network Monitoring (PNM) service that monitors their broadband connection 24 hours a day. Other network monitoring tools continuously monitor backhaul circuits, internal and customer services on servers, and all network equipment. With PNM, in case of an outage, a service ticket is created automatically and assigned to a Speakeasy technician who immediately begins troubleshooting and working to resolve the problem. These teams of people, working in tandem, are the reason Speakeasy confidently offers its products SLAs.

Speakeasy's Network Operations Group is a cadre of dedicated professionals who work around the clock to make sure the complex interactions of network elements and services are constantly available. The group is located in Seattle, Washington with the rest of the corporate headquarters and has very close ties to the DBM and support groups as well as the Network Engineering team.

The Speakeasy Network Engineering Team is responsible for the design and turn up of the network. They maintain all the code versions and perform the stability testing and validation of all the network features. They are also available around the clock to support network operations in the event a major outage taxes the knowledge and procedures of the group.

Author: Danny Pickford, VP Technology, Speakeasy, Inc.