White Paper

High-Availability Solutions for SIP Enabled Voice-over-IP Networks

The rapid evolution of voice and data technology is significantly changing the business environment. The introduction of services such as instant messaging, integrated voice and email, and follow me services has contributed to a work environment where employees can communicate much more efficiently, thus increasing productivity. To meet the demands of the changing business environment businesses are beginning to deploy converged voice-and-data networks based on Session Initiation Protocol (SIP).

SIP was originally defined in 1999, by the Internet Engineering Task Force (IETF) in RFC2543. The definition was the culmination of years of work in the IETF's MMUSIC Working Group to provide a mechanism to allow voice, video, and data to be integrated over the same network. SIP provides the foundation for building converged networks that support seamless integration with traditional voice networks, email, the World Wide Web, and next-generation technologies such as instant messaging and 3rd Generation Partnership Project (3GPP) mobile networks.

As businesses continue to increase their use and reliance on converged services, reliability and availability becomes increasingly important. This paper will introduce the key elements of a SIP network, further define the concept of high availability in SIP networks, and explore various techniques to increase the availability of SIP-based Voice over IP (VoIP) networks.

Introduction to SIP Networks

A SIP-based network is made up of several components:

- SIP User Agent—any network end-point that can originate or terminate a SIP session. This might include a SIP-enabled telephone, a SIP PC-client (known as a "softphone"), or a SIP-enabled gateway.
- SIP Proxy Server—a call-control device that provides many services such as routing of SIP messages between SIP user agents
- SIP Redirect Server—a call-control device that provides routing information to user agents when requested, giving the user agent an alternate uniform resource identifier (URI) or destination user agent server (UAS).
- SIP Registrar Server—a device that stores the logical location of user agents within that domain or sub-domain. A SIP registrar server stores the location of user agents and dynamically updates its data via REGISTER messages.

High Availability

High-availability solutions for VoIP networks address the need for users to be able to place and receive calls under peak-load call rates or during device maintenance or failure. In addition to lost productivity, voice-network downtime often results in lost revenue, customer dissatisfaction, and even a weakened market position. Various situations can take devices off line, ranging from planned downtime for maintenance to catastrophic failure.

High availability is not a specific technology, but a goal that is based on specific business needs. The complexity of a high-availability solution is determined by a company's availability needs and by the amount of system interruptions that can be tolerated by a business. High-availability solutions improve your VoIP network's availability, lower the costs associated with downtime by preventing outages, and reduce the impact of outages when they do occur. There are two key elements that contribute to availability in a VoIP network: capacity and redundancy. These concepts will now be explored further.

Capacity

Capacity is a measurement of the volume of traffic a network is engineered to handle. Voice networks are typically engineered to handle a target peak-load capacity, commonly measured in calls per second. Target peak-load capacities are specific to each business and industry, and are based on measured busy-hour call rates. For example, the traffic during the busiest hour on Mother's Day may be the target peak-load capacity for a residential voice network.

Redundancy

Redundancy measures the extra capacity, to be used only in the event of an equipment failure, that is placed in a network. When a primary node in a voice network is taken down for maintenance or failure, a redundant secondary device can take over the processing of the voice traffic.

Measuring Capacity and Redundancy

Engineers commonly use the notation n+k to describe the engineered capacity (n) and redundancy (k) of nodes in an availability solution. Consider a network which distributes across three SIP proxies, each rated with an individual capacity of 100 calls per second. Assuming that the proxies could act together in various logical combinations, there are several options for engineering capacity (n) and redundancy (k).

For example, each proxy could be deployed in a stand-alone configuration resulting in a total capacity (n = 3 proxy nodes) of 300 calls; however, there would be no redundancy (k=0) since if one proxy fails there has been no arrangement for one of the other proxies to pick up the traffic.

Since redundancy is critical it is more likely that the cluster will be engineered in a n+ k arrangement of (2 + 1) for a capacity of 200 calls per second. In this scenario, supporting the offered load of 200 calls per second only requires 2 out of the 3 proxies to process the call. The third proxy is available to take up the traffic (100 calls per second) in case one of the 2 primary proxies fails.

A third scenario would be to engineer the cluster of three proxies to support an offer load of 100 calls per second. This scenario only requires one proxy to handle the offered load of 100 calls per second so the additional two proxies are providing redundant capacity in case of a failure. This would be referred to as a engineered capacity and redundancy (n + k) of (1 + 2).

When engineering a high-availability solution for a voice network, one should determine the capacity (n) requirements first, and then add additional redundant nodes (k) to achieve the desired availability.

Designing a Highly Available Voice Network

Following is an example and depiction (Figure 1) of a highly available voice network that can handle two outbound calls per second.

In this example, each SIP gateway has four T1 trunk interfaces. At 24 channels per T1, a single SIP gateway can handle 96 concurrent calls. Assuming an average call length of three minutes, 96 calls over 180 seconds is ~0.5 calls per second. The SIP proxy servers are each rated at 100 calls per second.

Capacity of SIP gateway = 0.5 calls per second

Capacity of SIP Proxy Server = 100 calls per second

To handle the target capacity of two outbound calls per second, the network will need four gateways.

Figure 1
Voice Network Rated at Two Outbound Calls Per Second Capacity

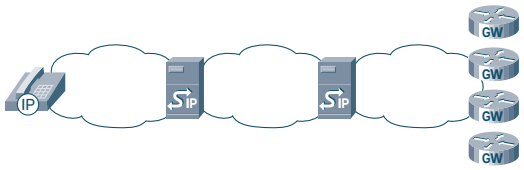

Figure 2 shows the addition of one redundant device at each point along the call path, enabling the voice network to handle two outbound calls per second, with 25 percent redundant capacity. It should be noted that the overall additional capacity of this particular scenario is limited to 25% because only one additional gateway was added. The addition of each proxy actually doubles the capacity of each proxy cluster. Thus, the overall capacity of this scenario could be increase to 50% by simply adding one more gateway for a total of six gateways.

Figure 2
Voice Network Rated at Two Outbound Calls Per Second with 25 Percent Redundant Gateway Capacity

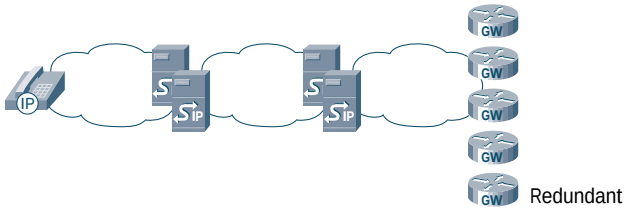

Placing Intelligence in the Network

There are several mechanisms that can be used to achieve load distribution and failover. These can be divided into two groups:

1. placing intelligence in the previous hop
2. placing intelligence in the redundant hop

Intelligence in the Previous Hop

Depending on the device in the previous hop, intelligence in the previous hop is typically provided through a backup proxy or DNS SRV (the concept of DNS SRV is defined later in this paper) in a user agent, through static routes or DNS SRV in a SIP proxy, and through dial peers or DNS SRV in a SIP gateway.

Backup Proxy, Static Routes, and Dial Peers

- *SIP user agents* can typically set a backup proxy to take over if the primary proxy fails to respond.
- Each static route in a *SIP proxy server* can typically be assigned a weight (for load balancing) and/or a priority (for redundancy failover).
- Redundant dial peers can commonly be set up in *SIP gateways*, should the primary dial peer fail to respond.
- Static routes can be configured for load balancing and redundancy, whereas a user agent's backup proxy and a SIP gateway's dial peers only provide redundancy failover.

DNS SRV

Most network administrators are familiar with DNS A records. These records associate the name of a particular *device* with the IP addresses of the network cards of that device. DNS SRV records associate the name of a *service* with the IP addresses of devices which provide that service (SIP, for example). DNS SRV records are commonly configured to point to the SIP proxies in the local domain which accept SIP packets (over UDP): "_sip._udp.company.com". The specification for DNS SRV can be found in RFC 2782, at: http://www.ietf.org/rfc/rfc2782.txt.

Each device listed in a DNS SRV record has an associated weight and priority. By adjusting the weight and priority of routes, system administrators can tailor the capacity and redundancy behavior of the SIP proxy server to their routing needs.

To take advantage of DNS SRV records, SIP clients must be programmed with the intelligence to parse and act upon DNS SRV records. Although most client vendors claim to support DNS, they typically mean DNS A records, not DNS SRV records. It is important to verify that a SIP client specifically supports DNS SRV records.

Intelligence in the Redundant Hop

Intelligence can be placed in the redundant hop by adding a pair of redundant load balancers or enabling Virtual Router Redundancy Protocol intelligence in the redundant-hop devices.

Virtual Router Redundancy Protocol

Virtual Router Redundancy Protocol (VRRP) defines a standard mechanism that enables a pair of redundant (1+1) servers on a LAN to negotiate ownership of a virtual IP address. One device is elected to be active and the other to be standby. If the active fails, the backup server takes over. An advantage of this scheme is that it achieves 1+1 redundancy without requiring any special intelligence in previous-hop SIP devices. However, this scheme only works for n=1 capacity and k=1 redundancy; it will not scale above 1+1. VRRP is described in RFC 2338, at: http://www.ietf.org/rfc/rfc2338.txt. Another protocol that accomplishes the same thing is Hot-Standby Routing Protocol (HSRP).

Load Balancers

Load balancers can be used to achieve redundancy and active load balancing. Load balancers can be configured as 1+1 redundant, using VRRP to negotiate active ownership of the virtual IP address of the redundant node. Load balancers are typically used when the redundant devices do not natively support VRRP themselves.

Load Distribution and Failover

The intelligence to distribute load and failover must be placed in devices in the network. When a call is placed through a highly available node in the network, the load balancing and redundancy intelligence can be placed either at the previous hop in or the redundant hop.

Between a SIP user agent and its outbound SIP proxies, DNS SRV or backup proxy intelligence can be placed in the calling user agent or a pair of redundant load balancers or VRRP daemon (VRRPD) intelligence can be placed in the SIP proxies (Figure 3). With low call rates (under 100 calls per second), placing VRRPD intelligence in the SIP proxies is preferable, due to the low implementation costs and the flexibility to serve less intelligent endpoints.

Figure 3
Illustration of Two Options for Location of Load Failure and Distribution Intelligence Between SIP User Agents and Proxies

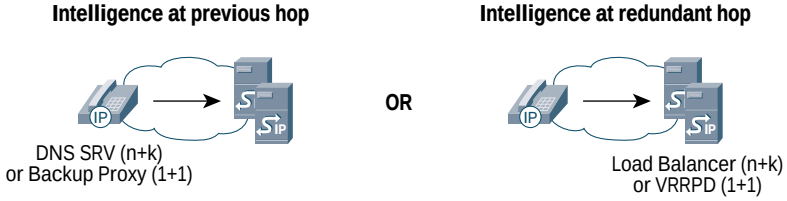

Between proxies, static routes or DNS SRV intelligence can be placed in the previous-hop SIP proxies or a pair of redundant load balancers or VRRPD intelligence can be added to the redundant SIP proxies (Figure 4). Using static routes is typically easiest.

Figure 4
Illustration of Two Options for Location of Load Failure and Distribution Intelligence Between Interconnecting SIP Proxy Clusters

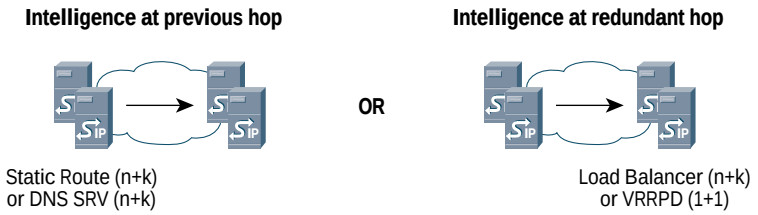

Between SIP proxies and outbound SIP gateways, static routes or DNS SRV intelligence can be used at the previous-hop SIP proxies or a pair of redundant load balancers or VRRPD intelligence can be placed at the SIP gateways (Figure 5). Using static routes in the previous-hop SIP proxies is typically easiest.

Figure 5
Illustration of Two Options for Location of Load Failure and Distribution Intelligence from SIP Proxies to SIP Gateways

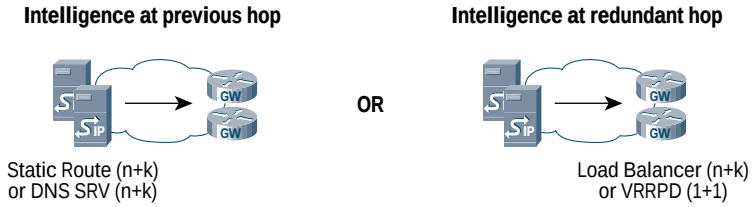

For inbound calls from SIP gateways to SIP proxies, either dial peers or DNS SRV intelligence can be used in the SIP gateways or a pair of redundant load balancers or VRRPD intelligence can be used in the SIP proxies (Figure 6). Dial peers are usually easiest, but may be unnecessary if the SIP proxies are already configured with load balancers or VRRPD intelligence.

Figure 6
Illustration of Two Options for Location of Load Failure and Distribution Intelligence from SIP Gateways to SIP Proxies

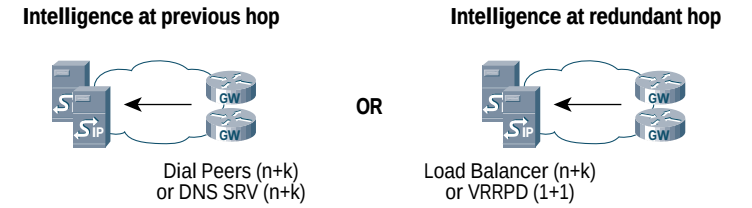

Priority and Weight

In redundancy schemes, routes to a primary and secondary server are negotiated based on their designated *priority.* A system administrator typically configures the priority of each route. For the solutions described in this document, the smaller number is considered to have the higher priority. For example, a route with a priority of 2 is secondary to a route with a priority of 1. When two or more routes have the same priority, traffic is load balanced across both servers.

In load-balancing schemes, new requests are distributed across available servers using a selection algorithm. A common selection algorithm for statistical load balancing is *round-robin.* A more general selection algorithm is *weighted random*, which distributes requests proportional to the *weight* assigned to each route. To load balance between routes to devices, the priority of the routes must be the same. A higher weight represents a proportionally higher device capacity. In many cases, the capacity of each device is

the same, so the weight associated with the route to each device is 1. If the capacities of the devices are different, the weights are adjusted accordingly. For example, if device A has three-quarters the capacity of device B, then the weight of the route to device A would be 3, and the weight of the route to device B would be 4.

By configuring the priority and weight of static route or DNS SRV entries, system administrators can tailor the capacity and redundancy behavior of their voice networks.

Designing Voice Networks for 99.999 Percent or "Five Nines" Reliability

With the knowledge of component-specific call success rates, it is possible to calculate aggregate call success rates. In the following examples, the impact of adding a redundant SIP proxy to a SIP proxy server farm is illustrated. In these examples, an individual SIP proxy call success rate of 99.96 percent is assumed.

Figure 7 illustrates a single SIP proxy in the SIP farm (1+0).

Figure 7
Call Success Rate without Redundancy

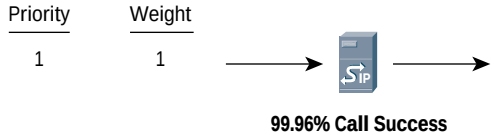

$$P_{success}(A_1) = 0.9996 = 99.96\%$$

Improve the call success rate through the SIP proxy farm by adding a redundant SIP proxy server. Figure 8 illustrates a 1+1 redundancy configuration of two SIP proxy servers.

Figure 8
Call Success Rate with One Redundant SIP Proxy Server

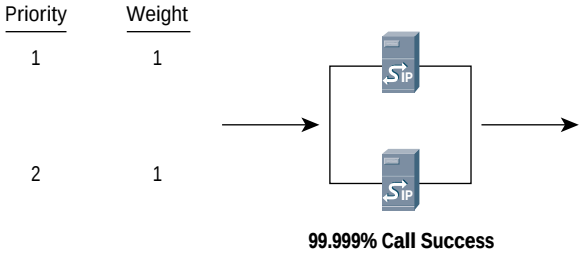

$$P_{success}(A_1 \cup A_2) = 1 - P_{failure}(A_1 \cap A_2) = 1 - P_{failure}(A_1) \cdot P_{failure}(A_2)$$
$$= 1 - (1 - P_{success}(A_1)) \cdot (1 - P_{success}(A_2))$$
$$= 1 - (1 - 0.9996) \cdot (1 - 0.9996)$$
$$= 99.999984\%$$

Putting It All Together

By placing the desired combination of high-availability routing intelligence in the network, we can complete the high-availability solution. In the following example, the endpoints and gateways remain "simple," placing the high-availability intelligence in the SIP proxies (Figure 9).

Figure 9
Highly Available SIP-based Network with Intelligence Placed in SIP Proxy Servers

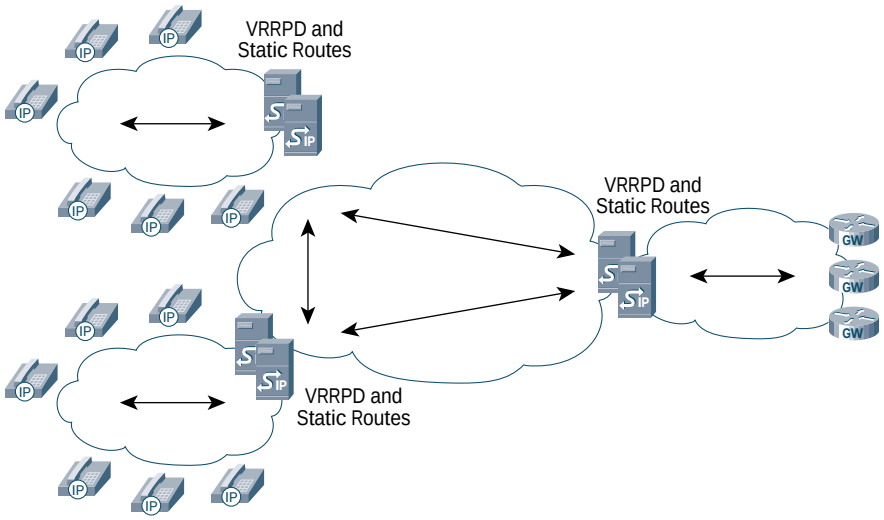

Conclusion

There are various availability solutions that are customizable to your business needs. Implementing a high-availability solution for your VoIP network will improve your voice network's uptime, and by preventing outages, will lower the costs associated with downtime.

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems Europe
11 Rue Camille Desmoulins
92782 Issy-les-Moulineaux
Cedex 9
France
www-europe.cisco.com
Tel: 33 1 58 04 60 00
Fax: 33 1 58 04 61 00

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
Capital Tower
168 Robinson Road
#22-01 to #29-01
Singapore 068912
www.cisco.com
Tel: +65 317 7777
Fax: +65 317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the
Cisco Web site at www.cisco.com/go/offices

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia
Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland
Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland
Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden
Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

MH/CI/0702