

agileTel Runs Active/Active with Linux and MySQL

October 2013

In a previous article,¹ we looked at the paucity of active/active implementations in the commodity server world. We review in this article one such successful example of redundant, geographically distributed systems using Linux servers and a MySQL database in an active/active configuration.




agileTel

agileTel (www.agitел.com), a business unit of NetWeave Integrated Services, provides small to medium businesses with a high performance, high quality Internet VoIP (Voice over IP) virtual phone system that provides the features of an enterprise phone service at no incremental cost. According to caller ID and time of day, incoming phone calls can be routed to extensions, to mobile phones, or to an automated attendant. Voicemails can be sent to users' email inboxes, and users can listen to their messages via their smartphones or computers. SMS text messages are sent to notify users that a voicemail has been received.

Least-cost routing provides the lowest cost for outgoing calls. A web-based call-accounting interface allows customers to track call usage. There is no on-site equipment except for the IP-based telephone handsets. Analog handsets can also be used through Analog Telephone Adapters (ATA's).

The agileTel system

The agileTel system uses a three-tier architecture – IP phones, virtual machines, and data centers:



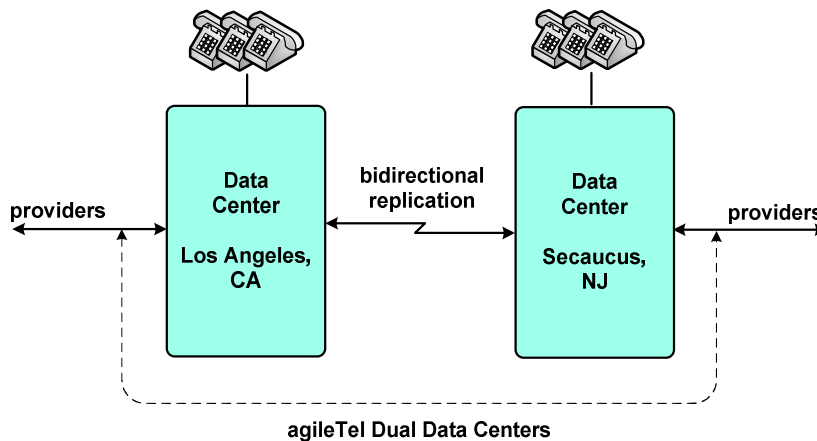
agileTel Three-Tier Architecture

- The IP phones are digital devices located on the customers' premises.
- Each IP phone is assigned to a Virtual Machine (VM) residing in a nearby cloud that handles the routing of the call.
- The VMs interface with agileTel data centers that provide the interface to telephone provider networks.

¹ *Active/Active on Commodity Servers*, *Availability Digest*, September 2013.
http://www.availabilitydigest.com/public_articles/0809/commodity_active_active.pdf

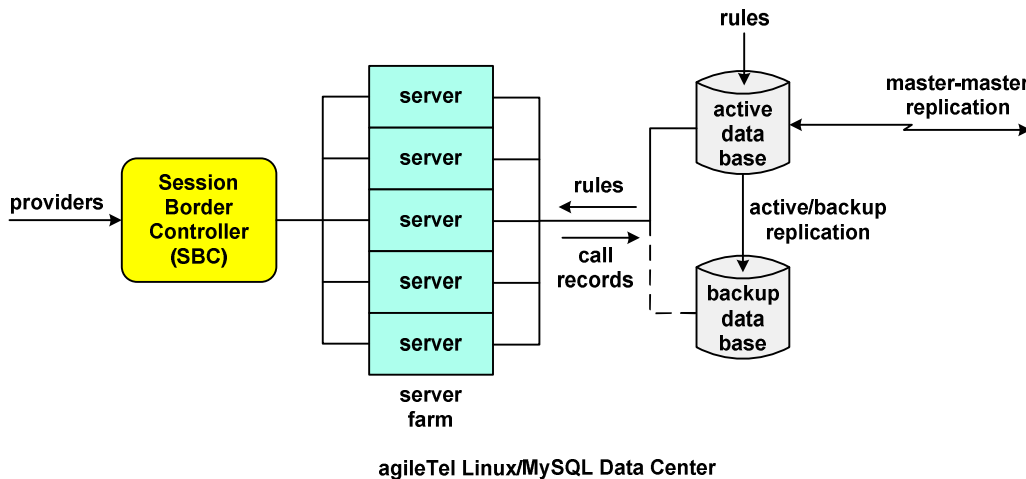
agileTel operates two U.S. data centers, one on the West Coast in Los Angeles and one on the East Coast in Secaucus, New Jersey. Providers such as Level3, Verizon, AT&T, and Sprint, route incoming calls through the PSTN (Public Switched Telephone Network) to one of the two agileTel data centers. The data centers share the call load. It is the data center's job to route the calls appropriately to the called agileTel parties.

The data-center databases are kept synchronized via bidirectional replication so that any call can be handled by either data center. Therefore, a provider may route a call to the data center of its choice. Should a data center fail, providers simply route all calls to the surviving data center.



An agileTel Data Center

The primary workhorse in an agileTel data center is a farm of Linux servers. They are responsible for connecting incoming calls from providers, for placing outgoing calls to providers, for providing the services available to the called parties, for maintaining the connections during the calls, and for terminating the calls.



When a call request is first received at a data center from a provider, it is directed to a Session Border Controller, or SBC. The SBC determines which server in the server farm should handle the call. This selection is made according to server load (the most lightly loaded server is used) and on a round-robin basis between servers handling the same load. If a server is unresponsive, the SBC will try another server. Servers in the local server farm take precedence, but the SBC can also route calls to the servers in the opposite data center. If the SBC cannot find an operational server, it will reject the call back to the

provider, who will try the other data center. The provider also tries the other data center if a problem exists with the SBC.

Each server has the rated capacity to handle 500 simultaneous calls. Currently, the data centers are configured with four to five servers each, giving a total call capacity for the agileTel system of about 5,000 simultaneous calls. This capacity is easily expanded by adding servers. The database servers are relatively lightly loaded in comparison to the server farm that manages calls. An active/passive database server pair is rated to handle approximately 20 servers in the server farm.

The SBC, call handling servers, and database servers are Dell rack-mounted servers running Linux. A typical server comprises four four-core CPUs.

The servers have access to a redundant MySQL database configured as an active/backup pair. The backup database is kept synchronized with the active database by MySQL data replication. Should the active database fail, the backup database is put into service in just a few seconds. This active/passive architecture uses the Linux-HA heartbeat system, with the active member publishing a virtual IP address. In the event that the passive member detects loss of the primary (2 missed heartbeats), it will publish the virtual IP address and take over primary database operations. When the former primary returns, the former passive member will detect it and relinquish the virtual IP address to the primary. The servers in the server farm are essentially unaware of these operations with the expectation that occasionally they may find that they have to reconnect to the database when the virtual IP switches.

The database stores two primary sets of data:

- *Rules* governing the routing for each called number, such as the VM to which to route the call, forwarding to a mobile phone, varying the endpoint according to the time of day, or routing to an automated operator (i.e., "Press 7 for sales.") The Rules Database is created and maintained by agileTel personnel via administrative terminals. agileTel provides flat rate calling plans for all domestic phone operations, so call rating is not a significant component as it only occurs for international calls.
- *Call records* written at the end of each call giving the details of the call (calling number, called number, time and date, duration, etc.) Call records are inserted into the Call Record Database by a server upon the completion of a call that it is handling.

In order to run the two data centers as an active/active pair in which a call can be handled by either data center, it is important that the two data centers have the same view of the database. This is accomplished by the use of the Tungsten Master-Master Replication engine from Continuent (www.continuent.com) to keep the active databases synchronized. Whenever a change is made to the Rules Database or an insert is made into the Call Record Database, that update is not only replicated to the active database's backup database, but it is also replicated from the active database to the active database in the other data center. The receiving active database will then replicate it to its own backup database.


Data conflicts (the simultaneous updating of the same data field at both data centers) are unlikely since the majority of database activity is call-record inserts.

Call Establishment

When a provider submits a call to a data center, the SBC receives the request. The SBC selects a server from the server farm to handle the call and passes the called number to the server. The server accesses the Rules Database for the entry corresponding to the called number to determine the appropriate VM to which to route the call.

The server then passes the call and the appropriate rules to a virtual machine (VM) located in a cloud that is near the called party. agileTel uses Rackspace and Linode clouds to host their VMs.


The VM provides the PBX functions for the customer and controls the IP telephones at the customer's site. Typically, there is one VM per customer, though very small customer installations may share a common VM and very large customers may have multiple VMs. Large customers have redundant VMs that replicate configurations between them every few minutes. Smaller customers do not have redundant VMs, but a replacement VM can be configured in a few minutes. The VMs use static IP addresses, and a set of phones can be rolled over to a new VM in two minutes via a DNS update.



The VM establishes an IP connection with the appropriate end device, and the call is now connected.


Connected Call

Once the two parties have been connected, the SBC is released from the call and the conversation proceeds. The conversational path comprises the provider, the selected server, the called party's VM, and the called party's IP phone. Analog-to-digital conversion is accomplished by the provider.



Call Disconnect


Either end of the call can initiate the disconnect. The server and the SBC coordinate the disconnect in either case to forward the hang-up signal to the opposite end of the call. The SBC will terminate the call with the provider. The server terminates the call with the agileTel end-device via its VM and writes a call record to the Call Record Database. The call record is replicated to the active database at the remote data center to keep the databases synchronized.



Outbound Calls

If an agileTel user wants to make a phone call, he is connected to his local VM. The VM will connect to a server in one of the data centers. If it is unable to do this, the VM will try the other data center.

The called number is passed to the server, which will choose a provider based on a Least Cost Routing (LCR) algorithm. The provider will complete the call.



Summary

There are three primary considerations when deciding whether to run a critical application as a continuously available active/active architecture:

1. Can the application be structured to run in a distributed environment with other active copies of the application?
2. Is there a reliable and efficient bidirectional replicator available?
3. Is the cost of going active/active warranted by the cost of downtime?

In agileTel's case, the answer to all of these questions is a "yes." The processing of each transaction (a call request) is independent of all other transactions, so that running multiple such transactions in multiple sites does not cause any conflict.

The Tungsten Master-Master bidirectional replicator is an excellent MySQL data-replication engine. Other third-party packages such as Shadowbase (www.gravic.com/shadowbase) and Oracle GoldenGate (www.oracle.com) support bidirectional replication between a variety of databases, including Oracle, SQL Server, MySQL, DB2, Sybase, and Ingres.

And finally, any downtime for the agileTel telephone system is simply unacceptable. If agileTel's system is down, its customers are cut off from the rest of the world.