

Building a Distributed Call Center

Who Am I?



- Leif Madsen, Asterisk consultant
- Co-author of Asterisk: The Future of Telephony
- Asterisk bug marshal
- Asterisk release manager
- Specialize in database integration and call center implementations



Overview

Topologies

- -Describe various call center setups
- -Start simple; increasing complexity
- -Call center types
 - Traditional
 - Hybrid
 - Database driven
 - Distributed

Apologies Now



- Initial presentation idea was a bit ambitious for 35 minutes
- Typically have 50 minutes of talk time
- Will follow up with blog posts and dialplan logic
- Already have OpenAIS and XMPP based labs developed
- Watch http://www.asterisk.org



Topologies and Discussion

Traditional Call Center







- Calls come from PRI into traditional PBX
- Delivered to agents via standard ACD (Automatic Call Distribution) rules
- Not distributable; hard (impossible?) to deliver calls to remote employees
- Typically need to be physically connected to the phone system

Hybrid System





Hybrid System



- Asterisk is the enabler; SIP end-points add new technologies to existing systems
- Still limited by physical connections
- Typically expensive (PRI) or unreliable (analog)
- All hardware needs to reside at call center facility
- Allows remote employees (yay!)

Pure Asterisk, Non-Distributed





Pure Asterisk, Non-Distributed



- No limitation on number of lines (within reason)
- No need for expensive lines at all; could be entirely SIP based
- Remote employees easily added to system
 - save on electricity costs
 - increase employee moral
 - serve more timezones
- Simple and efficient, but limited in expandability

Asterisk + Database





Asterisk + Database



- Slightly more complex, but easier to update
- Create own GUI system; need not be complex
 - func_odbc recommended; dynamic data, static dialplan
 - also like func_curl
- Start clustering; hot-swap via LinuxHA
- Keep /etc/asterisk in (r)sync
- Changes are nearly realtime

Asterisk + Replicated Database





Asterisk + Replicated Database



- Master-Master Replication with MySQL
- Perform load balancing
- Help prevent downtime with redundant live databases
- Reasonably easy to setup and maintain
- Failover with pen; res_odbc also fails over

Asterisk + Database, w/ Distributed Device State





Asterisk + Database, w/ Distributed Device State



- For our Swedish friends in the audience; this is sexy
- Multiple systems know status of devices remotely
- Allows agents to reside on multiple Asterisk systems
- Can login to queues across the cluster
- Status of agents are known on other servers; no multiple ringing or excessive attempts
- With OpenAIS (already available) you can only use over a low latency link

Distributed Device State Over WAN





Distributed Device State Over WAN



- New branch allows device state to be distributed over WAN links
- Uses Tigase XMPP server right now; plans to expand to other XMPP servers
- Allows agents to reside in different physical locations and answer calls from queues across the cluster
- Can provide redundancy by failing over calls from PSTN or ITSP to other servers

Multi-Queue, Multi-Site





Multi-Queue, Multi-Site



- Load distributed across multiple systems
- Useful where queues are heavily utilized
- Expand number of agents by adding additional servers
- Agent servers can perform other tasks
- Save processing power on queue servers for just queues





- For large queues, must be on same system
- No ability to share queue position
- Does it matter?
- Will callers know they are the 3rd caller on one system vs. multiple systems?



Started simple; increased in design complexity

Review

- Many ways to grow as your company expands
- Building a simple system now, with good forward planning, can be expanded
- Increase employee moral and save on hardware and electricity costs with remote agents





Questions?