

Voice apps can strain storage

http://storageMagazine.techtarget.com/magItem/0,291266,sid35_gci1179419,00.html

by: [Ray Lucchesi](#)

Issue: [Apr 2006](#)

As companies implement digital voice applications, storage managers need to prepare for these new data sources.

Voice over Internet Protocol (VOIP)

phone calling is accelerating the evolution from an analog to a digital world and bringing new storage challenges to the data center. Before the digitization of voice, analog voice traffic was recorded on tape and perhaps later transferred to CDs or DVDs. Some of these recordings might have been listened to--typically 30 to 90 days after their creation--and then the media was archived or jettisoned. Some shops kept their voice-recording media for a longer period of time, but as the number of CDs grew, storage management problems multiplied.

It's much easier today to store digital voice recordings as a file or row item in a large relational database; all of the tools and processes of modern data centers can be used to store, index, search, archive and expire this digital data. With so many different options to store voice that's been captured in a digital format, storage managers need to develop a plan to incorporate VoIP recording data into their data center storage infrastructure using a tier of storage that's cost-effective and appropriate for the value of the data.

How VoIP works

Telephone data over the Internet uses Session Initiation Protocol (SIP) packets to initiate phone call activity. VoIP data

can be recorded directly from VoIP systems, or indirectly using standard IP networking tools to replicate traffic destined for a telephone adapter via an IP or MAC address. In either case, data traffic ends up at a voice-recording system server, which reconstitutes the telephone session from the IP packets and creates a file or database row.

Most systems add meta data associated with a voice recording, such as a time-date stamp, and possibly Caller ID and target phone number. This additional information is crucial for retrieval and processing of the voice recording; it can be appended to the voice-recording database row or be in a separate file directory associated with the voice recording.

Voice recordings have historically been used for quality-control purposes (of phone-based services) and for possible legal proceedings; this is especially true for 911 and other emergency-response center phone sessions. Financial service firms routinely document customer order details in case an order is subsequently questioned. New compliance regulations mandate record retention, but compliance is mainly a function of the content of a record, not its physical form, says Mike Casey, vice president of services development and

VOIP Storage Requirements

VoIP recording consumes approximately 1MB of storage per minute of phone conversation. In general, most voice-recording systems use DAS, but many allow voice data to be recorded to SAN/NAS storage as well. According to Cara Diemont, marketing director for Data Dimension's Customer Interactive Solutions and managing editor of the firm's "Merchants Global Contact Centre Benchmarking Report," a mid-sized call center handles approximately 10 million calls per year. At that rate, 38,000 calls per workday would generate approximately 154GB of voice data each call-center day (assuming four minutes per call and recording all calls)--or approximately 750GB of voice data per week.

There are some adjunct files that are produced for each phone call as well. These don't amount to much data, but may represent a sizable number of files, considering there may be close to 1 million calls per month, which could quickly overwhelm a single file directory.

Industry statistics indicate that only 1% to 3% of stored voice recordings are listened to during the first 30 to 90 days; after that period, access rates drop precipitously and the voice-recording data should be moved to less-expensive tertiary or archival storage. Of course, it's also possible that many of the

marketing at Contoural Inc., a technology consulting firm in Mountain View, CA.

Some call centers record all telephone sessions, while others record only certain calls. Calls can be recorded in real-time by a telephone operator simply tagging a call session, or recorded based on the phone number called or activated automatically by other predetermined criteria. David Brandon, senior technical call-center consultant at Forsythe Solutions Group in Skokie, IL, says stress monitoring is one trigger that can be used to identify a voice recording for special archiving; if the stress level is high, the recording is tagged for later human review and stored in a special archive.

There are products that "word spot" a

recording, searching for a limited word vocabulary like "stop," "cancel," "halt" and so on. Paul McIntyre, product manager, customer interactive solutions at Dimension Data, says there are a few active pilots of new technologies that do real-time word spotting and stress-level assessment. As these products become more sophisticated in their ability to tag and find particular recordings, there will be more ways to data mine recordings. All of these advanced functions require online or nearline storage.

Storage considerations

VoIP sessions can be recorded at several different compression levels and saved as different file types, but a minute of voice recording generally consumes about a megabyte of storage (see "[VoIP storage requirements](#)"). Depending on the voice-recording system, the formats of the recordings may be proprietary or standard types, such as WAV or MP3 files. Proprietary formats typically compress better and are more secure, but they're available from only one vendor and may not be amenable to automated post-processing of the recording. A large call center might handle 10 million calls per year, with each call averaging about three to four minutes. Using four minutes as the length of an average call, and assuming that all calls are recorded, this would be approximately 40TB of data per year.

Storage used for direct call recording can run the gamut from DAS to SAN or NAS boxes, and everything in between. This would be considered primary storage in call-center parlance and used as a temporary holding place for the voice recording during the highly accessed stage of its lifecycle. Once this phase has passed--usually 90 days--the voice recording could be placed on secondary storage such as lower cost disk, an archive appliance, tape or optical media. EMC Corp.'s Centera, for example, is widely used for long-term voice storage because many voice-recording systems support the Centera API.

data looks an awful lot like the rest of your corporate data. The voice files might be larger and have more stringent real-time capture constraints, but the data needs to be online for a short period of time and then archived to a different tier of storage.

David Kaercher, vice president of core services at Allianz Life Insurance Company of North America in Minneapolis, says his company recently moved to VoIP and initially stored its voice recordings on tier-one (EMC Symmetrix) storage. As the volume of calls grew, tier-one became too expensive and the voice recordings were

STORAGE USED FOR DIRECT CALL RECORDING CAN RUN THE GAMUT FROM DAS TO SAN OR NAS BOXES, AND EVERYTHING IN BETWEEN

Historically, call-center disaster recovery planning has been separate from data center disaster recovery, but this is changing because VoIP call-center recordings are increasingly being considered mission critical. "When e-mail first came out, it wasn't deemed mission critical--look where it is today," says Robyn Danz, storage specialist at CDW Government Inc. "Voice is the new e-mail."

A call center generating 10 million files annually could easily overtax a single file directory. For call-center systems generating that volume of calls, it may be prudent to place the recordings into a large database instead of dumping all of the voice files into a file directory. Retrieval of a voice record is based on the meta data associated with the call.

From a storage perspective, digital voice

moved to tier-two (Network Appliance Inc. NAS) storage. Tier-two proved too expensive for long-term retention, so Kaercher wanted to move the voice data to an EMC Centera, which Allianz uses for e-mail archiving. But Kaercher couldn't convince the firm's VoIP vendor, Roswell, GA-based Witness Systems Inc., to support Centera's API. Today, Allianz is reconsidering its use of the Witness VoIP product and is planning on archiving voice recordings on an IBM Corp. tape library (see "Allianz Life's experience with Witness Systems").

Voice meta data

Call-center systems create a standard record for each telephone session. This record is typically short and contains just the specifics of each call, including time and date, operator number, call duration, phone number (from and to), wait queue time and so forth. This meta data is

Allianz Life's Experience with Witness Systems

Allianz Life Insurance Company of North America in Minneapolis reviewed various call-center products before deciding on Roswell, GA-based Witness Systems' voice-recording services. With a few exceptions, the firm was extremely happy with how well the system installed and performed. Like most organizations, Allianz Life began using DAS to retain its voice recordings. It wasn't long before the DAS arrangement became overtaxed and a storage upgrade was required.

Allianz tried tier-one storage, but found it too expensive for voice data. The company moved its VoIP data to tier-two storage, but found that it, too, wasn't a good fit for long-term storage of its voice-recording data. Allianz next turned to archive storage and felt that its EMC Centera, which hosted the firm's e-mail archives, would be a natural fit for its voice recordings.

When Allianz Life asked Witness Systems if it would consider supporting Centera's API, Witness declined, says David Kaercher, vice president of core services at Allianz. EMC said it was relatively straightforward to implement the API and estimated that the customization work to support its API would cost approximately \$25,000. Allianz later learned that Witness Systems was considering introducing its own archive product and had no desire to support Centera. Nonetheless, Allianz Life still wants to move data off its tier-two storage after 30 days and now plans to archive its voice recordings on an IBM Corp. tape library. This particular Allianz Life call center receives about 100,000 calls per month, which adds up to approximately 2TB to 3TB per year. The firm plans to retain this voice data indefinitely once it's moved to tertiary storage.

joined to the voice recording. In addition, periodic screen captures--every 100 milliseconds to one per minute--are sometimes saved during the call to correlate an operator's actions with voice and call detail records. This data is vendor-specific to the call-center software and is often used to audit the actions of telephone operators to determine, for example, that call scripts are followed properly.

Some systems also provide computer-to-telephony interface data, which can be

extracted for the call. This data might not appear on the operator's screen, and could include information such as operator name, login duration and customer account information not pertinent to the call. The voice recording, call detail record, screen captures and other computer-generated data can provide a good view of a specific call and require a great deal of storage.

The leading voice-recording systems on the market are Nice Systems Inc., Rutherford, NJ, Verint Systems Inc., Melville, NY, and Witness Systems; all are tightly coupled with call-center operations software. Recording voice becomes much easier if all phone traffic is funneled through an Ethernet switch to the voice-recording system. Voice recording is also available for

analog phone equipment. In the old days, this was done with what amounts to line or PBX taps using proprietary hardware/software systems.

IP PBX storage

IP PBXs often use internal disk for voice mail storage. Depending on the number of user mailboxes and the maximum amount of voice mails per mailbox, storage requirements can escalate quickly. A typical small IP PBX might have a 40GB drive supporting 20 to 40 voice mailboxes, and could store up to

17 hours of voice mails for 40 users. High-end PBX systems have options for much more voice mail storage. In some cases, the move to VoIP has merged voice mail and e-mail. You can receive your voice mails as an e-mail with a WAV file attachment and maintain your

begin to accumulate, calls can be analyzed and data mined to determine what type of response works best to drive up the yield from telephone operations centers. But real-time feedback of this sort is years away and will certainly boost storage requirements.

IN SOME RESPECTS, VoIP APPLICATIONS MAY MEAN JUST ANOTHER 40TB-PER-YEAR DATA STREAM FOR YOUR DATA CENTER TO HANDLE

voice mails within your e-mail repository with these systems.

Voice analytics are mechanisms used to extract additional information from a voice stream. Voice analytics used in real-time can provide hints to the operator on how to handle the person on the call. If, for example, the person on the call is stressed, the analytics may tell the operator not to try to sell them more products/services, but to simply help the caller. Once voice recordings

In some respects, VoIP applications may mean just another 40TB-per-year data stream for your data center to handle. But digital voice applications and storage have their own unique lifecycle and performance aspects and, if not properly planned for, can easily strain your already-stretched network infrastructure and SAN/NAS/archive storage. Voice recordings are beginning to look a bit like e-mail during the early years, but roughly 1,000 times larger per element

Ray Lucchesi is president of Silverton Consulting and since 2004 has helped startups to Fortune 500 companies improve storage product development and marketing. Ray has also helped end-users better select and use storage subsystems.

<mailto:info@SilvertonConsulting.com>
<http://www.SilvertonConsulting.com>

Storage magazine is a part of the [TechTarget](#) portfolio of enterprise IT-focused media.

Copyright 2001-2006, TechTarget. All Rights Reserved.