SUBSCRIPTION SERVICE

May 9, 2008

COMPARING CALL CONTROL IN MICROSOFT OCS 2007 AND IBM LOTUS SAMETIME UNIFIED TELEPHONY

COMPANY(S) MENTIONED:
IBM, Microsoft, Nortel, Siemens.

PRODUCT(S) / SERVICE(S):
IBM Lotus Notes, IBM Lotus Sametime Unified Telephony, Microsoft Exchange 2007, Microsoft Office Communications Server 2007, Nortel MCS 5100, Nortel Agile Communications Environment, Siemens OpenScape.

Wainhouse Research Score	3.5*
Our rating scale:	
Will affect all vendors/users	5.0
Will affect most vendors/users	4.0
Will affect some vendors/users	3.0
May affect a few providers/users	2.0
Will not affect providers/users	1.0

ANALYST(S):
E. Brent Kelly, Ph.D.
bkelly@wainhouse.com

OVERVIEW

He who owns the call control owns the presence engine, the user interface, and significantly influences the available endpoints in a unified communications solution. PBX vendors and software vendors, notably IBM and Microsoft, will be aggressively slugging it out in call control battles in an effort to gain a dominant position in an enterprise's communications spend. This research note examines several call control strategies, and focuses on the differences in architecture and function between the telephony call control available in Microsoft OCS 2007 and IBM Lotus Sametime Unified Telephony.

INTRODUCTION

Call control has traditionally been a function of the enterprise PBX. As "Communicator-like" clients start to proliferate on the desktop, enabling click-to-call, click-to-conference, and softphone capabilities, call control is becoming more and more a software function.

The telephony vendors continue to rely on their PBXs as the main call control engine, interfacing their PBX platform to their particular brand of communicator client. Many of these vendors have also integrated video devices with their PBXs to provide click-to-call and short-digit dialing between video endpoints and either telephones or other video devices.

In a homogeneous environment in which all phones are connected to the same brand of PBX, using the PBX as the main call control engine along with that vendor's PC client makes sense, and it actually has some cost and ease of installation advantages. However, in a heterogeneous environment, where there are PBXs from multiple vendors, companies face serious challenges providing users with the same functionality and user experience. In these instances, companies are turning to software vendors that provide unified clients and call control engines that span the PBXs in a heterogeneous telephony environment, and that give users not only consistent functionality, but the same functionality and the same user experience.

Although Nortel (MCS 5100 and the new Nortel Agile Communication Environment) and Siemens (OpenScape) both create call control engines and Communicator-like clients that work in heterogeneous telephony environments, IBM and Microsoft are clearly the leading contenders when it comes to multi-vendor PBX deployments and desktop clients.

MICROSOFT OCS 2007 TELEPHONY INTEGRATION ARCHITECTURE

Microsoft has built both presence and call control capabilities into its Office Communications Server 2007 solution that was released in October 2007. In an OCS solution, the Communicator client runs on the PC, call control resides in the OCS 2007 Server, and an OCS Mediation Server sits between the OCS Server and the enterprise PBX as illustrated in figure 1.

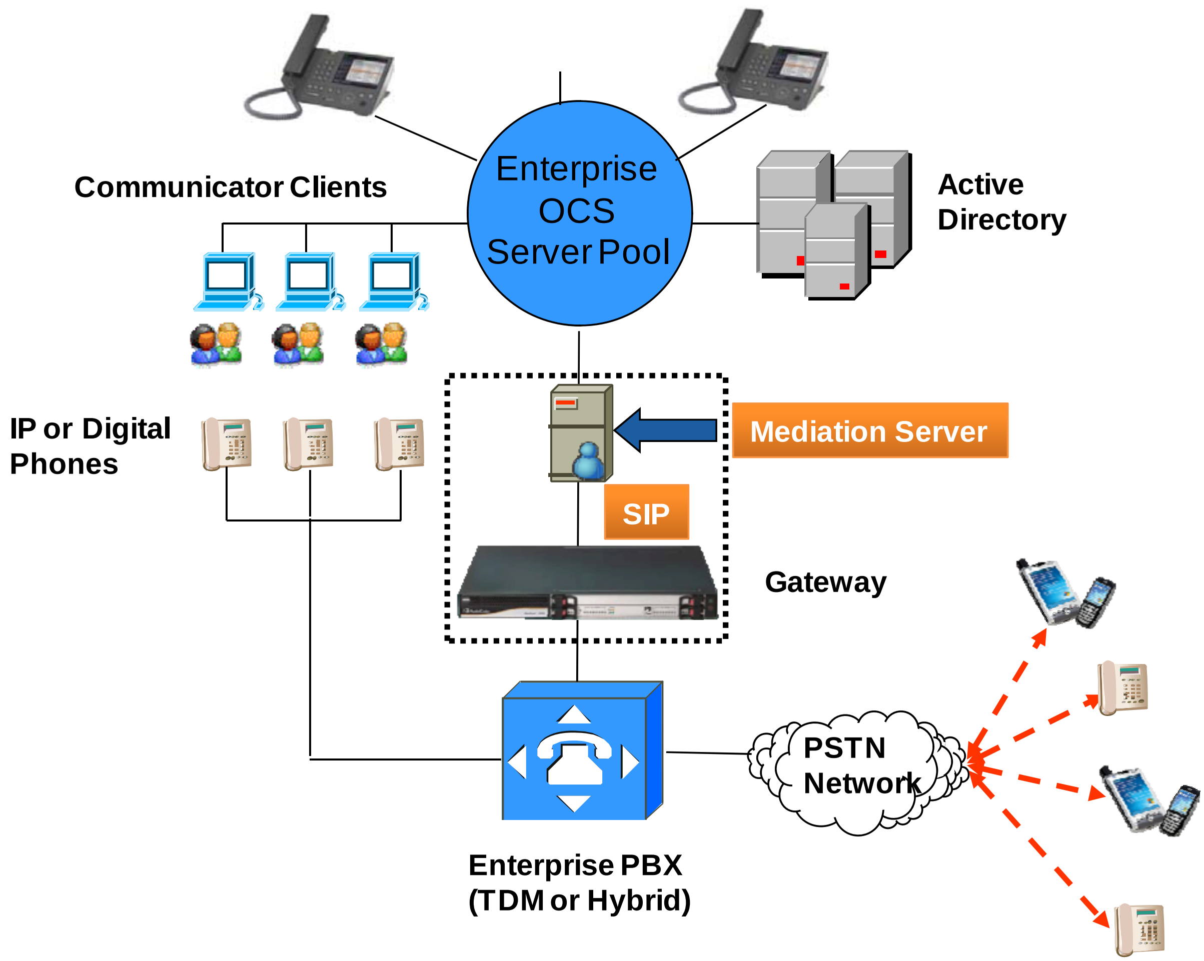

Figure 1: Microsoft OCS 2007 telephony integration architecture.

Mediation Servers are Figure 1 shows the Mediation Server residing between the OCS server and a gateway/PBX. The Mediation Server is the path for all calls between OCS and the PSTN. Thus, any location where an enterprise wants/needs to interface OCS calls with the PSTN requires a Mediation Server. Thus, Mediation Servers are distributed to the edge of an organization's network.

A gateway or PBX is always required with a Mediation Server; these devices provide the route to the PSTN. There are several use cases:

1. If an enterprise wants an exit route to the PSTN in a location where no PBX already exists, a gateway can provide PSTN connectivity. This may be useful in some least cost routing scenarios, for example.
2. If the enterprise already has a PBX at a location, and the PBX is already connected to the PSTN, then a gateway may or may not be required as follows:
    a. If the PBX supports direct connectivity to OCS, then no gateway is required.
    b. If the PBX is a TDM device, or it is an IP PBX, but it does not support OCS directly, then a gateway is required between the PBX and the Mediation Server. Although OCS runs SIP, the IP PBX manufacturers must modify their SIP stack to interface directly with OCS. Note that Nortel's CS1000 PBX is certified to interoperate directly with OCS, and several other manufacturers are working on certification.

A single Mediation Server can point at only one PBX or gateway; however, depending on the PBX or gateway capacity, multiple Mediation Servers can point at the same PBX or gateway. Mediation Servers can also run in tandem to provide redundancy: new call routing alternates between them. Mediation Server call capacity depends upon the power of the processor running the Mediation Server as shown in table 1.

Hardware	T1 's Supported	E1's Supported
Single processor, dual core, 2 GHz, Memory: 2GB RAM 2 x 1 GBit NIC	4 (96 calls)	3 (96 calls)
Single processor, dual core, 3 GHz Memory: 2GB RAM 2 x 1 GBit NIC	5 (120 calls)	4 (128)
Dual processor, dual core, 3 GHz Memory: 2GB RAM 2 x 1 GBit NIC	10 (240 calls)	8 (256 calls)
Dual Proc/ Quad Core 2.66 GHz, Memory: 2GB RAM 2 x 1 GBit NIC	18 (432 calls)	14 (448 calls)

Table 1: Mediation Server call capacities.

Mediation Servers perform a number of functions, among which are the following:

1. They transcode between Microsoft's proprietary audio codecs to standard codecs like G.711.
2. They provide intermediate signaling and call flows via a back-to-back-user-agent. With this capacity, the Mediation Server adds or removes Microsoft proprietary SIP header elements that are not supported by most PBXs and gateways.
3. They can act as an ICE[1] client to enable PSTN-originated calls to traverse intervening NATs and firewalls. This is used when remote users place or receive a call that is routed through the PSTN.
4. They enroll in the management, monitoring, and provisioning scheme Microsoft has provided.

Microsoft is working with all of the major PBX manufacturers to support interconnectivity with the OCS Mediation Server directly. Furthermore, it is also working with gateway manufacturers to support gateway compatibility. The Dialogic 4000 gateway has the Mediation Server embedded so that enterprises using it only need to deploy a single edge device.

[1] Interactive Connectivity Establishment. ICE has been designed to work with the SIP protocol to provide NAT and firewall traversal capabilities. In OCS, this is useful for remote users who use Microsoft Office Communicator over an Internet connection without a VPN.

IBM LOTUS SAMETIME UNIFIED TELEPHONY INTEGRATION ARCHITECTURE

To provide call control capability for Lotus Sametime, IBM has created a new server called the Telephony Application Server. This server interfaces with Sametime via the Sametime Virtual Places (VP) protocol[2]. In addition, a second server, the Telephony Control Server, is used to interface with PBXs and gateways. The Sametime Unified Telephony architecture is illustrated in figure 2.

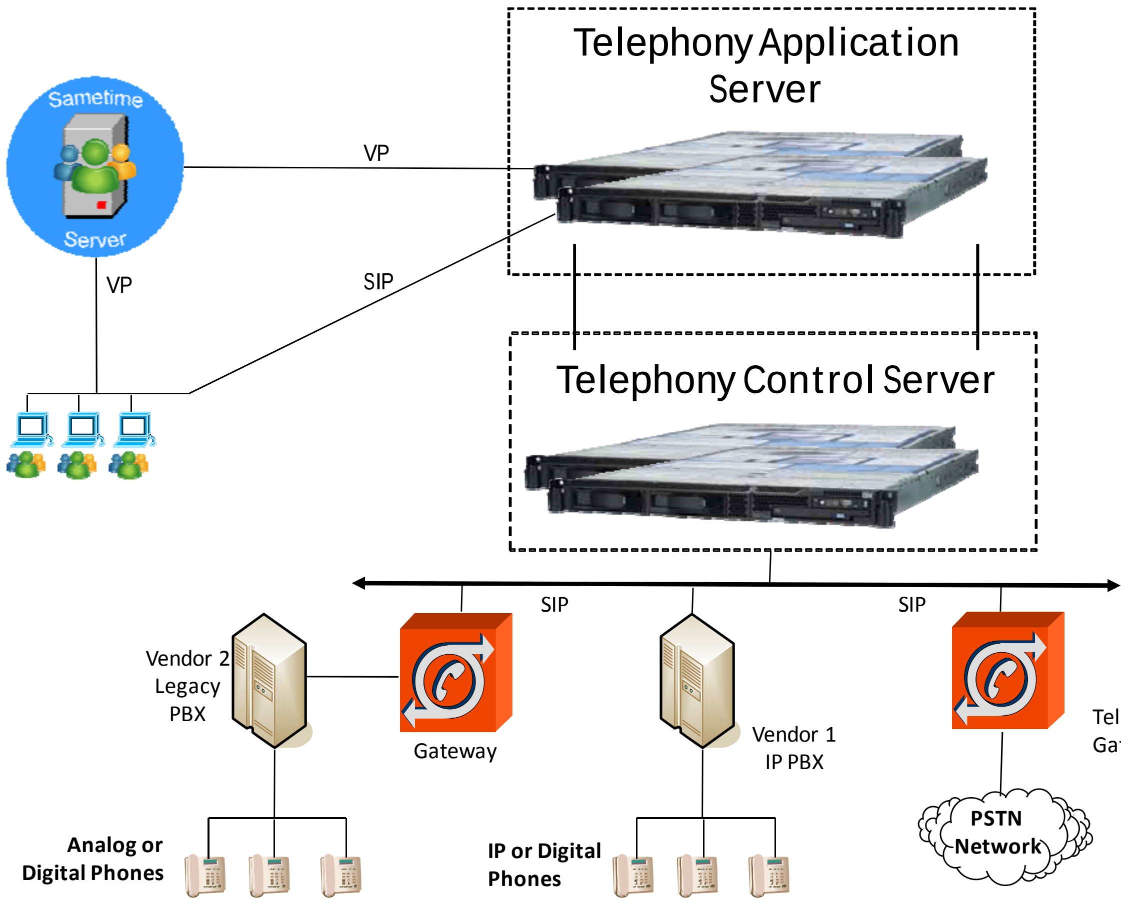

Figure 2: IBM Lotus Sametime Unified Telephony Architecture.

The Telephony Application Server (TAS) does all call routing for Sametime. It is a typical IT software application running on a "semi-appliance", and each server can handle 15,000 – 20,000 users. TAS serves are always deployed in an n+1 configuration for redundancy. These servers are clustered in the data center, not distributed to the network edge.

The Telephony Control Server (TCS) provides a PBX and gateway independent abstraction layer that allows TAS call routing instructions to be passed to nearly any PBX or gateway. This server has a SIP back-to-back-user-agent, so it supports all SIP PBXs and gateways, and it will support legacy PBXs via SIP through an appropriate gateway. Telephony Control

[2] Internally, Sametime uses Virtual Places as its session/connection protocol rather than SIP. This is historical as Sametime was developed before SIP was even a standard.

Servers are always deployed in pairs with pre-installed software on a hardened appliance. TCS servers support five 9's reliability, and pair can handle up to 100,000 users. TCS servers are located in the enterprise data center.

ANALYSIS

Several key differences between Microsoft's architecture and IBM's architecture immediately jump out:

1. Whereas Microsoft has placed call control capability within the Office Communications Server 2007 server pool, IBM has chosen to decouple call control from Sametime and run it on a special Telephony Application Server.
2. IBM has opted to develop the interfaces between its Telephony Control Server and all of the gateways and PBXs itself, while vendors working with Microsoft are required modify their PBXs and gateways to interface with the Mediation Server.
3. Microsoft has chosen to create a distributed model where Mediation Servers are located at the network edge near any exit to the PSTN. IBM has chosen a centralized model with TAS and TCS servers located in the data center.
4. IBM has built its servers with very high simultaneous call capacities that can interface directly with many PBXs and gateways while Microsoft has chosen to build low capacity Mediation Server devices that can interface with a single PBX or gateway.
5. Microsoft uses its high-quality, but proprietary, RT audio codec within OCS, which the Mediation Server must transcode to a standard codec when calling outside; the Lotus Sametime softphone uses a standard audio codec that requires no transcoding when calling outside.

There are a few differences in how the two products handle call flow. In locations where an enterprise has a PBX with PSTN trunk lines, OCS has developed an optional mechanism called dual-forking. As illustrated in figure 3, when an incoming call enters the system, the PBX rings the desk phone; simultaneously, it notifies OCS of an incoming call. OCS then also rings the Communicator client. Thus, both the phone and the Office Communicator client ring. If the desk phone is answered first, the PBX must signal OCS, and OCS turns off ringing to Office Communicator. The call is anchored in the PBX, and the PBX maintains call control. OCS polls the PBX using µaCSTA to get call status so that the phone presence displays correctly in Office Communicator.

If, on the other hand, the user would have answered the call with Office Communicator, OCS would signal the PBX, and the PBX would stop ringing the desk phone. In this case, the call would be anchored in OCS, and OCS would maintain call control.

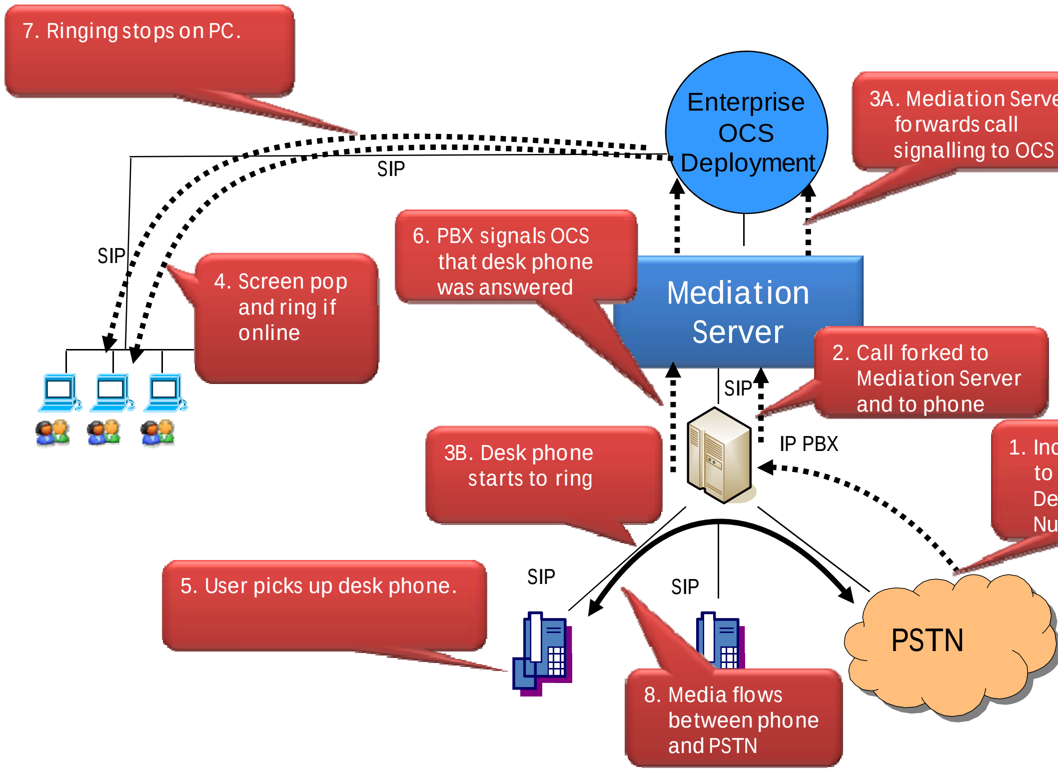

Figure 3: Microsoft OCS incoming call flow with dual forking[3].

Thus, there are two centers of call control with an OCS deployment: the PBX and OCS. Both need to continually communicate with each other. Microsoft uses µaCSTA to query the PBX for telephony device status (on-hook/off-hook). Because of this dual control situation, Microsoft has had to create a few SIP extensions (which the telephony vendors must implement) to avoid the phenomenon called hair-pinning or tromboning: a situation where a call is routed up through the PBX to OCS and because of user device preferences or call transfers, OCS routes it back through the PBX.

[3] This scenario is known as dual-forking in OCS. Sametime can ring multiple devices simultaneously as well, but the preferred model in Sametime is to choose which device to ring or signal based on presence and rules.

A typical incoming call flow for Sametime Unified Telephony is illustrated in figure 4.

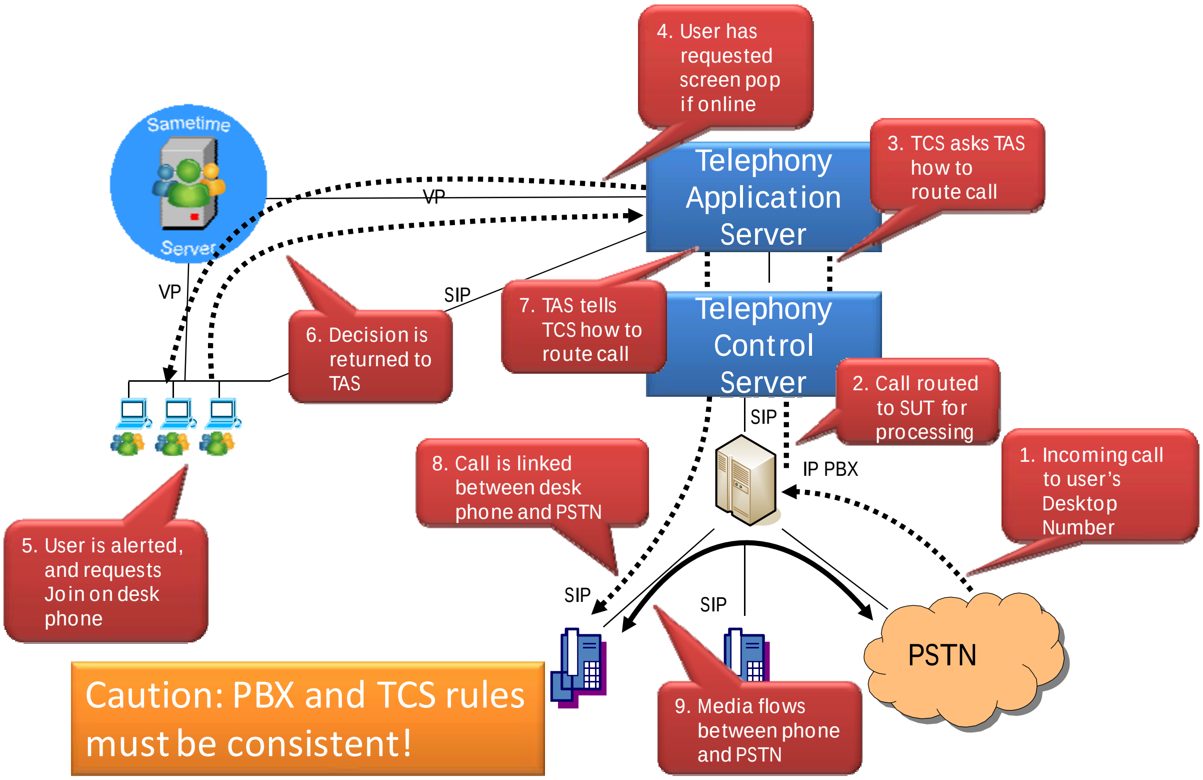

Figure 4: Sametime Unified Telephony incoming call flow.

In Sametime, all calls are always anchored in the Telephony Application Server. In order for Sametime to know a telephone's status (off-hook or on-hook), the PBX must forward all call control decisions up to the TAS (in the first release, Sametime Unified Telephony does not support µaCSTA for PBX telephone state queries). TAS can ring any device, either simultaneously or selectively based on rich presence and user rules. If a user's preference or selection is for the desk phone to ring, TAS instructs the PBX to ring the desk phone. If a user's presence indicates that the person is out of the office and not at home, then TAS will ring the person's cell phone, if the rules are so configured.

When a Sametime user dials from his phone handset, the PBX must send the call signaling to the TAS server in order for Sametime to know that the person is in a call. The TAS server will determine how to best connect the call based on the identity and profile of the person being called.

Another key difference between Sametime and OCS is the capacities of the Telephony Control Server in Sametime and the Mediation Server in OCS. This clearly influences server proliferation and may impact scaling. Comparing the multi-location deployment architectures for OCS and Sametime, clearly a Microsoft OCS solution will require more servers than Sametime when OCS is deployed in a large organization, and server propagation is exacerbated as more locations are added.

Consider, for example, a 10,000 person company with 100 branch locations, each with its own small- or medium-sized PBX, as in a bank with branch offices. Both the Sametime server cluster and the OCS Enterprise Server pool can easily handle this many users from an IM and presence standpoint: Sametime would require two servers (for Domino and Sametime Server) while OCS would require three (two for OCS and one for SQL Server). However, as we consider integrating these solutions with the telephony infrastructure, each branch office would require at least one and possibly two Mediation Servers in an OCS deployment – the number depends upon the call volume in each location. In a Sametime deployment, only two TAS servers and two TCS servers would be required. In a real world OCS solution, however, an organization would likely consolidate its PBX infrastructure before it would deploy so many servers, or it would drop the call off onto the PBX network and let the PBX network route the call as opposed to OCS.

For either Sametime or OCS, E-911 will be an issue. Both solutions support softphones, and it is not immediately clear how 911 calls would be handled for these users. There has been some discussion of using location services to help pinpoint a person's whereabouts, but these are presently primitive, and insufficient for E-911.

With OCS's proprietary voice model, Microsoft is hoping to grow a new business around devices that license Microsoft's voice codec and integrate directly with OCS. OCS does not support third-party SIP phones today, although Microsoft has allowed Tandberg and Polycom to integrate their SIP video endpoints with OCS. Although Sametime does not support any third-party SIP devices today, we predict that within a short period of time, the company will open up Sametime's TAS to registration by standards-based SIP phones, like those from Polycom and others. Thus, both companies could significantly spawn a movement away from traditional telephony devices and PBXs.

CONCLUSION

Both IBM Lotus and Microsoft have developed compelling call control capabilities in Sametime and OCS, respectively. Sametime Unified Telephony is targeted primarily toward the large distributed enterprise, and the product architecture is based on highly scalable and reliable centralized components. Because IBM has developed the interface with the PBXs and gateways, this solution tends to isolate the front end experience from back end telephony infrastructure decisions.

Microsoft OCS is targeted toward large, medium, and small businesses. Feedback from the field on deployed OCS solutions is very positive, and we are aware of several deployments of 10's of thousands OCS/Communicator users. The company has a proprietary voice model that may require a significant number of Mediation servers depending upon the size and geographical diversity of the enterprise's deployed PBXs. This solution also isolates the front end user experience from the back end telephony infrastructure decisions, with the caveat that the Mediation server requirement complicates an OCS deployment from a telephony integration point of view.

Both companies are likely to ultimately promote movement away from traditional PBX infrastructures, and both architectures can support telephony with no PBXs in the mix. Currently, Microsoft is farther down this path than is IBM Lotus.