Customers. Performance. Loyalty.

WHITEPAPER

Troubleshooting Enterprise VoIP

Moving from Despair to Repair

Using the Hammer Call Analyzer to Find and Repair Difficult IP Telephony Problems

Executive Overview

VoIP is here to stay, and nearly all enterprises are considering when and how to take advantage of this new technology. There are a variety of business motivations driving VoIP adoption in the enterprise, and there are real costs savings, productivity increases, and image benefits to be gained. These benefits are typically enabled by VoIP infrastructure on a converged network, but achieved through IP telephony applications such as messaging, conferencing, contact center, and mobility.

Realizing these benefits can be a challenge, and organizations may experience frustration, stress and even despair as they work to deploy VoIP in their environments. Simply put, VoIP is much more than just another application on the network, and most organizations have never managed an application with high availability and performance requirements like those of VoIP. With the right tools and training, however, organizations may become more skilled and responsive at troubleshooting and resolving issues — **repair**.

The best approach is to consider VoIP deployment as a lifecycle and **prepare**.

Careful planning throughout the entire lifecycle can help enterprises meet the challenges of VoIP deployment before they turn into quality of service issues. The best way to prepare for and combat the quality issues inherent in a converged network is to test the network thoroughly before rollout, and the best way to prepare for VoIP application rollout and troubleshooting is to baseline VoIP applications and then test across a range of conditions to anticipate and resolve problems before they arise in production.

This paper is Part Two of a three-part series on enterprise VoIP, which was developed to provide information about the factors affecting both the infrastructure and applications during VoIP deployments. It lays out a lifecycle approach for VoIP deployment and describes the multiple levels of measurement needed to manage VoIP applications. Tools and techniques are offered for baselining, proactive problem identification, and troubleshooting. Testing is not an option, it is essential, so specific examples of how testing addresses common VoIP issues are described to help enterprises confront these issues before they impact quality and erode the business benefits of VoIP.

The State of Enterprise VoIP

Part One of this series focused on the state of enterprise VoIP deployments, common pitfalls, and the benefits of taking a lifecycle approach to deployment. An overview of Part One is included here for background; for additional information, refer to Part One.

For enterprises, VoIP adoption has become a matter of when and how, not if. A recent survey showed 52 percent are already deploying VoIP at some level, 46 percent have plans for deployment, and only 2 percent have no plans to deploy VoIP. The compelling business benefits on which VoIP deployments are justified are evolving to be increasingly application-driven. They provide reduced operating costs, simpler network administration, improved team collaboration/productivity, and better communications with customers, partners and suppliers.

Like any major new technology, VoIP brings with it a unique set of challenges as well as opportunities. The concept of VoIP is so simple that it's common to underestimate just how complex the reality of VoIP is. VoIP is still an emerging technology, so there are a plethora of new standards, moving feature sets, and competing approaches, which makes interoperability a major challenge. Voice is a tricky medium for IP: it is very sensitive to the inherent characteristics of IP networks, including latency, jitter, and packet loss. Voice quality, availability, and security are critical, but they are difficult to measure and manage. (Quality also depends heavily on the characteristics of the phones, equipment and the incoming signal, since echo, distortion, and noise matter as much as packet-level impairments). Yet the expectations are high: five '9's availability, toll quality, clean interoperability and applications, and dial tone even when the lights go out. Most companies have never managed an application with such high availability and performance requirements.

Experiences of Despair

Pitfalls of IP Telephony

Since enterprise VoIP deployments first began in earnest around 2001, there have been a number of high-profile disasters. Dow Chemical, Carnegie Mellon, Merrill Lynch, and others began wide-scale deployments only to back out or change direction due to quality and acceptance issues. In cases like these, the cost of failed deployment can be millions of dollars — and this doesn't include the inevitable disruption to business operations. To call the feelings of those in this situation despair may be putting it mildly. Less visible but more common are projects which run into trouble and tough it through, but lose the business benefit along the way. Project managers and vendors usually don't want to talk about these cases of quiet desperation, but they are far too frequent.

For those who have successfully deployed VoIP and IP telephony applications, the process is generally a learning experience. The combination of new technology, new organizational structures, and operating at several layers simultaneously make each deployment project a challenge. Though each project has specific issues, there are a number of common pitfalls:

- Lack of organizational readiness
- Underestimating VoIP as just another application on the network
- Expecting telephony-grade performance without benchmarking
- Depending on a single-shot network assessment
- Lack of a lifecycle view
- Focus on infrastructure only rather than applications

Though this scenario — lots of complex pitfalls and challenges, lots of interdependent planning and testing required — may seem daunting, it is quite possible to have a clean deployment and realize significant business benefits. It is important to understand that there is no complete process available today, you should seek out the best practices available, and to equip your team with the best available knowledge and tools.

Moving to Repair

For those enterprises that have experienced despair in their deployment process, the way out is clear: fix the current problems and plan the next phases with more wisdom. In most cases, this means equipping the team with better knowledge and tools.

Multiple Levels of Measurement

Measurement is crucial for improvement; it's an old adage that you can't manage what you can't measure. At any stage in a deployment it is important to capture baselines and to have the capability to measure at multiple levels. Table 1 shows a view of the levels involved in a VoIP deployment, some important measurements, and some useful tools.

Table 1: VoIP Management and Tools

LAYER	MEASUREMENTS	TOOLS
• Packet Layer	• Packet loss • Jitter • Delay	• Packet Sniffers and Testers • Element-Level Network Management • Switch/router Consoles/statistics
• Call Layer	• Call Completion Rates and Latency • Voice Quality • Equipment Status	• Call and Protocol Analyzers • Infrastructure Testers • Equipment/Server Consoles
• Application Layer (format)	• Application Completion Rates and Latency • Cross-Platform Status • Provider SLAs	• Application Testers • Transaction-Level Management • SLA Management

By measuring the operations of currently deployed IP telephony applications regularly, IT staff can gain a good background for troubleshooting, problem resolution, growth, and further deployments. A positive side effect of a measurement program is that staff become exposed and trained in how the infrastructure and applications work.

Troubleshooting Tools and Techniques

Resolving VoIP problems, especially intermittent ones, is very difficult. A skilled troubleshooter needs a variety of tools at his or her disposal. One indispensable tool is an analyzer capable of tracing and troubleshooting VoIP at the call level. In exploring how this kind of tool can aid in problem troubleshooting and resolution, we'll use some specific examples with regard to a specific tool.

The Hammer Call Analyzer™

The Hammer Call Analyzer is a tool that enables network engineers and administrators to visualize signaling and voice quality problems in VoIP networks. Figure 1 shows some of the user interface screens from the Hammer Call Analyzer. It allows network engineers to select an individual message from any call, and instantly see all related messages from the same call even if the trace contains hundreds of calls. Organizing all the packets automatically results in huge time savings for the troubleshooting engineer. In addition, the Hammer Call Analyzer displays waveforms and the Stream Quality Signature (an enhanced display of jitter) for any call. These features allow engineers to visualize problems in the exchange of messages between the various devices and to solve them quickly.

Figure 1: Hammer Call Analyzer screens.

Scenario A

Call Completion Failures

In many deployment scenarios, there are situations with calls that don't complete: the user may complain about not reaching a given destination, or may retry and find an alternate route, leaving the problem to be encountered by the next user.

Common causes of this kind of symptom are:

- Gateway and switch misconfiguration
- Interoperability issues between equipment
- Capacity limitations
- Performance issues and delays triggering timeouts
- Feature interaction issues such as conflicting call-forwarding settings

Capturing a trace of a failed call is invaluable. A trace contains all of the packets, both media and signaling, that the analyzer saw during a capture session — for example, all the packets for all calls during a specific time period. With the Hammer Call Analyzer, a trace can be easily interpreted and failure codes decoded for fast problem resolution. An example screen of a call capture in real time is shown in Figure 2. This is displayed in a ladder diagram format, which allows visualization of scenarios with many calls, and quick drill-down to trace into a specific call with both signaling and media. Since it is a VoIP-specific tool, the Hammer Call Analyzer can provide insight based on a deep understanding of all the protocols involved with VoIP.

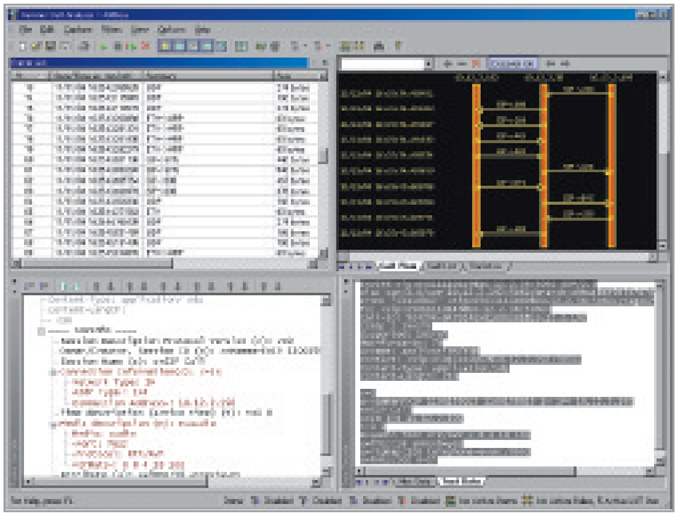

Figure 2: Hammer Call Analyzer capture screens

If problems are intermittent, trigger and search capabilities allow capture of elusive failures. A common example of this is to set a trigger on call failure codes, which are usually associated with a call setup failure. Triggers allow capture to stop, and can be set so the surrounding context is collected along with the specific event. This can include pre-trigger and after-trigger information, as shown in Figure 3.

An example trigger setup is shown in Figure 4. A similar capability allows search within any captured set of traces. A troubleshooter can then examine factors like:

- Are calls around the problem call experiencing similar problems?
- Are there significant packet-level issues? Are messages being retransmitted at timeout thresholds?
- How many calls per second are coming in? Is this near the capacity of the proxy server or gateway?
- Are there specific tags or information elements associated with problem calls?

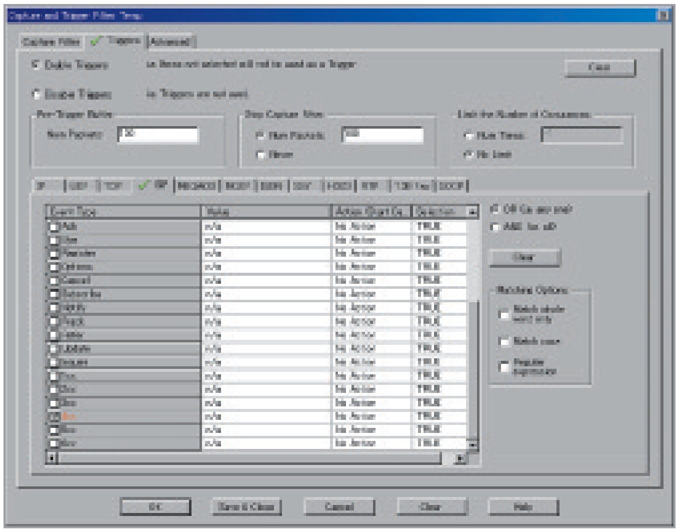

Figure 4: Example trigger setup.

Scenario B

Routing Issues Across Domain Borders

Using Session Border Controllers (SBC) is de facto practice with most VoIP service providers, and is the increasingly common within enterprises. They provide the ability to isolate different parts of the VoIP network and help ensure security and reliability of each domain. For enterprises considering outsourcing part of their IP telephony applications (such as IVR, ACD, or even parts of the IPT infrastructure), tracking calls across session borders is a nightmare. Call legs in each domain have completely different IDs, as shown in Figure 5.

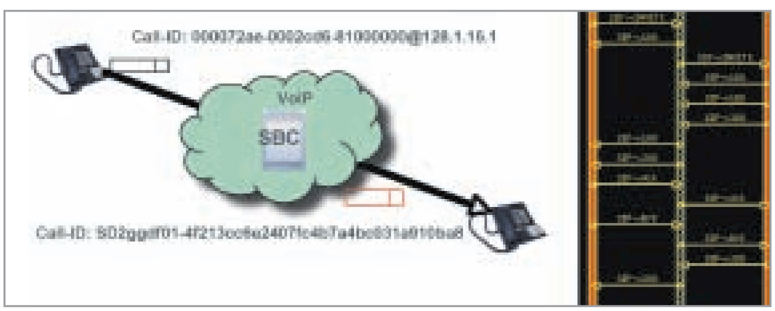

Figure 5: Tracing calls across session borders.

Tracing calls across these borders is essential for troubleshooting across different service providers, and it requires tools that are protocol-aware. With the Hammer Call Analyzer, the Call Merge Map and Name Map are available for this situation and allow a troubleshooter to see how calls traverse network boundaries. Even though the Call ID changes across session boundaries, calls can still be shown easily on the same ladder diagram.

Scenario C

Voice Quality Issues

Complaints about voice quality are dreaded by many IT departments that have deployed VoIP, because the causes are varied, the problems are commonly intermittent, and the users' perceptions are very subjective. Voice quality measurements are themselves complex. (See Voice Quality Measurements, page 9).

With the Hammer Call Analyzer, voice quality measurements are readily available. These measurements include details of packet-level statistics including a distribution of jitter (which is too often regarded as a single number). A classic issue is a misconfiguration on only one of the multiple routes taken by a stream. This exhibits itself as a bimodal jitter distribution, as shown in the lower right of Figure 6.

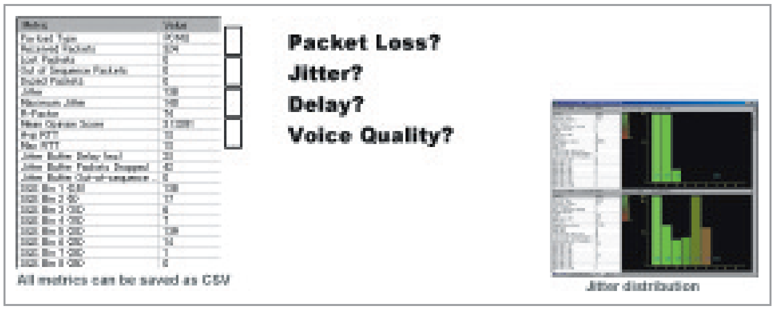

Figure 6: Example jitter distribution and MOS measurement.

By drilling deeper into the media stream, examining the jitter and other metrics of a specific call over time, and listening to the media, a troubleshooter can resolve a variety of common issues. For example, Figure 7 below shows a case where voice packets are dropping because jitter increases suddenly at one point. It is simple to use the Hammer Call Analyzer to measure and determine that this affects voice quality significantly. By checking the other network traffic from this set of captured data, the troubleshooter found that the user was downloading large files at this point, which impacted voice packets because QoS parameters were set incorrectly.

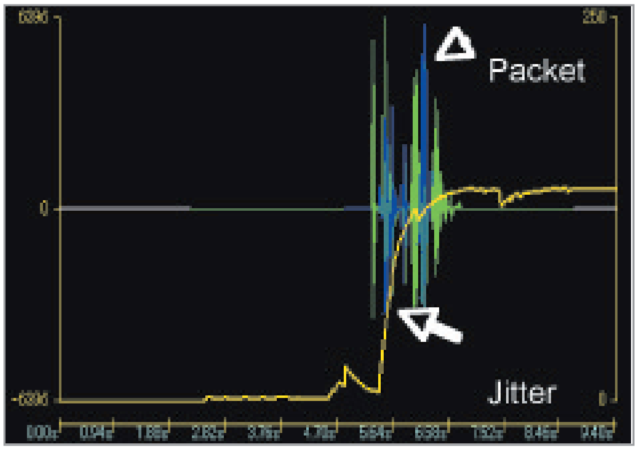

Figure 7: Example with sudden increase in jitter.

A crucial but commonly overlooked voice quality issue is echo. Echo does not show up on any packet-level measurement; instead cal-level voice quality measurements like TELR (Talker Echo Loudness Rating — ITU Recommendation G.122) are necessary. Figure 8 shows an example of a TELR measurement using the Hammer Call Analyzer.

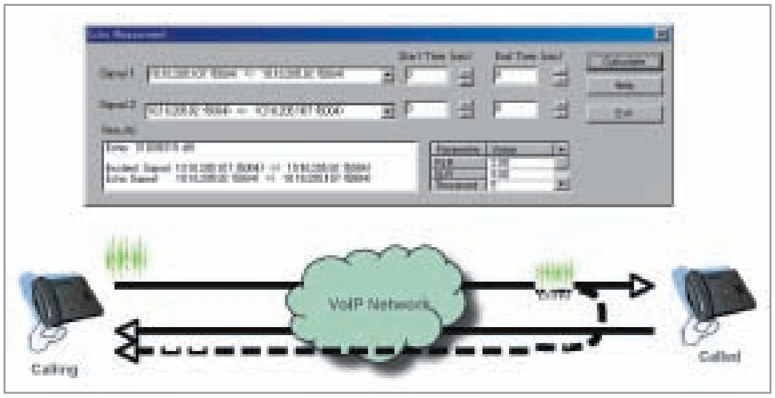

Figure 8: Use of TELR to measure echo.

Since the biggest determiner of voice quality is the jitter buffer found in the VoIP telephone, it is important to be able to evaluate how effectively the jitter buffer is in any given environment. The Hammer Call Analyzer allows a troubleshooter to try different parameters and evaluate the resulting voice quality metric (as well as listening to the call). See Figure 9 for an example.

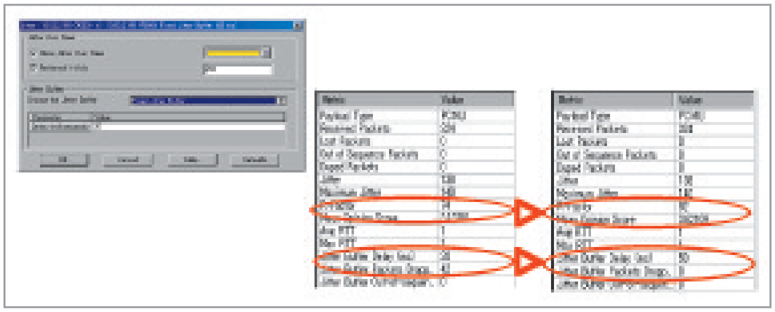

Figure 9: Determining the quality impact of changing jitter buffer size

Voice Quality Measurements

Unfortunately, there is no perfect measure for voice quality, despite years of work by thousands of scientists. The ideal measure is a single number, which is simple, objective, and can be applied on any call including live calls. The state of the art today requires multiple kinds of voice quality measurements.

PESQ

PESQ (Perceptual Evaluation of Speech Quality — ITU-T Recommendation P.862) is the best measurement available. PESQ can provide a rapid and repeatable result in a few moments, offering an objective measurement tool that predicts the results of subjective listening tests on telephony systems. PESQ uses a sensory model to compare the original, unprocessed signal with the degraded signal from the network or network element. The resulting quality score is analogous to the subjective Mean Opinion Score (MOS) scale traditionally used to rate telephone speech quality. A MOS score ranges from 1 for an unacceptable call to 5 for an excellent call, with 4 being toll quality. A typical range for Voice over IP would be from 3.5 to 4.2.

The advantage of the PESQ approach is that it measures the effects of many different impairments and their interactions. Unfortunately, PESQ measurements are not available on live calls; they must use test calls with prerecorded audio to measure the audio characteristics of a given call.

Passive Monitoring

Non-intrusive or passive Monitoring examines a stream of voice traffic and produces a transmission quality metric that can be used to estimate a MOS score. This has the advantage that all calls in a network can be monitored without any additional network overhead but the disadvantage that the effects of some impairments are not incorporated.

Passive measurements such as the industry-standard ITU G.107 E Model, the endpoint-based XRTP-XR, and Telchemy's VQmon provide quality monitoring through observation of the RTP stream and incorporate effects such as packet loss burstiness. This produces an R Factor which can be used to estimate a MOS score. G.107 is embedded into many VoIP gateways and other end systems, but must be queried and interpreted appropriately.

P.563 is a new passive monitoring algorithm that analyzes the voice stream in order to estimate call-quality scores. It is much more computationally complex (over 1000x) than VQmon and recently standardize, and hence is not found within any current network equipment.

Empirix's tools include the best available voice quality measurements, including PESQ, PAMS (an older version of PESQ), and VQmon, as well as echo measurement TELR (Talker Echo Loudness Rating — ITU Recommendation G.122) and many other metrics.

Scenario D

Pinpointing Security Holes

VoIP security is a deep subject, and it is still evolving. Nevertheless, there are a number of common security issues that can be found via inspection. These include:

- Use of global IP addresses across session borders (causing internal addresses to be visible even if this is not desired; for example, creating addresses for DOS attacks).
- Improper security settings in gateways (causing media to be sent in clear text when it is meant to be encrypted).
- Firewall settings allowing probes to reach telephony servers (many of which have insecure administrative logins).
- Malformed signaling messages, especially those designed to cause buffer overflows in telephony servers and other equipment (akin to the worms that infect Windows machines).

Clearly it is better to prepare for VoIP security, but it is also important to be able to troubleshoot and repair these security problems. Figure 10 below shows a scenario captured with the Hammer Call Analyzer which broadcasts a global IP address through a session border controller, causing a security hole.

Figure 10: Example call with security issue.

Testing Tools as a Means to Repair

Testing tools are another invaluable addition to the VoIP deployer's arsenal. While we will cover useful testing techniques in another paper, it is important to remember that it is never too late to test. Many deployment projects have begun without baselining or network testing ahead of time, only to discover the need and do it mid-project.

The baseline scenarios that can be captured through automated testing of VoIP infrastructure and applications are invaluable. They create a captured record of normal operation for a given environment, help find remaining issues before users do, and are one of the best ways to train IT staff in VoIP operations and troubleshooting.

Tools as Part of Training

Learning VoIP on the Job

VoIP is new for both the data and voice people in an IT organization; it requires skill and commitment from both communities and can provide a growth opportunity for everyone. Training is critical for everyone involved. There are numerous good sources for training on VoIP infrastructure and IP telephony applications. This training can be greatly amplified by availability of VoIP-specific tools ahead of time.

Simply tracing real calls, watching the signaling, and measuring the media, quickly teaches network engineers VoIP protocols, networking, and measurement. By benchmarking and testing existing voice applications (either prior to or during a VoIP migration project) is a great way to learn how these applications are put together and how they are experienced by the user community.

Prepare and Avoid Despair

Being Ready for Deployment

Of course, if one can avoid issues in VoIP deployment, one should. The vast majority of issues are in fact avoidable using proper planning, training, processes, and tools. No amount of preparation will provide a 100 percent trouble-free rollout, so being ready to repair is important as well. However, if you view deployment as a lifecycle with distinct phases and are ready for each phase ahead of time, deployments can be done smoothly and significant amounts of time and money saved.

VoIP-specific testing is a critical, often overlooked, tool to check a variety of key factors, including:

- Ability and expertise to configure the network properly
- Interoperability issues
- Capacity
- Management tools

Part Three of this series discusses testing techniques in more depth, and describes specific scenarios and techniques for benchmarking, assessment, and testing using a specific tool — the Hammer VoIP Test Solution for Enterprises — as an example.

Network architects and managers should address call quality and performance management problems when they plan and deploy their IP networks, but they should be aware that these problems also frequently occur during normal day-to-day network operation. Many VoIP-related problems are transient in nature and can occur at many places along the network path. The techniques covered above for troubleshooting will be extremely useful even for the best-prepared organizations.

Conclusion

Enterprise VoIP deployments are complex and tricky projects which can reap substantial business benefits. Too many organizations have been caught in the despair of projects running into major trouble midstream without sufficient planning, training, or tools. This situation is recoverable; more importantly it is avoidable. Some simple principles, tools, and processes can be used to reduce headaches and improve business benefits wherever you are in the deployment lifecycle.

Prepare. The ideal situation is to prepare ahead of time. This includes organization, training, tools, and planning. Benchmarking very early, with network testing and current application measurement, is very powerful in guiding plans, and lays the groundwork for optimization after cutover. It can also be a great training experience. Planning ahead for all phases of deployment may seem overwhelming, but it pays off. Testing is an essential process at all stages of deployment; it is not an option.

Repair. Organizations will inevitably find themselves needing to repair. (Those who are well prepared have much less repair work and will be more effective when they do have it because the tools, training, and processes are in place.) Troubleshooting is difficult with VoIP, and VoIP-specific tools are indispensable. Testing plays a major role here as well, both for isolating and recreating problems and for determining how critical they are through regular benchmarking.

Despair. Too many organizations find themselves in despair during their VoIP deployment projects. This may sound dramatic, but quiet desperation is widespread. For organizations in despair, there is hope. First, get good at repair, and stop the bleeding that happens daily in these situations. Second, go back and prepare. This will make ongoing operations manageable and make the next deployment phase a positive experience.

The business benefits of VoIP are real, and the adoption of this technology is rapid and widespread. Realizing these benefits can be a challenge, however, and organizations may find themselves in despair too frequently and too easily. By being aware of common pitfalls, viewing VoIP deployment as a lifecycle, and using the best tools and practices, deployments can go smoothly and deliver significant business benefits.

WP-1-TRBLSHTEVOIP04.05

Corporate HQ | 20 Crosby Drive | Bedford, MA 01730 | USA | tel: +1 781.266.3200 | fax: +1 781.266.3201 | www.empirix.com
European HQ | Building A, Trinity Court, Wokingham Road | Bracknell, Berkshire RG42 1PL | United Kingdom | tel: +44 (0)1344 66 8080 | fax: +44 (0)1344 66 8172