VIRTUAL DESKTOP PERFORMANCE AND QUALITY OF EXPERIENCE

UNDERSTANDING THE IMPORTANCE OF A DISTRIBUTED DATA CENTER ARCHITECTURE

EXECUTIVE SUMMARY

Cloud services, such as virtual desktop infrastructure (VDI), bring enterprises efficiencies and cost savings. However, before moving to VDI, enterprises must understand the network conditions required to ensure end-user quality of experience (QoE) for real-time applications. They must also understand the relationship between VDI performance and data center architecture. This knowledge will help enterprises choose a cloud service provider that can meet their performance and QoE requirements.

To win enterprise customers, cloud service providers must also understand how data center location affects service performance and QoE. This understanding will help ensure that they build a cloud solution that can deliver QoE their customers are willing to pay for. It will also help them educate customers about the benefits of a distributed data center architecture compared to the centralized architecture offered by their competition.

TABLE OF CONTENTS

INSIGHTS AT A GLANCE

- In a virtual desktop infrastructure (VDI) service, different office suite tasks have different levels of sensitivity to network impairments. Those that are most bandwidth-intensive are most sensitive to decreases in bandwidth and to increases in latency and packet loss.
- Office suite activities that require more bandwidth tend to be more highly correlated to end-user quality of experience (QoE) than those that require low bandwidth.
- End-user QoE for VDI office suite activities is highly correlated to changes in latency and bandwidth. Users quickly become dissatisfied as network conditions degrade.
- End-user QoE is more closely tied to the effects of latency than to the effects of bandwidth. This means that simply increasing bandwidth is not enough to ensure end-user QoE. Latency must be kept low.
- End users' willingness to pay drops dramatically as bandwidth drops. Even at 6 Mb/s, users feel that the VDI service offers less than half the value of a desktop service.
- A VDI service that offers an experience that is almost indistinguishable from a desktop experience for office suite applications requires 6 Mb/s or more of bandwidth and less than 20 milliseconds (ms) of round-trip latency.
- A VDI service that has minimal effect on office suite performance and is deemed acceptable by most users requires 6 Mb/s of guaranteed bandwidth and no more than 40 ms of round-trip latency.
- Although VDI applications will operate under minimal network conditions, end-user QoE and performance are negatively impacted once round-trip latency exceeds 80 ms or packet loss reaches 10^{-2} (1 percent).

RATING VIRTUAL DESKTOP QOE

To help enterprises understand the relationships between VDI performance, end-user QoE and the effects of data center location, Alcatel-Lucent performed a series of subjective tests. During these tests, 30 individuals, all of whom were very comfortable using computers and office suite applications, executed a variety of typical knowledge worker related activities. Researchers chose activities that were representative of routine office work as well as activities that might be more network dependent. The list of activities included performing tasks with Microsoft® PowerPoint®, Word® and Internet Explorer®, and watching streamed video. Links to more information on these tests are provided in the Resources section, below.

Each participant executed the activities on a laptop where the office suite (referred to as "productivity suite" in the tests) applications and video were installed locally on the hard drive. They then executed the same activities on virtualized machines across a network that simulated the cloud. Network conditions — bandwidth, latency and packet loss — were varied to represent both distributed and centralized data center architectures.

Test participants were asked to rate their QoE for the office suite activities according to the Mean Opinion Score (MOS) ratings listed in Table 2[1] [1-3]. They provided a separate rating for video QoE. They had no knowledge of the differing network conditions that were applied for each set of tests.

Table 2. Test participants rated QoE according to a Mean Opinion Score

Rating	Opinion	Description
5	Perfect	Just like using applications on a laptop.
4	Good	I can perceive imperfections but still work effectively under these conditions.
3	Fair	Affects my work day.
2	Annoying	Nearly impossible to work effectively.
1	Impossible	Impossible to work under these conditions.

In addition to providing a MOS rating for each test condition, participants were asked to place a dollar value on the combined video and productivity suite experience. Their instructions were to assume that a laptop and software would cost 30 United States dollars per month, then rate the value of the virtual machine based on their recent video and productivity suite experience.

[1] The methodology for subjective rating follows the spirit of the ITU-T methods for subjective evaluation of audio and video quality using Mean Opinion Score (MOS). Currently, no standard exists for subjective scoring of application performance.

FINDING BANDWIDTH BARRIERS

Tests that varied the amount of allowable peak bandwidth (i.e. ceiling) showed that raising the ceiling to more than 6 Mb/s of bandwidth does not improve QoE for the productivity suite activities (the grey curve in Figure 1). However:

- Limiting bursts up to 3 Mb/s makes imperfections more noticeable and reduces QoE.
- Limiting bursts up to 1.5 Mb/s causes users to express more widely distributed opinions about QoE.

Video MOS (the blue curve in Figure 1) drops more rapidly than productivity suite MOS as bandwidth falls. The green graphs at the bottom of each column in Figure 1 show a histogram of individual MOS ratings. At a ceiling of 6 Mb/s, 23 percent of test participants indicated that video performance impacted their effectiveness or productivity.

These results were achieved with a round-trip-time (RTT) of 20 ms, which is the delay that can be expected from a locally situated data center.

Figure 1. 6 Mb/s of bandwidth provides very good QoE for productivity suite tasks

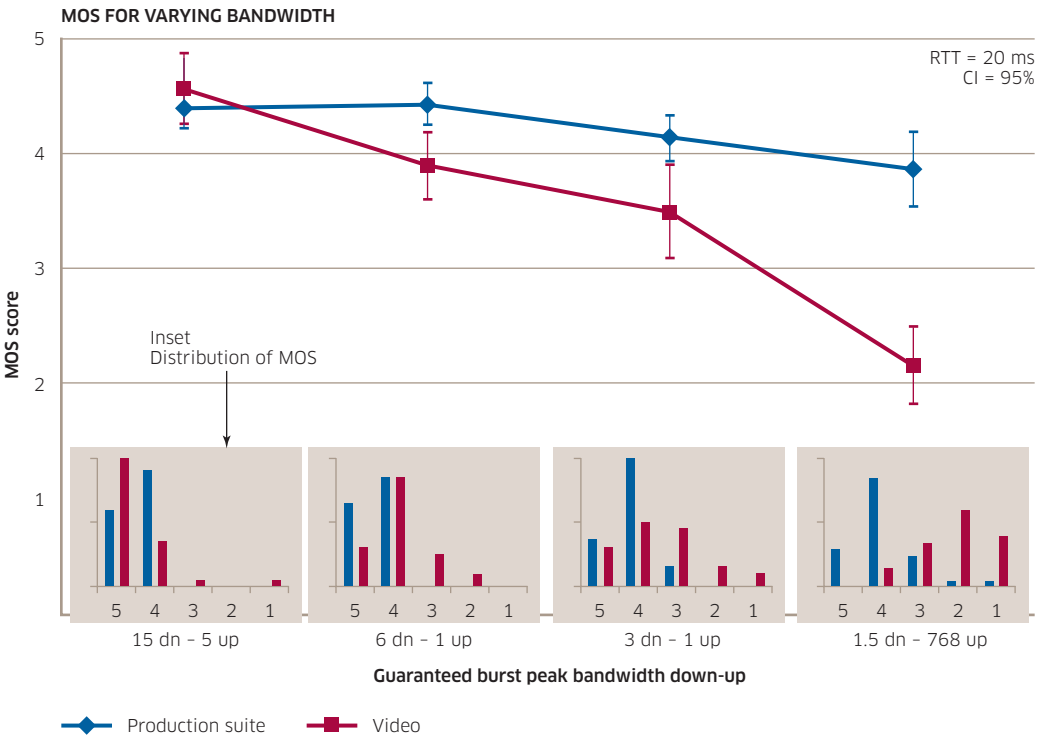

Users' willingness to pay for a VDI service, such as the one tested, drops dramatically as bandwidth drops (Figure 2). Willingness to pay appears to be correlated to the video curve in Figure 1, indicating that users more heavily weighted the video experience.

Figure 1 shows that the video MOS rating was consistently lower than the productivity suite MOS for the same network conditions. This affected the value score. Figure 2 shows the histogram of the users' value scoring. Here we see a much more varied set of opinions compared to the MOS ratings:

- At limits of 3 Mb/s and 6 Mb/s user scores vary nearly evenly from 0 to 30 dollars.
- At a limit of 15 Mb/s the average MOS rating for productivity and video is about 4.5, or 10 percent below perfect, while willingness to pay drops 33 percent from the best case.

As a result, we can infer that, despite implying that the experience would not affect their work day, users are still hesitant to pay for the equivalent of a standalone PC (30 United States dollars per month). Even if they don't affect a user's work day, noticeable performance differences make the service feel inferior and drive down the user's willingness to pay. This leads us to an important conclusion: that increased network performance leads to higher MOS and, more importantly, increased revenue.

Figure 2. The monetary value of the VDI service drops as bandwidth drops

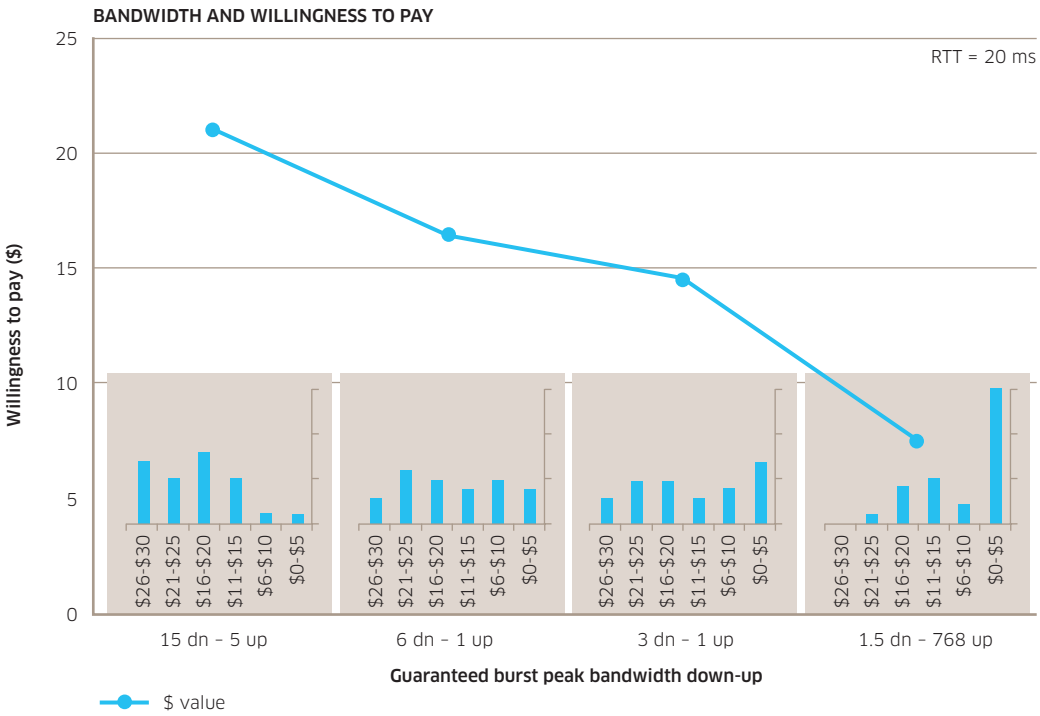

TESTING LATENCY LIMITS

Tests that varied the amount of latency showed that delays of more than 40 ms create a perceptible difference in QoE for the productivity suite (the blue curve in Figure 3). Latency for centralized data centers varies depending on the user's location and the number of service providers between the user and the serving data center. Latencies of 40ms to 120ms are not uncommon when the connection to the serving data center traverses regions. In the green graphs in Figure 3, note that:

- At 20 ms of latency no test participants felt their productivity was affected.
- At 80 ms of latency about 20 percent of test participants felt their productivity was affected.
- At 120 ms of latency more than half of the participants felt their productivity was affected.

Figure 3[2] shows the curve for a 5-inch x 3-inch video frame (the red line) and for an 8-inch x 5-inch video frame (the green line). Larger video frames require additional bandwidth so are more demanding on the network. Note that:

- QoE for the 5-inch x 3-inch video frame fell below a MOS rating of 4 at 40 ms and quickly declined as latency increased.
- QoE for the 8-inch x 5-inch video frame was marginal even with minimal delay, and consistently remained below the 5-inch x 3-inch video MOS ratings as latency increased beyond 40 ms.

Figure 3. MOS ratings drop as network latency increases

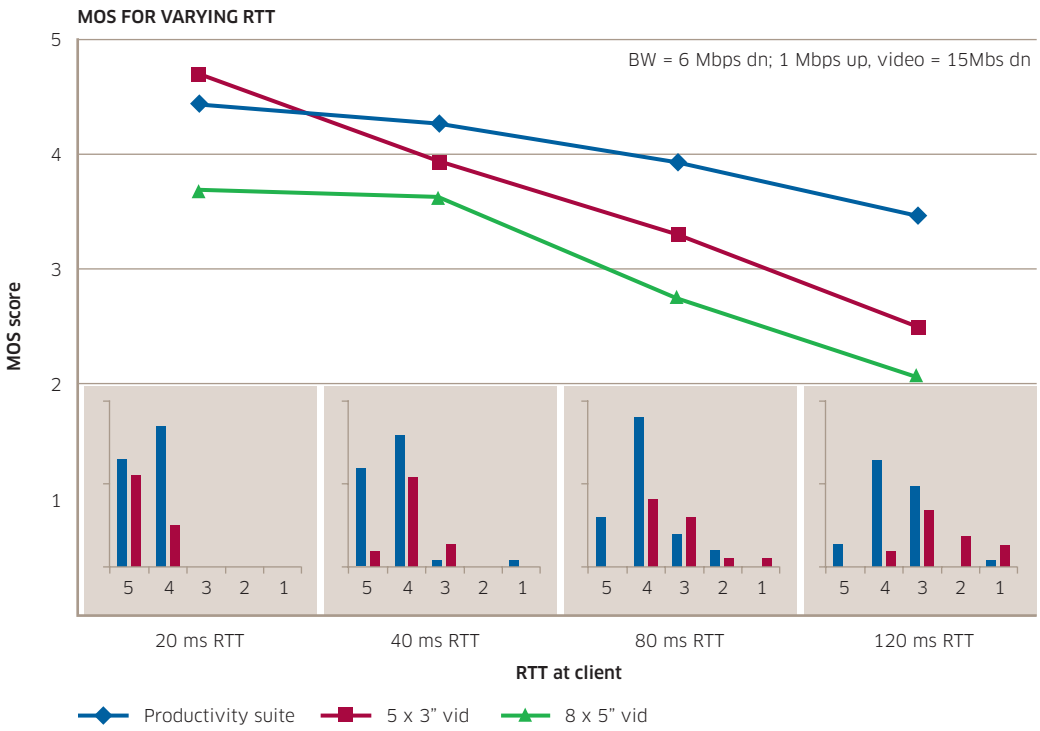

We can conclude that, compared to productivity applications, video applications put more demand on the network to achieve the same QoE. In addition, we note that larger video frames produce a lower QoE than smaller frames when used under the same network conditions.

[2] The red columns on the bar chart reflect the user ratings for the 5-inch x 3-inch video. The ratings for the 8-inch x 5-inch video were omitted for clarity.

ENGINEERING PACKET LOSS FOR QOE

Tests that varied the amount of packet loss showed that packet loss in the range of 10^{-2}, or 1 percent, affects bandwidth and QoE. As packet loss increased to 1 percent (not uncommon on the Internet), MOS dropped from 4.27 to 3.8 — about a half-point drop in QoE. At a 1 percent packet loss, many more users felt the experience would impact their ability to work.

Because the amount of packet loss in the network can be engineered by adapting the amount of bandwidth allocated during busy periods, service providers should tune their network to ensure packet loss does not reach unacceptable levels.

CORRELATING MOS, BANDWIDTH AND LATENCY

A look at average bandwidth peaks for individual productivity suite activities during the latency tests and the corresponding MOS ratings revealed that bandwidth curves fall in a consistent slope as latency round-trip increases from 20 ms to 120 ms (Figure 4). Similarly, the MOS curve tracks closely to the bandwidth curves. This implies that users can sense the drop in peak bandwidth and perceive it as lower QoE. This is reflected in lower MOS ratings.

Figure 4. MOS falls as latency increases

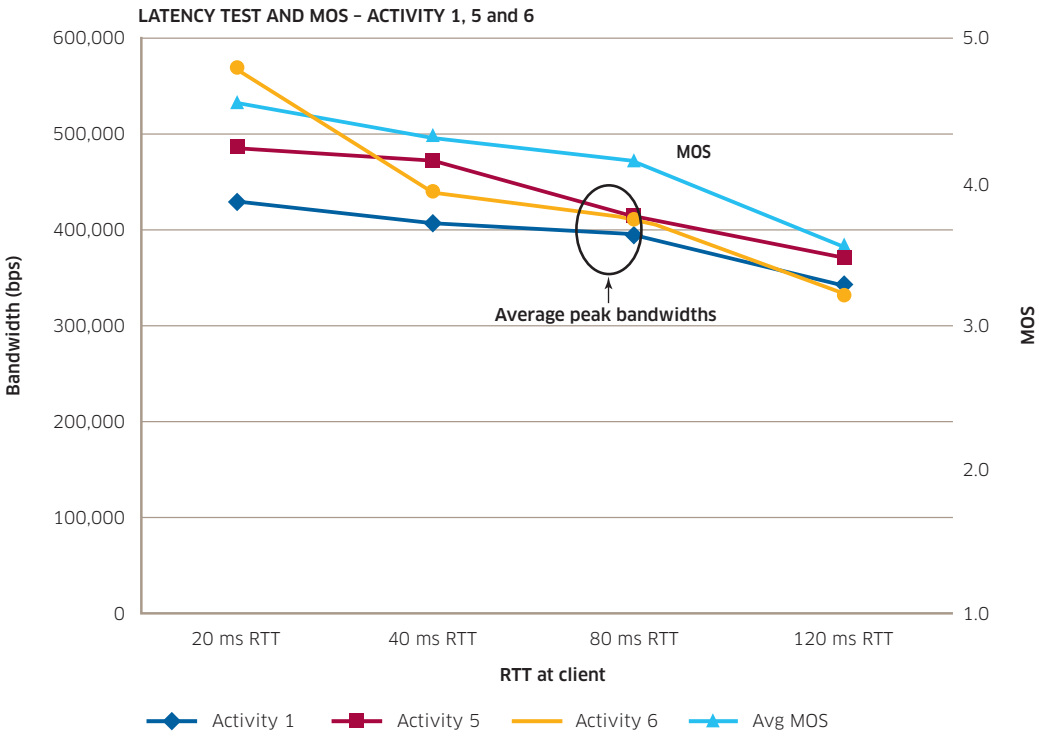

Calculations to determine the correlation[3] between the peak bandwidth curves and the MOS curve showed that higher bandwidth activities are closely correlated with the MOS curve while lower bandwidth activities are not (see examples in Table 3). This implies that high-bandwidth activities are more dependent on network performance for higher QoE.

Table 3. High-bandwidth activities are more closely correlated to MOS than lower-bandwidth activities

Activity	Description	Maximum bandwidth required	MOS Correlation[4] to latency
1	View a presentation slide show.	330 Kbps	1.0
5	Scroll a web page.	454 Kbps	0.96
3	Resize a presentation window.	3.5 Mbps	0.94
6	Copy and paste from a web page to a document.	624 Kbps	0.93
2	Edit a presentation slide.	52 Kbps	0.58

Comparing the correlation of the latency tests to MOS with the correlation of the bandwidth tests to MOS showed that changes in latency are more closely tied to end-user QoE than changes in bandwidth (Figure 5). This means that end users experience the effects of delay more than changes in bandwidth.

Figure 5. Latency is more closely correlated to MOS than bandwidth

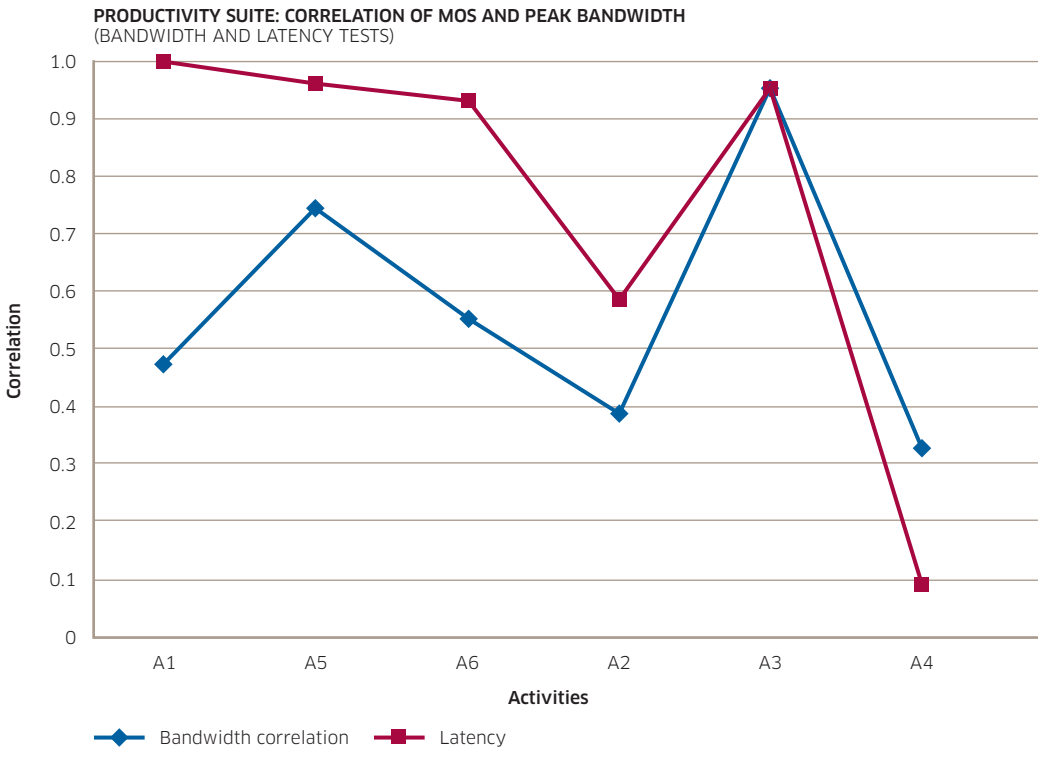

[3] Correlation is defined as $c = \frac{E((Y-\mu_Y)(X-\mu_X))}{\sigma_X \sigma_Y}$

[4] Two data sets that rise and fall exactly in proportion are perfectly correlated and have a correlation of 1. Two data sets that are not correlated at all have a correlation value of 0.

OTHER CONSIDERATIONS

This study describes some of the considerations that impact user QoE and willingness to pay. However, network-side and IT-related factors not considered in this study can also impact QoE. On the network side, these factors include network jitter, choice of VDI transport layer protocols (RDP, RGS or PCoIP™) and dynamic behavior of multiple desktops on the same network. On the IT side, these factors include Virtual Machine (VM) configuration (including CPU and memory allocation), storage performance in case of desktop boot storms and data center network performance. Specific desktop user workload and session profiles may also differ from those considered in this study.

Alcatel-Lucent welcomes opportunities to work with service providers on custom studies and assessments of specific network and IT environments and virtual desktop use cases. The aim of these studies is to enable service providers to deliver the best possible virtual desktop experience to end users.

DELIVERING REAL-TIME CLOUD SERVICES WITH QoE

With its close relationship to QoE, low latency (40 ms or less) can be a significant differentiator in VDI offerings. Low latency requires a distributed data center architecture that brings the cloud closer to end users.

Communications service providers are in an ideal position to build distributed data centers. They already have:

- A huge natural footprint. By adding information technology (IT) infrastructure components to their distributed central offices and managing the network and the cloud as a whole, service providers can offer a new class of “carrier cloud” services that ensure QoE for real-time applications.
- Trusted relationships with enterprises. A 2011 Alcatel-Lucent survey of more than 3800 IT decision makers found that trust is one of the most important purchasing factors for cloud services.

Alcatel-Lucent partners with service providers to help them develop a cloud services strategy based on market knowledge and to build a carrier cloud that delivers differentiated cloud services.

RESOURCES

To learn more about the carrier cloud and the benefits of a distributed data center architecture, please contact your local sales team or visit **www.alcatel-lucent.com/cloud**.

For information on IT decision maker survey results and willingness to pay per cloud feature, please contact the Alcatel-Lucent Market and Consumer Insight team at **mcinsight@alcatel-lucent.com**.

For details on tests conducted for the virtual desktop QoE study, please see Contacts section below.

REFERENCES

[1] ITU-T "P-800 Methods for subjective determination of transmission quality." 1996.

[2] ITU-T "P.910 Subjective video quality assessment methods for multimedia applications." 2008.

[3] ITU-T "P.920 – Interactive test methods for audiovisual communications." 2000.

[4] Data Center Modeling – Bell Labs Internal 2011.

[5] "Sunray Bandwidth vs. Citrix", https://blogs.oracle.com/ThinGuy/entry/sun_ray_bandwidth, 2005

CONTACTS

Ramesh Nagarajan, **ramesh.nagarajan@alcatel-lucent.com**
RJ Vale, **rj.vale@alcatel-lucent.com**

ABBREVIATIONS

IT	Information Technology
ms	Millisecond
MOS	Mean Opinion Score
PCoIP	PC-over-IP
QoE	Quality of Experience
RDP	Remote Desktop Protocol
RGS	Remote Graphics Software
RTT	Round-trip-time
VDI	Virtual Desktop Infrastructure

Alcatel·Lucent