riverbed

Think fast.™

WHITE PAPER

Riverbed Response to Miercom 090815B

Cisco and Miercom attack Riverbed again (Confusion runs rampant)

HERE WE GO AGAIN…

Cisco and Miercom: persistent, or slow learners?

Cisco has recently been distributing Miercom report 090815B, a "Lab testing summary report" dated September 2009. The new report shows that both players are basically acting the same way they have on previous reports (see appendix). In the new report, as in previous ones, Cisco paid Miercom to "evaluate" WAAS against Riverbed® Steelhead® appliances, and the resulting report declares that Cisco is the winner based on tests that look reasonable at first glance but which on analysis prove to be at best misleading, at worst completely ridiculous.

This document works through that analysis.

Wait a minute, something is strange here…

At a high level, of course, the report is completely predictable: Miercom claims Cisco WAAS is somehow superior to Riverbed Steelhead appliances, and there are some tests that allegedly support this claim. But entertainingly, we don't even have to get into the details of the claims or the evidence to notice that something is not quite right. Even though the report is dated September 2009, Miercom is comparing currently-available Cisco WAE-7371 appliances with Riverbed Steelhead 6020 appliances, whose end-of-sale was announced in January 2009 and were last offered for sale in April 2009. Of course, Riverbed continues to support Steelhead 6020 appliances, and hundreds are working as vital parts of customer networks every day. But it is notable that Miercom can't even be bothered to make an apples-to-apples test – with currently-available products – further reducing the report's value to anyone trying to make a purchase decision.

An interesting example of Miercom technique is how they attempted to minimize this problem. Look at this excerpt:

> *We spot checked other Cisco WAAS appliances to verify that the product architect and operating system are consistent within the product family line.*
> *Appliances supported for Cisco and Riverbed:*
> *Cisco WAAS appliances:*
> - *WAVE 274,474,674, WAE 512,612,7341,7371, NME-302,502,522*
>
> *Riverbed RiOS appliances:*
> - *SH-50,100,200,250 L,250M,250H, 550, 1050, 2050, 5050, 6020, 6050*

The first sentence makes a slightly strange statement about "spot checking" Cisco WAAS for consistency across the product line. If we are paying attention we might wonder, how can that matter? The rest of the excerpt makes an incorrect (and oddly-formatted, but reasonable-seeming) list of both Cisco and Riverbed models. This artful juxtaposition of "checking for consistency" with the two lists of models means that a casual reader is very likely to think that Miercom also checked the *Riverbed* appliances for consistency across the product line. But, brilliantly, Miercom never said any such thing.

Key Cisco/Miercom claims

The key claims in the document are that the Cisco WAAS device is superior in four areas:

- Better behavior under high load
- More accurate reporting
- Non-disruptive configuration changes
- Modular architecture

We consider the evidence for each of these in turn.

Is Cisco WAAS better behaved under load?

The first test in the Miercom report allegedly shows that a high-end Riverbed appliance can't handle optimizing 3000 concurrent TCP sessions while the WAAS appliance can handle 4000. This result immediately seems surprising for two reasons: first, the supposedly-failing Riverbed appliance (the Steelhead 6020) is rated at 40,000 concurrent connections and is deployed in

customer networks around the world. It would seem that a failure to handle less than 10% of the rated capacity would have led to widespread and widely-known problems. Second, the supposedly-successful Cisco appliance is rated at 50,000 concurrent connections, and so its alleged success is likewise happening at only a small fraction of its capability.

Presumably if the Cisco device had been able to handle 50,000 concurrent connections, that comparison would have been even more damning for Riverbed… so perhaps we can reasonably infer that the Cisco appliance couldn't handle much more than 4000 connections. In another part of the report, Miercom only claims 6000 concurrent connections for the same 50,000-connection appliance. What's really going on here?

Using a crazy test gives crazy results

Although it's not obvious from the report, the test that causes the problem apparently uses simultaneous download of large quantities of large objects via HTTP. When thousands of 1MB+ objects are simultaneously fetched over HTTP, and Riverbed HTTP optimization is enabled, a software flaw in the HTTP-specific optimization code causes problems very much like the ones shown in the report. Although it is technically correct to count "TCP sessions" because each HTTP session is carried over a TCP connection, it's misleading because it gives the impression that any kind of TCP traffic would have this problem – whereas in fact it's limited to optimized HTTP.

Is this a serious problem? We certainly don't want bugs in our software, and we work hard to fix them. We found and fixed this particular bug in RiOS® 5.5.5. But even if you have a version of Riverbed software with this particular flaw (problem number 47222), it likely won't cause any problem in a real-life network.

First, it turns out that there are few, if any, enterprise applications or customers that have this behavior of massive, simultaneous download of megabyte-sized objects over HTTP. No-one at Riverbed or at any of Riverbed's 6500+ customer deployments ever complained about this particular problem until Cisco found it and decided to make it a centerpiece of their marketing efforts. In responding to Cisco, we have replicated and isolated the problem – but to our knowledge it still hasn't turned up at any customer site.

Second, for any customer who is concerned about the possibility of this problem, the easy workaround is to disable HTTP optimization. Disabling HTTP optimization is not ideal, and might seem like a real Riverbed handicap against WAAS. But it turns out that WAAS 4.1.3 does not have any HTTP-specific optimization of its own. What Cisco calls "HTTP optimization" is just a form of TCP-connection reuse (which has its own problems, judging from Cisco WAAS release notes). So disabling Riverbed HTTP optimization puts Riverbed into a mode that matches Cisco's lack of optimization for HTTP. And again, it's worth underscoring that there are many Riverbed customers using HTTP optimization in production deployments without problems.

It's funny to realize that Cisco's most-alarming "breaking Riverbed" test depends on finding a flaw in a capability that their product lacks completely. Perhaps we will next see Cisco marketing material claiming that they deliberately avoided introducing HTTP optimization… because they didn't want to expose customers to the risk of having additional software bugs.

Putting it in perspective

Based on this analysis, we would argue that Cisco and Miercom have spent a disproportionate amount of time and energy on a single recently-fixed Riverbed bug, which they have over-interpreted into a far-reaching theory about the supposedly superior modularity of Cisco WAAS. The bug involves such an unusual workload that it was literally never seen before Cisco manufactured the right kind of torture test. Although we appreciate Cisco's investment in finding our bugs, we would caution them that "people who live in glass houses should not throw stones." As we'll see later, the bug-fixing situation in Cisco WAAS is nothing to boast about.

Is Cisco WAAS reporting more accurate?

The second highlighted claim on the first page is that "Reporting for Cisco WAAS was proven accurate; performance reporting matched measured efficiency." The implication is that Riverbed was somehow deficient in this area. The reality is that the only difference between Cisco and Riverbed in this area is whether TCP headers are reported or not.

There are reasonable arguments for both approaches. From its very first product, Riverbed has delivered optimization that is based on application-protocol-specific examination of payload data across multiple streams, without respect to packet boundaries. This approach was an important differentiator from earlier packet-based compression devices. Riverbed's strength in breadth of application-level optimization has led to a natural emphasis on measurements of optimization in terms that were relevant to application or storage. Accordingly, data related to packetization has been relegated to low-level reports only; the most common displays of optimization measure the output bytestream against the input bytestream, not the input packets vs. the output packets. In contrast, Cisco comes to the market from its strengths in routing and switching and understandably takes the view that packets are the only reality.

In the real world, the difference between the two reporting techniques is usually insignificant, exactly because it's desirable for headers to be small compared to payloads. In the screen shot that Miercom shows, they have generated a workload where more than half the bytes transmitted consists of headers. While such traffic may well exist, Miercom greatly exaggerates its importance.

Getting into the gory details

It's worth looking closely at the numbers in Miercom's screen shot so as to understand what's really happening here. On the LAN, they have 442945 packets while on the WAN they have 870519 packets. Already, we can see that something unusual is happening. Usually a WAN optimization device handles fewer packets on the WAN than on the LAN for two reasons. First, there are usually fewer payload bytes to transmit if compression and deduplication are working. Second, the payload data that arrives in multiple small incoming LAN packets can often be coalesced into a single larger packet on the WAN. Here, in contrast, the WAN has nearly twice as many packets as the LAN. How can that be? There are 53423956 bytes on the WAN carried in 870519 packets, which gives an average packet size of slightly more than 61 bytes. Meanwhile, there are 138061743 bytes on the LAN carried in 442945 packets, for an average packet size of somewhat less than 312 bytes. The packet size on the LAN is more than 5 times larger than the packet size on the WAN! While Miercom omits the details of their workload and configuration that would allow a diagnosis, such a small packet size might suggest that there is something wrong: perhaps a problem with the WAN MTU set too small, a small server MTU with Nagle timers disabled, or similar errors.

The bottom line is that the only way that Miercom can create a compelling difference between the Riverbed and Cisco reporting is by setting up a crazily bad network. A network engineer seeing the traffic in Miercom's highlighted screenshot would be trying to fix the network, not pondering differences between vendor WAN optimization reports.

Is Cisco WAAS better at non-disruptive configuration changes?

Miercom's third highlighted claim is that "Configuration changes were proven in testing to be nondisruptive with the Cisco WAAS appliance" with the implication that Riverbed was somehow deficient in this area.

The only specific example cited is the disabling and re-enabling of CIFS optimization. For Riverbed, this particular configuration change requires restarting the optimization service; for Cisco, it does not.

Even taking this comparison at face value, Miercom's approach to this issue is silly. Both products are complex with hundreds of features and options: picking a single test case where one product happens to behave better than the other is a weak basis for claiming something about the overall comparison of products. One could just as easily pick some part of product configuration that is more difficult for Cisco than for Riverbed. For example, many customers observe that it's impossible to do a quick tactical two-box deployment of WAAS on a problematic connection – every WAAS deployment needs a full-blown WAAS central manager, with all of its overhead.

And as usual, Miercom takes a reasonable and accurate data point – the specific case of enabling/disabling CIFS optimization – and wildly exaggerates its significance. Their summary of this issue is

> *Overall, we deem the configuration changes as non-disruptive for the Cisco product, whereas the Riverbed RiOS product required service resets after configuration changes.*

If you just read that one sentence above, you would think that Miercom had tested this issue extensively and that *any* configuration change to a Riverbed device requires some kind of restart. In fact, of course, the vast majority of typical

configuration changes to a Riverbed device don't require a restart. Again Miercom implies that Riverbed is somehow unstable, misleading, unwieldy, and unusable – leading to the puzzle of why more than 6500 organizations have deployed Riverbed in their networks (well ahead of the number who have deployed WAAS) and why Riverbed has won 9 times out of 10 in direct head-to-head customer evaluations against Cisco WAAS.

Is Cisco WAAS more modular?

For the three claims we've already considered, there is a pretty good alignment between the summary on the front page and the evidence in the following pages. The evidence turns out to be weak or nonsensical, but at least there's something presented. For the last claim, however, this relationship breaks down completely. The Miercom claim on the front page is that the supposed advantages of WAAS are due to superior modularity:

> *Cisco WAAS modular architecture proved in testing to offer superior scalability and stability, high performance, high availability and improved fault management*

However, the only testing Miercom described that relates to scalability, stability, performance, availability, or fault management is the already-analyzed simultaneous-HTTP-download test (with its highly implausible workload). Students of Miercom technique should note how much mileage they are getting out of a single unconvincing test exploiting a single bug in a competitive product. Not only did they already use this to point out how dangerous Riverbed products are to your network, now they are also able to establish that Cisco's architecture is superior in all these areas… and all from just one test! Of course, one serious problem with this line of argument is that all the supporting evidence went away when the Riverbed bug was fixed… too bad.

Miercom declares that satisfied customers are wrong

The closest we can find to evidence for the supposed superiority of WAAS modularity is a section on "Operating System Revision Cycle." This short section is worth analyzing closely:

> *Miercom conducted an audit of the software revision changes for both the Cisco WAAS and the Riverbed RiOS. In a one year period April 2008 to March 2009, there were 100 reported releases of the Riverbed RiOS from version 3.X to 5.X, whereas, Cisco had eight releases 4.0.x to 4.1.x. The cycle of updates required for the Riverbed Steelhead product is deemed overly excessive.*

Let's assume for a moment that Miercom is accurately counting releases for both vendors, because there are interesting aspects to this section even without considering the data. The first question is: Does Miercom really think that appliances need to be upgraded on each release? (The appliances don't, but that's what you might think from the phrase "updates required.") The next question is: why does Miercom use the hilariously exaggerated phrasing "overly excessive?" A third question is: on what grounds does Miercom judge Riverbed software practice to be excessive? Are they acquainted with the incredibly elaborate release structure of Cisco IOS including trains, major/minor versions, releases, builds, and rebuilds? Does anyone even know how many variants of IOS are released per year?

Clearly, software releases are not an end in themselves, they are a means to customer satisfaction. Miercom did not "audit" Riverbed customer satisfaction, but Riverbed performs semi-annual customer surveys. Riverbed conducted such a survey in October 2008 – which we choose because it falls in the period (April 2008 to March 2009) for which Miercom examined RiOS releases. The summary appears in a Riverbed press release from November 21, 2008 available on the Riverbed web site (reference at end). Riverbed received ratings well above 8 on a 10-point scale in a variety of areas – *including software quality*. The director of the research company conducting the survey is quoted as follows:

> *In the multitude of customer satisfaction studies we have conducted for technology companies over the past thirteen years, results have shown that the Riverbed customer base is among the most loyal and enthusiastic. The overall scores reported are some of the highest that we have seen.*

We can see that whatever Riverbed was doing, it led to happy customers. Apparently customers like what (Cisco-paid) Miercom labels as "overly excessive."

Miercom raises the specter of Riverbed bugs (bad idea, as it turns out)

In classic Miercom fashion, a completely unwarranted conclusion is presented as though it follows from the previous weak argument:

> *Due to the Riverbed monolithic architecture, any small bug can be critical and crash the Riverbed core WAN optimization process "SPORT".*

Rather than seeing the Riverbed software development process as an essential part of our highly-effective focus on solving customer problems, Miercom imagines that it shows a Riverbed vulnerability to crashing. We will take up the idea of "monolithic architecture" in a moment, but first let's consider the question of comparing WAAS bugs to RiOS bugs.

In the most recent release of WAAS (4.1.3b) Cisco acknowledges the presence of 72 "open caveats." (These open caveats are split across three sub-releases, so a careful count is required to get the total). Meanwhile, in the most recent release of RiOS (5.5.4d) Riverbed acknowledges the presence of two "known issues." Of course, the two companies likely have different standards for how to categorize and describe problems, so these two numbers are not directly comparable. However, it's useful to note the net improvement/degradation over the last couple of minor releases.

In moving from WAAS 4.1.3 to 4.1.3a, Cisco removed two open caveats and added 20, for a net increase of 18 notable bugs. In addition Cisco "found and fixed" 11 bugs not previously documented between the two releases. From WAAS 4.1.3a to 4.1.3b, Cisco again removed two but added 14, for a net increase of 12 notable bugs. In addition Cisco found and fixed four bugs. The overall summary is that Cisco added 30 notable bugs over the two most recent point releases.

RiOS 5.5.4b and 5.5.4c have the same two "open issues" Riverbed notes in RiOS 5.5.4d, so Riverbed shows no net change in notable bugs. From 5.5.4b to 5.5.4c, Riverbed "found and fixed" four bugs not previously documented. From 5.5.4c to 5.5.4d, Riverbed found and fixed two bugs. The overall summary is that Riverbed did not add or remove any notable bugs over the two most recent point releases.

It's still not possible to directly compare these repair rates. We don't have enough detail in the release notes from the two companies to determine whether Riverbed has shorter release cycles, slower bug fixing processes, or just has fewer bugs. But the net direction of improvement/degradation seems clear: while both companies have some "finding and fixing" activity, only Cisco is finding many more notable bugs than they fix in each release.

Looking at the "found and fixed" bugs, we see that Cisco had 15 across two releases while Riverbed had six. Even if we then assume that Cisco's bug assessment process publishes at finer granularity than Riverbed (with Cisco publishing 15 issues for every six that Riverbed publishes), that would still lead us to expect Riverbed to have had a net increase in "open issues" of (30 * (6 / 15)) = 12. In fact, Riverbed had a zero net increase in open issues. An actual Riverbed zero vs. a Cisco-predicted 12 is a significant gap that sheds some light on anecdotal complaints about WAAS software quality.

In a surprising contrast to Miercom's attempt to imply that Riverbed has a software quality problem, the recent release note data actually suggests that Cisco has a significant and ongoing problem with an increasing number of notable software defects in WAAS.

Monolithic software: the pot calls the kettle black

Returning to the first part of the Miercom quoted claim, the only sensible observation about Miercom's "monolithic" claim is that there are advantages and disadvantages to how different software systems are structured. Cisco has chosen to structure WAAS as a large collection of relatively small communicating processes, while Riverbed has a smaller collection of larger processes. Both competing systems are built as collections of processes within a Linux operating system, so it is misleading for Miercom to call Riverbed "monolithic." Perhaps the most widespread example of a truly monolithic software system in networks is Cisco IOS, which has all code in a single shared address space and non-preemptive scheduling. Accordingly, it is somewhat ironic to have Cisco (via Miercom) claiming that Riverbed's much-more-modular system is somehow unsuitable for deployment in networks. Does the WAAS group likewise recommend that customers should avoid Cisco routers and switches?

Miercom chooses to highlight a natural advantage of WAAS's approach, which is that you would expect features to be more independent in their behavior. However, a corresponding natural disadvantage to that structure is that sharing of state is more difficult and expensive across the finer-grained process architecture, and can impair performance. Of course, Miercom does not highlight those issues in their report.

Conclusion

All of the major claims made by Miercom turn out to be either incorrect or misleading (technically correct but not as significant as Miercom claims). We can only marvel at the level of desperation there must be in the WAAS business unit to feel that these lame smears are the best attacks they can make on Riverbed.

Why do companies like Miercom exist, and why do companies like Cisco hire them? Apparently customers have learned through hard experience that many vendors exaggerate or stretch the truth. Some vendors then seek to overcome this problem by enlisting "objective" testers such as Miercom. It may also be that these vendors have difficulty with hands-on evaluations and hope to substitute these "independent" evaluations for actual use of their products in customer networks. It appears that this arrangement works, at least some of the time, or else Miercom and similar tester-for-hire companies would be out of business.

It is also puzzling that a large and reputable vendor like Cisco would need to follow the same slightly-pathetic path as struggling startups like Silver Peak. One might expect that Cisco would not seek this sort of dubious external endorsement. A possible insight into this situation comes if we look at the Gartner Magic Quadrant for WAN Optimization Controllers (URL at end of this document). There we can see that Cisco and other Miercom clients all share the problem of falling well outside the "Leaders" quadrant. In addition to the Magic Quadrant, we list some additional Gartner reports on WAN optimization that the interested reader may find useful. Although no analysis is perfect, Gartner's Magic Quadrant and similar reports are not funded by any of the vendors – in contrast to the way that Miercom works.

One might have also expected someone at Cisco to be intelligent enough to see that the Miercom reports commissioned by those companies did not make any significant difference to their competitive position against Riverbed.

The success of Riverbed is encouraging for those who believe business should be about quality products instead of spin and hype. In somewhat more than five years of shipping products, and without sponsoring dubious competitive testing, Riverbed has won over 6500 customers – many of whom have become enthusiastic advocates for Riverbed. Riverbed's continued success is a testament to the number of customers with the courage to analyze issues for themselves.

Frankly, at Riverbed we would rather spend our time building better products instead of responding to badly-done reports by rent-a-tester companies. We hope customers feel likewise and will encourage competing vendors accordingly.

APPENDIX

Miercom: Somehow, the sponsor always comes out best

Miercom is a company that is paid to run tests. The tests they perform lead to results that favor the sponsor. This relationship has held up for every test they have done so far in the WAN optimization market.

For WAN optimization alone, prospective customers can enjoy reading previous Miercom reports that favor:

- Silver Peak: report 051207;
- Citrix (formerly Orbital Data): reports 060501 and 060725;
- Blue Coat (formerly Packeteer): report 070914; and
- Cisco: reports 060925, 061013, and an unnumbered "Test Review" dated 30 November 2007.

Full URLs are provided at the end of this report.

It is worth citing a few of the ways in which Miercom praises these competing products. See if you can tell which one came from which report:

- "the Silver Peak product consistently outperformed its counterpart"
- "while all three products perform optimization, only Orbital Data can be considered transparent to the existing network infrastructure and in-place security policies."
- "The Packeteer iShaper 400 provides enterprise customers with a powerful collection of tools and services in a single platform"
- "Cisco WAAS demonstrated excellent performance with compression, speed, and throughput."

Yes, each of these positive assessments is from the report sponsored by the named vendor.

Meaningful? No.

It is worth being clear about how Miercom's approach breaks pretty much every rule that would be in place for a meaningful comparison of products.

1. No access to raw data: Miercom simply provides summaries making the points that they want to make. Sometimes Miercom does not even provide the same measurements for the units being compared – for an example, please see the Riverbed white paper "Desperation is not Pretty" rebutting a previous Cisco/Miercom report.

2. Non-reproducible tests: the tests are so vaguely described that it would be impossible to reproduce them based on the Miercom documents alone. Miercom makes a practice of omitting significant aspects of their tests from their published reports.

3. Uneven treatment: the sponsoring vendor is free to deliver different equipment, configuration changes, or entirely new tests, with Miercom restarting their testing process as required.

4. Inconsistent tests: Miercom has tested competitive products with roughly similar functionality multiple times and has used different tests each and every time. This is not defensible by claiming that they are learning over time which tests are better or worse: tests are simply created for one client and discarded when it's time for the next client.

Miercom is consistent in some ways

Although Miercom does not use any consistent test methodology, there are other consistencies in their reports that are worth noticing.

Because competing vendors know that Miercom is working on behalf of the sponsoring vendor, they consistently decline Miercom's invitation to assist. This arrangement at least avoids the appearance that the competing vendors endorse Miercom's

work, but it also means that Miercom is free to proceed on the basis of misunderstanding or misconfiguring the competive systems. Situations that might be the basis of a support case for a real customer (such as lower-than-expected performance) are instead published as though they are the best achievable results.

Although the latest report is a consistent example of Miercom's efforts on behalf of Cisco, it is unusual in two ways that suggest Miercom is being even less straightforward than they have been previously. First, this report is missing the usual disclaimer explaining that Riverbed was invited but declined to provide assistance. Instead, this report is presented as "an excerpt from an independent Industry Assessment on Wide Area Network (WAN) Optimization products that Miercom conducts annually." Second, there are references to problems "as expected, and explained in Riverbed technical notes." Although it would be plausible to think otherwise, Miercom did not conduct an independent assessment across the industry with Riverbed assistance. And although it might look like Cisco only paid for these relevant excerpts, in fact the Miercom tests and analysis are very similar to materials that Cisco circulated on its own months earlier. As indicated previously, Riverbed has seen no evidence that Miercom is independent and objective. Riverbed provided no assistance to these tests.

References

References to Miercom reports

Silver Peak: report 051207
http://www.miercom.com/dl.html?fid=20051207&type=report&PHPSESSID=c8e84098c29b5e9d1f687962c2833950

Orbital Data (now Citrix): report 060501
http://www.miercom.com/dl.html?fid=20060501&type=report&PHPSESSID=429457593f087300802a3f3691289596

Packeteer (now Blue Coat): report 070914
http://www.miercom.com/dl.html?fid=20070914&type=report&PHPSESSID=507081d1bfb92c4dfdfedf78268a4258

Cisco: report 090815B
http://miercom.com/dl.html?fid=20090915&type=report
Cisco: report 0701215d
http://miercom.com/dl.html?fid=20071215&type=report
Cisco: report 070529
http://miercom.com/dl.html?fid=20070529&type=report
Cisco: report 061013
http://www.miercom.com/dl.html?fid=20061013&type=report
Cisco: report 060925
Report 060925 is not available on the Miercom site. However, it was widely distributed by Cisco:
http://www.cisco.com/web/CA/channels/pdf/Miercom-on-Cisco-WAAS-Riverbed-Juniper-competitive.pdf

References to Gartner Reports

“Magic Quadrant for WAN Optimization Controllers, 2006” by Andy Rolfe and Joe Skorupa. G00141742, October 11 2006.
“Magic Quadrant for WAN Optimization Controllers, 2007” by Andy Rolfe and Joe Skorupa. G00153256, December 14 2007.
“Cisco WAAS Update: Client Feedback Suggests to Proceed With Caution” by Mark Fabbi and Joe Skorupa. G00160302, September 17 2008.
“Magic Quadrant for WAN Optimization Controllers, 2009” by Andy Rolfe, Joe Skorupa, and Severine Real. G00165875, June 30 2009.

All reports available to Gartner subscribers at www.gartner.com. Reports may also be available for purchase individually.
We encourage readers to be familiar with Gartner’s Copyright and Quote Policy:
http://www.gartner.com/technology/about/vr_copyright.jsp

Reference to Riverbed press release on customer satisfaction

“Riverbed releases results of semi-annual customer survey” Riverbed press release, November 21, 2008.
http://www.riverbed.com/company/news/press_releases/press_112108.php

About Riverbed

Riverbed Technology is the IT infrastructure performance company. The Riverbed family of wide area network (WAN) optimization solutions liberates businesses from common IT constraints by increasing application performance, enabling consolidation, and providing enterprise-wide network and application visibility – all while eliminating the need to increase bandwidth, storage or servers. Thousands of companies with distributed operations use Riverbed to make their IT infrastructure faster, less expensive and more responsive. Additional information about Riverbed (Nasdaq: RVBD) is available at www.riverbed.com.

riverbed
Think fast.™

Riverbed Technology, Inc.
199 Fremont Street
San Francisco, CA 94105
Tel: (415) 247-8800
www.riverbed.com

Riverbed Technology Ltd.
Farley Hall, London Road
Bracknell, Berkshire RG42 4EU
United Kingdom
Tel: +44 1344 401900

Riverbed Technology Pte. Ltd.
391A Orchard Road #22-06/10
Ngee Ann City Tower A
Singapore 238873
Tel: +65 6508-7400

Riverbed Technology K.K.
Shiba-Koen Plaza Building 9F
3-6-9, Shiba, Minato-ku
Tokyo, Japan 105-0014
Tel: +81 3 5419 1990