


White Paper

Critical Elements of Testing Wireless Mobile IP

February 2004



Spirent Communications, Inc.

26750 Agoura Road
Calabasas, CA
91302 USA

Email: productinfo@spirentcom.com

Web: www.spirentcom.com

North America

+1-800-927-2660

Europe, Middle East, Africa

+33-1-6137-2250

Asia Pacific

+852-2166-8382

All Other Regions


+1-818-676-2683

Copyright

© 2004 Spirent Communications, Inc. All Rights Reserved.

All of the company names and/or brand names and/or product names referred to in this document, in particular, the name "Spirent" and its logo device, are either registered trademarks or trademarks of Spirent plc and its subsidiaries, pending registration in accordance with relevant national laws. All other registered trademarks or trademarks are the property of their respective owners.

The information contained in this document is subject to change without notice and does not represent a commitment on the part of Spirent Communications. The information in this document is believed to be accurate and reliable, however, Spirent Communications assumes no responsibility or liability for any errors or inaccuracies that may appear in the document.



Critical Elements of Testing Wireless Mobile IP

This White Paper provides a discussion of testing wireless Mobile IP.

Contents

Introduction	2
The Need for Benchmarking	2
Data Plane Performance.....	5
Control Plane Performance	7
Other Features	8
Summary	9

Introduction

CDMA mobile wireless technologies have enjoyed rapid proliferation, particularly in North America and Asia. Consequently, CDMA data-overlay services, such as CDMA2000 1xRTT, are enjoying similar market success.

One of the most compelling attractions of these data services is that operators can quickly double the effective throughput of the network¹ without additional frequency spectra. One particular benefit is that operators can add data services to an existing radio network with minimal incremental capital investment in infrastructure.

On the other hand, any technology that emerges so suddenly is likely to expose unforeseen limitations and problems lurking behind pilot services and implementations. In that regard, CDMA2000 data services are no exception: They have appeared so quickly that many operators and network equipment manufacturers have been left virtually devoid of any empirical data on data network performance. They can only guess whether their production network infrastructure is truly ready to support one or two orders of magnitude increase in subscriber population and application traffic loads.

The Need for Benchmarking

The widespread adoption of CDMA2000 data services means that simple functional and interoperability testing are no longer adequate to ensure network stability and performance. Instead, operators must explore a new set of benchmarks that include such metrics as data connection density, aggregate data throughput, data connection set-up time, packet latency, and packet loss.

To operators and providers traditionally focused on voice services, these new data benchmarks are effectively uncharted waters. Moreover, such benchmarks must reflect the level of performance for which subscribers are willing to pay. Finally, to ensure continued customer satisfaction, the operators also must be able to validate this performance level regularly as network loads and traffic profiles evolve.

Finding the Bottleneck

Conventional wisdom generally faults the air interface as the major bottleneck in wireless data network performance. However, this is a misconception.

Most CDMA2000 data networks now support data rates over 100 kb/s² thanks to the successful implementation of 1xRTT enhanced data services. Moreover, 1xEV-DO is already appearing in various subscriber “islands” and will likely become commonplace

-
1. 1xRTT allows voice and data to share the same spectrum, effectively doubling the number of signals that each mobile station can transmit and receive.
 2. 1xRTT supports a theoretical rate of 153 kb/s. Even compensating for network overhead and delays, this can frequently translate into over 100 kb/s of throughput upstream and downstream (i.e., an aggregate throughput of over 200 kb/s).

before the end of 2004. This means that the air interface is already supporting hundreds of kb/s. For example, a simple calculation will show that 1,000 subscribers at an average throughput of 150 kb/s (averaged across a collection of 1xRTT and 1xEV-DO connections) would exceed the capacity of even a Fast Ethernet uplink.

In such cases, how would the network behave? Would the network decline new connections, or force all connections to suffer equally? Would it share bandwidth equally among users or would some users garner a disproportionate quantity of network resources? Would latency and packet loss be equitable across all users? Will the network become unstable under heavy load? Will a faster uplink, such as Gigabit Ethernet, alleviate or exacerbate the problem?

These are not academic questions but rather genuine concerns for operators who are considering CDMA2000 data service offerings. The best way to answer these questions is through careful laboratory and field benchmark testing.

Mobile IP and Simple IP

The Spirent CDMA2000 Performance Test System supports both Mobile IP and Simple IP, including Mobile IP inter-PDSN handoff performance

In the case of CDMA2000 data networks, benchmarking should include both of the two leading IP data technologies – Mobile IP and Simple IP. Mobile IP allows a device to roam among several different IP networks while maintaining a single, “virtual” IP address.³ CDMA2000 networks most often use Mobile IP for IP Mobility, but some instead rely upon its close relative, coined “Simple IP.”⁴ Together, Simple IP and Mobile IP form the cornerstones of CDMA2000 wireless Internet access (see [Table 1](#)) and so define the protocol requirements of the wireless data network.

Table 1. Comparison of Mobile IP and Simple IP

	Simple IP	Mobile IP
Mobility	Within one PDSN	Between multiple PDSNs
Standards Governing Body	3GPP2	IETF
Complexity	Simple	Complex

Differences aside, both technologies have gained considerable traction for data connectivity and therefore warrant careful benchmarking.

3. Of course, Mobile IP is distinct from “IP Mobility.” The latter is a superset of the former and includes a number of wireless Internet protocols and technologies, including WLAN (“Wi-Fi”), GPRS, CDMA2000 and W-CDMA, among others.
4. Simple IP allows roaming within a single PDSN; Mobile IP supports roaming between PDSNs.

PDSN-FA and HA

*The Spirent
CDMA2000
Performance Test
System tests the
HA and FA in
isolation as well
as in combination*

Benchmarking the performance of the CDMA2000 data network must include benchmarking both (a) devices in isolation and (b) devices in combination. This is particularly critical for core network infrastructures where thousands of users are converging on the Packet Data Serving Node (PDSN)⁵ and on the Home Agent (HA). This translates into a requirement for a trio of tests: (a) PDSN alone, (b) HA alone, and (c) PDSN and HA in combination.

PDSN Alone

Testing the PDSN requires emulating the wireless network on one interface, using the Radio-Packet (R-P) protocol, and the HA on the other. This mandates that the test tool transmit traffic from the HA to the PDSN (downstream) but measure traffic directly from the PDSN to the Internet (i.e., bypassing the HA for upstream traffic). Such a test should also include roaming between PCFs⁶ attached to the PDSN. The diagram in [Figure 1](#) shows a PDSN-only test.

*Test
Recommendation
Measure Mobile
IP inter-PDSN
handoff times*




Figure 1. PDSN-Only Test

HA Alone

Testing the HA requires emulating Internet hosts on one interface (IP only) and emulating the PDSN on the other (Mobile IP). Such a test may also include roaming between PDSNs.

The diagram in [Figure 2 on page 5](#) shows the logical placement of devices in an HA test. Note that traffic from the PDSN to the Internet bypasses the HA and therefore is effectively irrelevant to the HA test.

5. In most cases, the PDSN is also co-located with the Foreign Agent (FA), so the PDSN is often called PDSN-FA. In this document, PDSN refers to the combination of PDSN and FA.

6. Packet Control Functions (PCFs) are the CDMA2000 equivalent of base station controllers.

*Test
Recommendation
Benchmark
PDSN alone, HA
alone, as well as
PDSN and HA in
combination*




Figure 2. HA-Only Test

PDSN-FA and HA in Combination

Finally, testing the PDSN-FA and HA in combination requires mapping R-P connections between the PCF and the PDSN-FA to send traffic from the PDSN-FA destined for the HA. Essentially, it involves combining half of the PDSN test with half of the HA test. The diagram in *Figure 3* helps to show the logical connections in this test. Note that the PDSN-FA and HA under test manage the “triangular” flow between the Internet and the PCF.




Figure 3. Combination PDSN-FA and HA Test

Data Plane Performance

The previous discussion focused on the network configurations that should be included in benchmark testing. In this section, we will discuss the specific data plane metrics that one should measure.

First, the test tool should exercise the device under test’s data forwarding performance to determine how it will perform under both current and future traffic loads. There are two critical factors – throughput and latency – that can be measured in isolation and in combination with one another, and each of which has a direct impact on the overall performance of the device in actual deployment.

Throughput

*The Spirent
CDMA2000
Performance Test
System measures
packet loss at
various sizes and
in various system
configurations*

In theory, throughput describes the maximum, zero-loss, steady-stream packet generation at different packet sizes. The problem with this metric is that the packet loss is relative. For instance, a test that measures zero packet loss after 60 seconds does nothing to guarantee that no packet loss will occur after 120 seconds. Moreover, the packet loss metrics are generally higher with smaller packets, so the margin of error for such metrics decreases with decreasing packet size. Finally, the concept of packet loss presumes that there is a monotonic threshold of loss below which performance is acceptable and above which it is unacceptable. The limitations of such metrics highlight the need for another method of testing – using application traffic – that we will address after a brief discussion of latency.

Latency

*The Spirent
CDMA
Performance Test
System measures
both latency and
application
response time*

The second performance component is latency. While CDMA2000 latency benchmarking might seem far more straightforward than the loss-dependent benchmarking described in the previous section, it is actually less so. Specifically, one must measure latency at a particular packet size and at a particular traffic load, often at the maximum zero-loss threshold described in the *“Throughput”* section. Therefore, before one can measure latency, one must determine the maximum zero-loss throughput. Since this value may vary from one device to another, latency metrics will reflect different network loads from one device to another. For this reason, so-called “packet blasting” latency metrics are best measured at several traffic loads (e.g., 10% to 90% of the theoretical maximum load at increments of 10%). Such testing naturally requires very flexible automation on the part of the test equipment (or a great deal of the testers’ free time!).

*Test
Recommendation
Measure latency
across a full
range of packet
sizes and traffic
loads*

Note that this discussion specifically excludes jitter as a metric. The reason is simple: Jitter is meaningful in a network in which end stations cannot buffer traffic for the longest acceptable delay. That is, if the maximum tolerable end-to-end delay is 250 ms (i.e., after 250 ms, a packet is considered lost), and if the end station can buffer 250 ms worth of traffic, jitter has no effect. However, if the end station can only buffer 125 ms of data, then the jitter (maximum end-to-end delta latency) cannot exceed 125 ms without packet loss. However, the commoditization of memory often makes it easier to increase end station memory than to engineer the network for low jitter. Therefore, we shall exclude jitter from this discussion.

Application Traffic

*The Spirent
CDMA2000
Performance Test
System can
generate up to 4
gb/s of data
upstream and
downstream*

An equally valid but much different approach to measuring network forwarding performance and latency involves using application traffic instead of, or preferably in addition to, throughput and latency metrics. Such an approach should include emulation of protocol dialogues, including packet header formatting and request/response transactions that emulate the real application.

Today, CDMA2000 users run a variety of applications over the wireless network, often expecting the wireless network to provide a ubiquitous, albeit lower-bandwidth, complement to existing wired services. This trend is reflected in the rapid growth of the PC card market for mobile wireless devices, including both CDMA and GSM, as well as emerging offerings for next-generation services such as 802.16.

*Test
Recommendation
Benchmark
application
throughput with
and without other
traffic loads
across the device
under test*

This mandates that test tools provide application traffic metrics across the CDMA2000 core network environment. Although generating near-wire-speed application traffic is difficult, good test equipment can approach wire-speed performance. Moreover, rather than establishing a monotonic, pass/fail packet loss threshold, TCP-based transactions dynamically converge to a steady-state throughput for a given network.

To exercise the functionality of the wireless infrastructure, the test tool should drive high-volume application traffic, such as HTTP and FTP, through the network to determine how well the application performs both in the absence of, and in the presence of, other connections and potentially other data traffic.

Summary

Ideally, one would perform both real-world application traffic tests and packet streaming tests. Both approaches have strong proponents, and the best CDMA2000 test will exercise both of them.

Control Plane Performance

*The Spirent
CDMA2000
Performance Test
System emulates
over 6 million
simultaneous
mobile
subscribers*

The discussion of CDMA2000 performance so far has focused exclusively on data throughput. Still, there are two important elements of control plane performance that directly impact application performance – elements that every CDMA2000 test tool must support.

Rapid growth in the popularity of CDMA2000 data services, including 1xRTT and 1xEV-DO, coupled with the promise for eventual evolution to 1xEV-DV, promises further growth in subscriber populations. This means that operators and equipment vendors alike must design ample performance and scalability headroom for future subscriber growth. This translates to benchmarking the effects of possibly tens or even hundreds of thousands of users accessing the network simultaneously; ideally, the test tool will also scale into the millions of emulated mobile nodes.

*The Spirent
CDMA2000
Performance Test
System emulates
setting up 32,000
Mobile IP and/or
Simple IP tunnels
per second*

In addition to the number of static connections that the network can support, operators must also determine the speed at which the network can establish and tear down connections. Such metrics are critical in determining how the network will respond to service restoration after an outage and to demographic events (e.g., breaks in activity during a soccer/football game). This mandates that the test tool provide session activation/deactivation rates that exceed the likely maximum system capacity by at least 20%. This equates to thousands and even tens of thousands of connections per second.

Other Features

While control plane and data plane performance metrics are among the most critical network benchmarks, there are three other factors that also directly impact how effectively one can leverage those benchmarks. They are: (a) ease-of-use, including automation, (b) support for IPv6, including Mobile IPv6, and (c) diagnostic features to facilitate troubleshooting.

Ease of Use

The Spirent CDMA2000 Performance Test System provides a menu-driven GUI to simplify configuration and automate testing

Once the customer has evaluated basic performance and scalability, he/she should immediately turn his/her attention to options for more granular configurations and metrics. Test tools that offer highly-granular, feature-rich configuration options can quickly become cumbersome to configure. Thus, a test tool's ease-of-use and, to a greater extent, test automation become increasingly important in the more feature-rich CDMA2000 test tools. Since ease-of-use is, to some extent, subjective, it is incumbent upon the customer (manufacturer or operator) to familiarize himself/herself with the tool and determine, based upon his/her particular requirements, the relative ease or difficulty of the test tool.

IPv4, IPv6

The Spirent CDMA2000 Performance Test System currently supports IPv4 and will support IPv6 later in 2004

Any serious Mobile IP or Simple IP test tool must include IPv6, including simultaneous support for both IPv4 and IPv6.⁷

Evaluating mobility infrastructure should include benchmarking the relative performance of IPv4 and IPv6 in isolation as well as in combination. Notably, IPv6 adoption in Asia – a hotbed of CDMA2000 development – is growing quickly. Although most of the North American CDMA2000 arena is still operating comfortably on IPv4, the inevitability of IPv6 in North America makes validation of IPv6 infrastructure performance unavoidable.

Diagnostics

A test system that emulates millions of subscribers connecting to the network at a rate of tens of thousands of tunnels per second is likely to exceed the limits of the system under test by a wide margin. While this desirable for a test tool, it also adds a requirement for rich diagnostics to determine the cause of connection failure, data loss, transit delays and other anomalies.

7. IPv6 eliminates the Foreign Agent and instead enables IP hosts to address the mobile node's Foreign Address directly.

*The Spirent
CDMA2000
Performance Test
System logs
events for
troubleshooting
and detailed
transaction
analysis*

One of the key metrics for operators and service providers is the response of the network infrastructure during periods of peak loading, especially at levels beyond what the infrastructure is designed to support. That is, although capacity planning presumes some level of under- or over-provisioning, it is an inexact science: Subscribers – not calculators – dictate load.

By comparison, the traditional telephony model – with which many wireless operators are familiar – prescribes that, during periods of unexpected system load, callers would receive a “circuit busy” signal. However, although access to new entrants would be delayed, at no time would the network actually fail, and calls in progress would continue to normal completion. Wireless operators lack empirical information on how their networks will respond under similar peak loading conditions yet they need this information desperately.

Summary

Benchmarking CDMA data services has long since transcended the boundaries of simple conformance and interoperability. Today’s rapid expansion in subscriber population, coupled with the aggressive broadening of wireless data applications, is stretching the core wireless infrastructure beyond the breaking point.

Now, more than ever, operators need reliable data for capacity planning – data that accurately and realistically profiles subscriber load. They need to know how many subscribers the network can reliably support and under what conditions. Only through the acquisition of such data can the operator hope to sustain customer satisfaction and retain customers.

Such benchmarking requires both time and capital investment. However, for the operator ready to engage this lucrative CDMA2000 data service market – and for the equipment manufacturer building out that network – the benchmark metrics are absolutely required to control the risk of service failures. A network that delivers services reliably to a growing satisfied customer base will easily show a substantial return on investment.

Spirent Communications has a broad portfolio of test tools to benchmark these data services accurately and easily. If you would like to speak with one of our representatives today, please call 1-800-SPIRENT and request information on the CDMA2000 Performance Test System.

