



NetMotion Mobility XE™ Report

A Broadband-Testing Report

First published February 2009 (V1.0)

Published by Broadband-Testing
A division of Connexio-Informatica 2007, Andorra La Vella

Tel : +376 633010
E-mail : info@broadband-testing.co.uk
Internet : [HTTP://www.broadband-testing.co.uk](http://www.broadband-testing.co.uk)

©2009 Broadband-Testing

All rights reserved. No part of this publication may be reproduced, photocopied, stored on a retrieval system, or transmitted without the express written consent of the authors.

Please note that access to or use of this Report is conditioned on the following:

1. The information in this Report is subject to change by Broadband-Testing without notice.
2. The information in this Report, at publication date, is believed by Broadband-Testing to be accurate and reliable, but is not guaranteed. All use of and reliance on this Report are at your sole risk. Broadband-Testing is not liable or responsible for any damages, losses or expenses arising from any error or omission in this Report.
3. *NO WARRANTIES, EXPRESS OR IMPLIED ARE GIVEN BY Broadband-Testing. ALL IMPLIED WARRANTIES, INCLUDING IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT ARE DISCLAIMED AND EXCLUDED BY Broadband-Testing. IN NO EVENT SHALL Broadband-Testing BE LIABLE FOR ANY CONSEQUENTIAL, INCIDENTAL OR INDIRECT DAMAGES, OR FOR ANY LOSS OF PROFIT, REVENUE, DATA, COMPUTER PROGRAMS, OR OTHER ASSETS, EVEN IF ADVISED OF THE POSSIBILITY THEREOF.*
4. This Report does not constitute an endorsement, recommendation or guarantee of any of the products (hardware or software) tested or the hardware and software used in testing the products. The testing does not guarantee that there are no errors or defects in the products, or that the products will meet your expectations, requirements, needs or specifications, or that they will operate without interruption.
5. This Report does not imply any endorsement, sponsorship, affiliation or verification by or with any companies mentioned in this report.
6. All trademarks, service marks, and trade names used in this Report are the trademarks, service marks, and trade names of their respective owners, and no endorsement of, sponsorship of, affiliation with, or involvement in, any of the testing, this Report or Broadband-Testing is implied, nor should it be inferred.

TABLE OF CONTENTS

TABLE OF CONTENTS III

BROADBAND-TESTING IV

EXECUTIVE SUMMARY 1

INTRODUCTION: THE VALUE OF KEEPING CONNECTED 2

NETMOTION MOBILITY XE: A PRODUCT AND SERVICE OVERVIEW 3

 What Exactly Is Mobility XE? In A Nutshell... 3

 How Does It Work?..... 4

 And Roaming? 8

NETMOTION MOBILITY XE: PUT TO THE TEST 12

 Roaming 12

 Data Compression 15

 Web Acceleration 17

 Policy Management..... 18

 Network Access Control (NAC) 21

 NetMotion Mobility XE Versus Legacy VPN Client 24

SUMMARY & CONCLUSIONS 25

TABLE OF FIGURES

Figure 1 – Mobility Server Management Interface3

Figure 2 – Mobility Client Software10

Figure 3 – Switching Between Networks12

Figure 4 – Switching To Mobile Data Network13

Figure 5 – Data Compression Test16

Figure 6 – JPEG Acceleration Test.....17

Figure 7 – Compression + Acceleration Test18

Figure 8 – Policy Rule Example.....20

Figure 9 – NAC Rule Creation.....22

Figure 10 – Client Session Details23

Figure 11 – Firewall Disabled NAC Example24

Figure 12 – Legacy VPN throughput compared to Mobility XE.....25

BROADBAND-TESTING

Broadband-Testing is Europe's foremost independent network testing facility and consultancy organisation for broadband and network infrastructure products.

Based in Andorra, Broadband-Testing offers extensive labs, demo and conference facilities. From this base, Broadband-Testing provides a range of specialist IT, networking and development services to vendors and end-user organisations throughout Europe, SEAP and the United States.

Broadband-Testing is an associate of the following:

NSS Labs (specialising in security product testing)
Limbo Creatives (bespoke software development)

Broadband-Testing Laboratories are available to vendors and end-users for fully independent testing of networking, communications and security hardware and software.

Broadband-Testing Laboratories operates an **Approval** scheme which enables products to be short-listed for purchase by end-users, based on their successful approval.

Output from the labs, including detailed research reports, articles and white papers on the latest network-related technologies, are made available free of charge on our web site at [HTTP://www.broadband-testing.co.uk](http://www.broadband-testing.co.uk)

Broadband-Testing can also provide technical writing services for sales, marketing and technical documentation, as well as documentation and test-house facilities for product development.

Broadband-Testing Consultancy Services offers a range of network consultancy services including network design, strategy planning, Internet connectivity and product development assistance.



EXECUTIVE SUMMARY

- With NetMotion Mobility XE we set out to prove the benefits of having continuous connectivity at the application layer, regardless of the data network connection.
- Running over Wired, WLAN and cellular data networks, we showed that the NetMotion Wireless solution allows a computer or PDA/smartphone user to maximise their work opportunity by making use of any available data network connection.
- The working model is pure client-server, easy to deploy and simple to understand.
- While complex beneath, on the surface – to the end user – the Mobility XE service is completely transparent.
- However, NAC and Policy Management options enable users to be controlled to a very finedegree.
- Using these tools, all traffic and data conditions can be optimised and maximised, users and applications prioritised. For example, the Mobility client will automatically move the user to the highest bandwidth, most cost-effective or whichever network type has been prioritised, as and when it sees a new network option – again transparent to the user.
- With a combination of NetMotion Mobility's data compression and protocol optimisation we saw up to a 500% improvement in performance when downloading from and surfing the Internet.
- Using NetMotion Mobility Policy Management we were able to deny access to certain applications depending, for example, on what the available bandwidth/network conditions were.
- Using NAC we were able, for example, to deny a user access to the Internet if their personal firewall was disabled.
- Initial testing versus Cisco and IPsec VPNs suggests that the NetMotion Mobility XE technology is almost 300% faster than its legacy equivalents. Broadband-Testing is following up this initial testing with a second report focused on Mobility XE performance versus legacy/alternative solutions.

INTRODUCTION: THE VALUE OF KEEPING CONNECTED

Performance isn't everything.

OK – so fast is better than slow and faster is better still. But, as with the tortoise and hare fable, consistency gets you there quicker in the end than does sporadic huff and puff. In other words, in much the same way that a 1970's Toyota Corolla would often cover 1,000 miles more quickly than an equivalent age Ferrari, because you didn't have to keep stopping for petrol, or the AA man, so always being connected to your applications and data is often more valuable than high performance in batches with nothing in between.

Such is the line of thinking behind the kind of always on connectedness that NetMotion Wireless delivers with its Mobility XE software solution. In other words, whatever your connection type – wired, WLAN, WWAN, mobile – your laptop/PDA/smartphone stays in sight of your applications and data, maintaining a virtual connection even when no physical connection is available.

Intriguingly, analysts still talk about the idea of ubiquitous enterprise mobility – for an employee to be always connected wherever and whenever – as something that *will* happen, not something that actually has. Take Forrester for example, an analyst group that follows this space very closely. Last year Forrester released a report where it sees the combination of public cellular and Wi-Fi technology as the harbinger of a new mobile network that will be a combination of short-range unlicensed technologies operated by their users and users IT staffs and carrier-based solutions for connectivity outside of the enterprise.

While this much may be true, why do we need to wait for it to happen when we already have vendors such as NetMotion Wireless offering this? But to state that this is a “going to happen” event when the technology prevails, strikes a false note. To quote from a Forrester report released last year: “While the widespread availability of solutions coming from familiar vendors is at least five years off, understanding what makes up the ubiquitous infrastructure and how to plan for the multi-network future is a must for enterprises looking to realise returns on IT investments happening now and during the next five years.” Clearly Forrester hadn't spoken with NetMotion Wireless then!

So, in this report we examine what the NetMotion Wireless technology is, what it does and whether it delivers what it promises. Read on...

NETMOTION MOBILITY XE: A PRODUCT AND SERVICE OVERVIEW

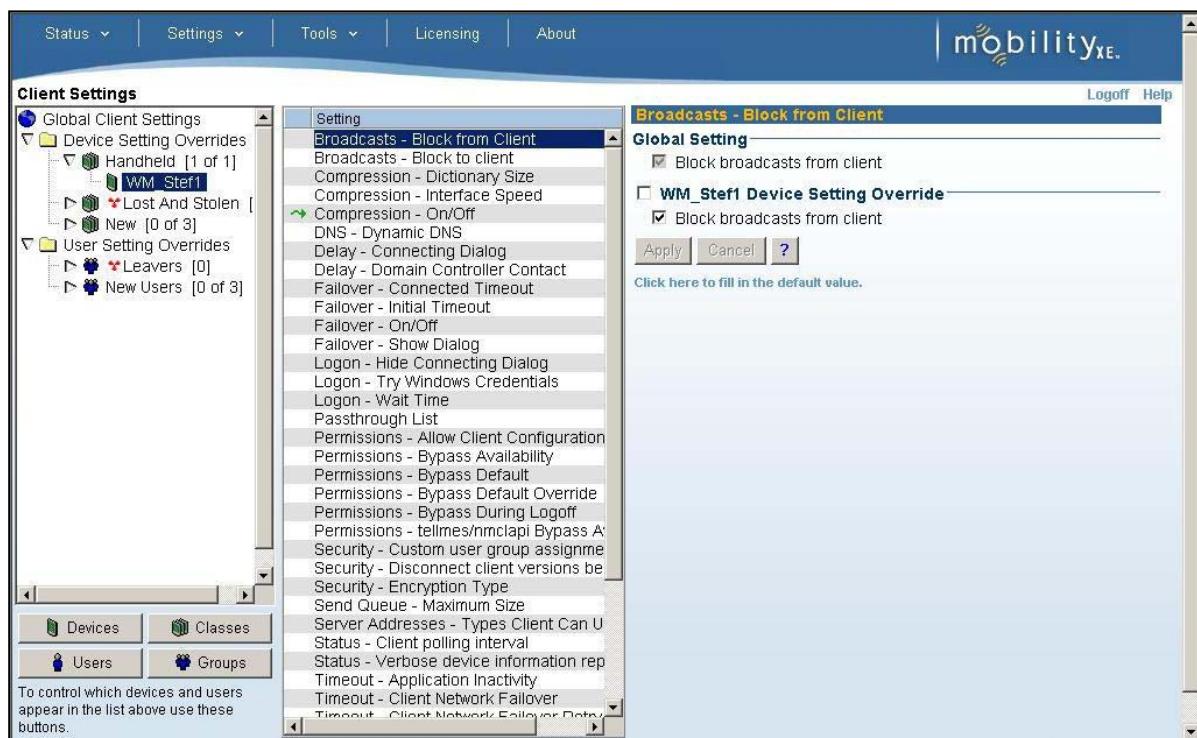
What Exactly Is Mobility XE? In A Nutshell...

Mobility XE is a software-based, mobile VPN that works within a standard network infrastructure and is designed to be highly scalable and resilient. For the latter it supports both active/active and active/passive redundancy.

There are two primary components – the Mobility server and the Mobility client. Communication between the two is via RPCs (Remote Procedure Calls) and the IMP (Internet Mobility Protocol) which runs on top of UDP (User Datagram Protocol).

The Mobility server manages wireless network connections for mobile devices and runs on either Windows Server 2000 or Windows Server 2003. It acts as a transport-level proxy

for each mobile device running the Mobility client, maintaining the state of each client and handling the session management required to maintain continuous connections to running applications. When a mobile device becomes unreachable, suspends operation, or moves



to a different network, the Mobility server maintains the connection to the client's peer applications by acknowledging receipt of data and queuing requests. The Mobility server also manages network addresses for the mobile devices.

Figure 1 – Mobility Server Management Interface

A Mobility client device is defined as a standard mobile device or off-the-shelf computer

running Windows 2000, Windows XP, Windows Vista or Windows Mobile. Each Mobility client receives a virtual IP address on the wired network, typically obtained from DHCP (Dynamic Host Configuration Protocol), or assigned from a range of addresses reserved for this purpose on the Mobility server. Mobility also supports the ability to statically assign a virtual IP address to an individual device or user.

Multiple Mobility servers can operate as a server pool with failover and load balancing capabilities. The Mobility server also provides tools and metrics that a systems administrator can use to configure and manage remote connections and troubleshoot

remote connections. A web-based interface allows system administrators to configure

Mobility XE settings and manage the server from another PC on the network.

How Does It Work?

The Mobility client software resides at the transport driver interface layer (TDI) on supported Microsoft platforms and performs indirection and redirection of application network calls. When an application wishes to use the network, the TDI calls are intercepted, the parameters are marshalled, and the call is forwarded for execution on the

Mobility server. It works transparently with operating system features to allow client-side

application sessions to remain active when the device loses contact with the network.

Each Mobility client has a virtual IP address on the wired network, obtained from DHCP or assigned from a range of addresses reserved for this purpose on the Mobility server. In addition, static virtual IP addresses can be assigned to specific devices or users.

For each active client, the Mobility server relays data directed to the client's virtual address to its current, actual address (the point of presence - POP - address). While the POP address of a Mobility client may change when the device moves from one coverage area to another, the virtual address remains constant for the duration of the user session.

Mobility XE's Remote Procedure Call (RPC) protocol and Internet Mobility Protocol (IMP) form the technological backbone that connects the Mobility server to each mobile device. A remote procedure call is a way of allowing a process on a local system to invoke a procedure on a remote system. With Mobility XE, the client's network calls are sent to the server for remote execution. If Mobility operated at the Winsock layer these would be calls such as open socket, bind, connect, send and receive. Because Mobility operates at the TDI layer, the TDI equivalent of these calls is what is sent to the server for remote execution.

The application on the local system is not aware that the procedure call is executed on a remote system. Mobility XE's RPC-style approach means that it allows the mobile device

to go out of range or suspend operation without losing active network sessions. Because session maintenance does not depend on customizing or rewriting applications, off the shelf applications will run without modification in the wireless environment. The RPC protocol is encapsulated by the IMP) which is encapsulated in UDP. IMP compensates for differences between wireline and less reliable networks by adjusting frame sizes and protocol timing to reduce network traffic. This is important when bandwidth is limited, there is high latency, or when battery life is a concern.

Mobility XE also strengthens data security by encrypting all traffic between the Mobility server and clients, and allowing only authenticated devices to connect to the Mobility server. When a Mobility client connects to a Mobility server for the first time, the server registers the mobile device's permanent identification (PID) number, a unique number that the client will use for all subsequent connections. This registration occurs only on the first connection and does not require any action by the user or administrator. The identification number is stored in the client system registry and in the Mobility warehouse (LDAP).

The Mobility server stores the PID based on the computer name. As long as the client computer name does not change, the server can restore the PID to the client device even if the client registry is lost. This may happen, for example, if the client device's hard drive

is re-formatted in order to re-install the operating system. When the Mobility client is

re-installed and re-connects, the server then searches for a match on the device name. If

it finds a match, it will restore the same PID to the device.

When the Mobility client establishes a connection to the Mobility server, it requires the user to authenticate. If authentication is successful, the server and client create the secure VPN tunnel that will be used for the duration of the session. The mobile worker can use his or her standard Windows logon credentials to authenticate to the network. The Mobility server authenticates that user against the enterprise's domain using NTLMv2 for native Microsoft authentication, RADIUS authentication, or RSA SecurID authentication.

For example, in a Microsoft deployment, a three-way handshake occurs between the

Mobility client and server:

- The client sends a list of the supported authentication types. This packet includes the NTLMv2BLOB.
- The server responds with an NTLMv2challenge.
- The client completes the authentication with the response to the challenge.

After authentication, the server and client exchange signed ECC (elliptic-curve

cryptography) public keys and related cryptographic materials to perform the

Diffie-Hellman key exchange. Symmetric keys are derived from the public keys; these are

not transmitted. This applies to all supported authentication methods.

A vital element of the Mobility software is the session persistence at the Application Layer. Unlike IPsec VPNs or SSL VPNs, the Mobility XE VPN does not require a fixed local address. The tunnel is maintained between the Mobility server, which is at a fixed

address, and the Mobility client, which can have an ever-changing POP address. By

mutual agreement, the client and server maintain the secure tunnel until one endpoint issues a disconnect: this could happen when the user logs off, the administrator quarantines the device, or due to a configurable link inactivity timeout.

The tunnel remains available and application sessions persist in any number of scenarios:

- Suspending operation on the mobile device and later resuming it.
- Moving to a different location on the network.

- Connecting a mobile device over slow, bandwidth-challenged, or high-latency networks.
- Encountering interference from microwaves, stairwells, elevator shafts — anything that interferes with radio signals.
- Changing network interfaces (for example, from a WLAN to a WWAN card).
- Moving across gaps in coverage.

The configurable timeout ensures that the resources on the Mobility server consumed by inactive sessions are not consumed indefinitely. But, according to NetMotion, in test scenarios, devices have been suspended in the middle of an application transaction, awakened a week later, and the transaction resumed exactly where it left off.

Here is why it is deemed that application persistence is so important. Network applications are written to use application-level interfaces, such as Winsock (the Windows sockets API). A single call to the application-level API may generate several outgoing or incoming data packets at the transport (IP) layer. In existing mobile networks, if one of these packets is lost, the state of the entire connection can become ambiguous and the session may be dropped.

The IMP overcomes these issues. The protocol emulates wireline behaviour and handles the reliable delivery of data while dealing with various scenarios, including connections lost in the middle of a transmission and roaming. The logical session is maintained and held open, even if the device becomes unreachable or is suspended.

From the standpoint of the application on the mobile device, the application simply waits until it receives a response unless it maintains its own timeout. Mobility XE accomplishes this whenever the server is unreachable, by setting the operation to a pending state. The Winsock interface then returns the appropriate response to the application, whether it is using blocking calls, asynchronous calls or overlapped I/O. The application encounters the same responses and behaviour it would encounter on a wired network.

Mobility does not 'test' to see if a device is reachable before sending—it simply sends the required data and looks for a return acknowledgement. If the return acknowledgement

isn't received, it will re-send for a number of tries, then back off, conclude the device is

unreachable, and re-send at a later time. The IMP state machine keeps track of packets

sent, those acknowledged, and those that need to be re-sent, preserving the integrity of

the entire session.

To determine whether an inactive mobile device is reachable, the Mobility system uses

keep-alives: the Mobility client periodically sends frames to the Mobility server. The

frequency of these keep-alive frames is user-configurable and may be decreased to

reduce traffic on bandwidth-constrained networks. If the Mobility server fails to receive

expected keep-alive frames during the configured timeframe, it indicates that the device

is unreachable in the server's management console. Keep-alives are sent only in the

absence of application or other traffic.

So what if a device is unreachable? Mobility's use of RPC preserves the state of interrupted data transfers and holds pending data in queue. *Mobility Preserves the VPN Tunnel, Application & Data even when the client is unreachable.* When the application server is transmitting data to the Mobility client and the client becomes unreachable due to a disruption in the network or the suspension of the device, Mobility's flow control algorithms notify the Mobility server's RPC layer to stop accepting data from the application server.

The Mobility server's TCP buffers fill, resulting in the TCP window size adjusting to zero. This, in turn, notifies the application server that data transmission should be paused until the Mobility client is able to receive data. In this state, the Mobility server and application server exchange TCP acknowledgements to keep the connection alive indefinitely or for a length of time configured by an administrator.

When the client becomes reachable again, the RPC layer receives a flow control indication notifying it that the data transfer may continue. The RPC layer then resumes taking data from the TCP stack, causing the TCP connections receive buffers to empty. When buffer space becomes available, the TCP stack sets the connections receive window to a

non-zero value indicating to the application server's TCP stack that it is again ready to

receive data and the transfer continues.

And Roaming?

Mobility uses DHCP to detect and facilitate “roaming”—when mobile devices move across a subnet boundary or to a different network.

The user does not need to restart the system or close existing network connections to obtain a new address. Mobility allows mobile devices to cross subnet boundaries without losing their network connection or application sessions. To do this, the Mobility client must be able to detect that it has moved to a new subnet and must be able to obtain a new point of presence address from a DHCP server.

The Mobility client uses two methods to detect that it has moved to a new subnet. One method is based on information from the network card driver, and is the preferred method. The other method is based on DHCP beaconing, and involves periodic DHCP discover/offer exchanges. Both methods of subnet roaming use DHCP services, at least in part, to perform roaming. The Mobility client uses DHCP services to acquire its IP address and other optional configuration parameters for the new network or subnet.

On a WWAN, the network infrastructure handles roaming; beaconing is not required. Beaconing is automatically turned off if either of the following is true:

- The Mobility client is connecting over a dial-up (DUN) or PPP (Point-to-Point Protocol) network adapter.
- The Mobility client’s network adapter acquires a static IP address and does not need to change its IP address in order to roam in the wireless WAN (*e.g.*, roaming is enabled in the WWAN’s underlying infrastructure).

Mobility provides a setting that modifies beaconing-based roaming detection in a super-netted network. This is a network where two conditions exist:

- A single physical network supports two or more logical networks, and
- DHCP servers assign IP addresses to Mobility clients on more than one logical network

If both of these conditions apply and if the network environment cannot be readily modified to simplify IP address allocation to Mobility clients, enabling super-netted

roaming on the Mobility server will improve performance of beaconing-based roaming by

preventing frequent attempts to roam when it is not necessary.

A Mobility client can maintain a connection and application sessions when moving between different network media. If network interface cards for different types of network connections have been installed and properly configured on the mobile device, Mobility XE provides application persistence as the client moves between a LAN, a wireless LAN, and a wireless WAN, or any other type of IP-based network. No additional configuration is required on the Mobility client or server.




Figure 2 – Mobility Client Software

The operating system and network hardware on the mobile device determine how much user intervention is required for Internetwork Roaming. A Mobility client with multiple active network interfaces may be able to switch from one interface to another automatically. For example, if a client device with both WLAN and WWAN cards goes out of range of any wireless LAN access point, communications may automatically transition to the WWAN medium.

On a client device with multiple network interface cards simultaneously active, Mobility XE uses whatever interface the operating system is using. The operating system’s interface selection may depend on the order in which the interfaces become active, and may not always result in the use of the “best” interface. The Mobility server modifies the metric for a defined route in the client operating system based on the speed of the network interface, so the mobile device uses the available interface with the greatest bandwidth.

If the Roaming—Use Fastest Interface option is disabled, the mobile user may have to manually stop and start interface cards, depending on the operating system.

Sometimes, network connections are available via two interfaces, but the Mobility Server can be reached only via one of these networks. Mobility XE's Client Network Failover allows the client to connect even when the preferred interface is available, but cannot connect to a Mobility Server.

How Mobility XE Provides Link Optimisation

There are many characteristics of TCP/IP behaviour that make it less than optimal in a wireless environment. To address these, Mobility XE is designed to provide optimum

performance over intermittent and bandwidth-challenged network links. Its architecture

includes enhancements that allow network traffic over IP to deal more effectively with momentary loss of connectivity from a mobile device, whether due to coverage outages or external factors, such as power management or user intervention. It makes the most efficient use of the given bandwidth using many advanced features that reduce the "chattiness" of transport protocols:

- Selective acknowledgments.
- Data and acknowledgment bundling.
- Message coalescing.
- Reduced and synchronised retransmissions.
- Fragmentation optimizations.
- Data compression.

- Error-reduction algorithms.

- Web acceleration.

Working in concert with one another, these advanced features provide for the most efficient movement of data. In addition, when "Use Fastest Interface" is enabled (the default), Mobility XE automatically switches to the fastest bandwidth network connection when multiple connections are active.

Some of these key elements of Mobility XE feature in our "put to the test" section which follows...

NETMOTION MOBILITY XE: PUT TO THE TEST

Roaming

For the roaming testing we carried out two lots of testing; in the LAN environment and in the mobile world. We switched between wired LAN and WLAN, as well as between different WLANs with different bandwidth/power settings, as well as between mobile cellular operators with no issues.

The screenshot below shows how we maintained connectivity switching between networks as we disabled then enabled various combinations; a single response took longer than standard but the connection remained intact at the protocol and application layers.

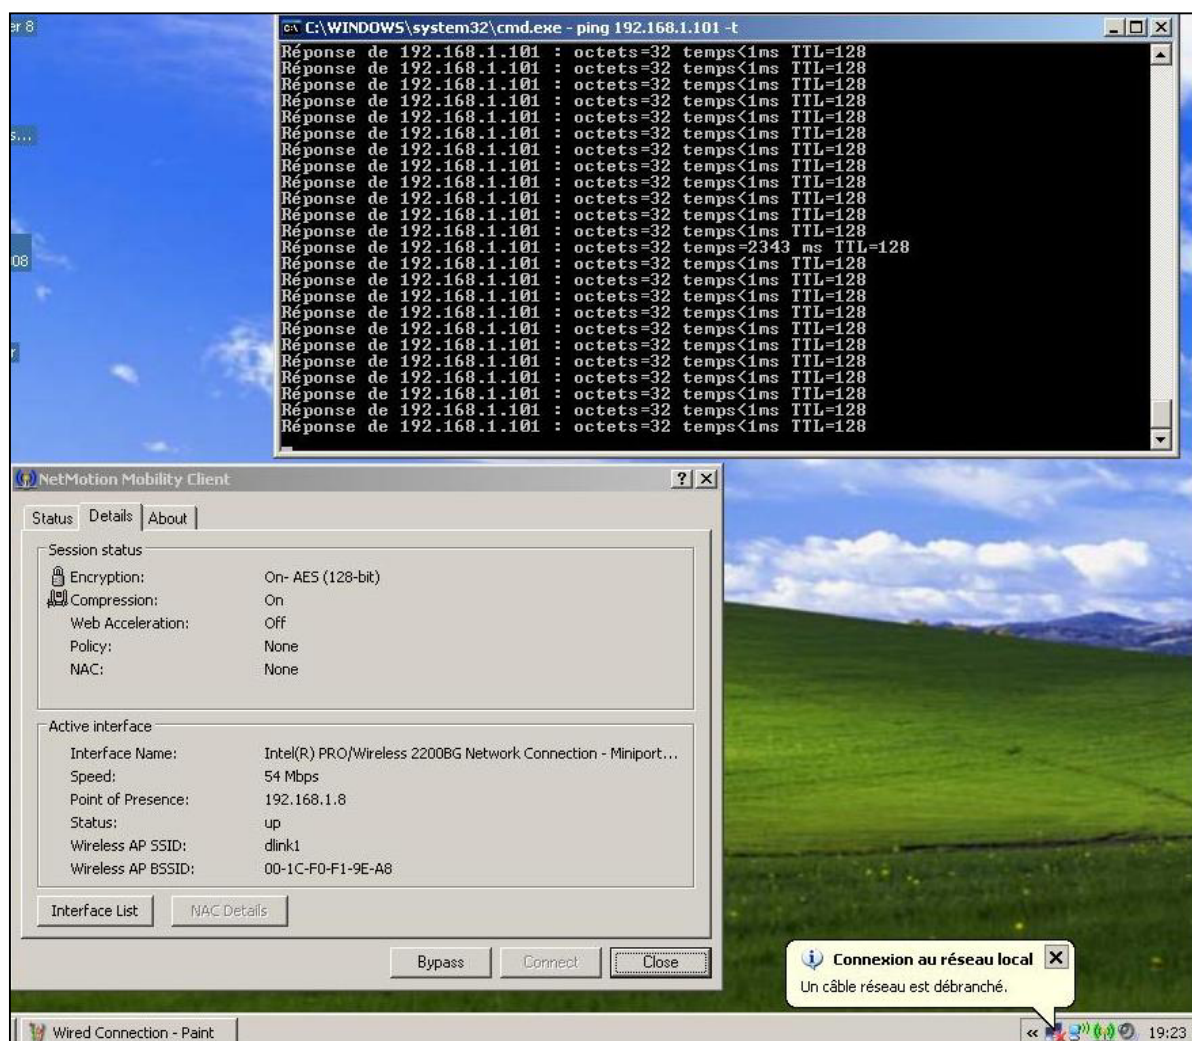


Figure 3 – Switching Between Networks

We also confirmed mobile data operation with online Internet browsing maintained when switching to a mobile data source.

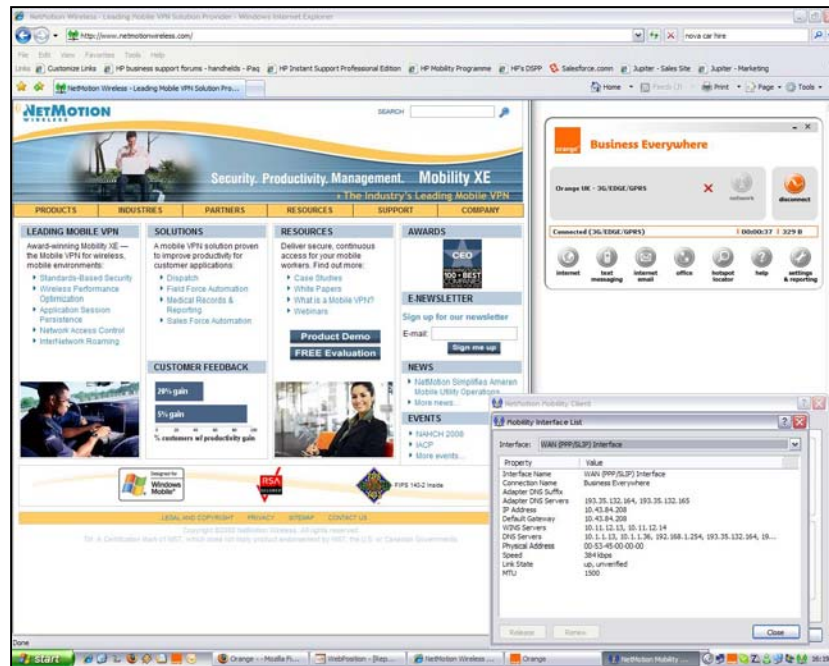


Figure 4 – Switching To Mobile Data Network

UDP vs TCP/IP

TCP/IP's packet sizes are not always optimum for wireless transmission. In a wireless environment, error rates rise as transmission power drops, making errors much more likely in marginal coverage areas. Sending large packets in these situations increases the

probability that an entire packet has to be thrown away and re-sent. Smaller packets may

increase overall efficiency by decreasing the number of re-sends. Many network

administrators are unaware that even WLAN's can have substantial numbers of dropped packets and packet errors.

The UDP protocol is much more appropriate for use over wireless networks. It avoids the overhead and inefficiencies of TCP, which was not designed with wireless networks in mind. Instead, Mobility's Internet Mobility Protocol (IMP) rides on top of UDP, and implements its own methods for dynamically adjusting both packet sizes and timing parameters for the network conditions. IMP, used in conjunction with UDP, handles the far greater variety of transmission speeds and connection conditions encountered in wireless networking. Using this approach also allows IMP to apply its own compression

and link optimisations, which can double the throughput over bandwidth-constrained networks.

Because the UDP protocol is connectionless, Mobility's Internet Mobility Protocol handles the job of guaranteeing reliable data delivery. It uses its own algorithms for selectively acknowledging packets, handling timeouts, detecting dropped packets and retransmitting them. It is far more sophisticated than TCP/IP—and indeed, it has to be. It not only has to verify delivery amid the uncertainties of a wireless environment, but also has to do so with minimal overhead and with as few retransmitted frames as possible without over-consuming limited bandwidth.

Data Compression

With Mobility XE it is possible to customise compression functionality and determine when it should be turned on, and whether it should be enabled globally (for all users and devices), for a mobile device class, a specific device, or an individual user. Or you can configure Mobility to automatically switch compression on or off, based on the current

interface speed. This means users can roam between high-bandwidth 802.11b LANs and

lower-bandwidth GPRS networks and automatically get the best performance possible.

How Does It Work?

Mobility XE employs the standard algorithms outlined in RFC 1951 (LZ77 Deflate/Inflate). Only the application payload of each frame is compressed — the transport headers are not modified. This allows Mobility to operate through any policy enforcement equipment, such as firewalls and network address translators (NATs). It is applied to all

application-level data that traverses the Mobility tunnel. No modification or re-

configuration of the application is necessary to take advantage of this functionality.

As Tested

On a WLAN with 1Mbps Uplink/4Mbps Downlink ADSL Internet access we accessed a text file on the performance.toast.net website, and downloaded (using Yahoo as the host) the file – 341KB – loaded in 4.454 seconds. With compression then enabled and the local cache flushed, we repeated the download and saw a total download time of 1.906 seconds and an effective reduced file size of course.




Figure 5 - Data Compression Test

Web Acceleration

To speed up web browsing on slow networks, Mobility gives you the option of compressing images. The level of compression is configurable, and you don't have to sacrifice Mobility XE's mobile VPN security. When enabled, all HTTP traffic on the designated ports will have the images compressed at the configured level.

How Does It Work?

Web acceleration is available in two different places:

Policy Management module: Using policies you can selectively turn web acceleration off and on, change the level of compression, or change the HTTP ports, based on the current network characteristics, on a specific application, or any of the other available conditions.

Client Settings: Web acceleration is available in the core product without additional licensing.

As Tested

We used the performance.toast.net website again and downloaded a JPEG graphic of the US air force Blue Angels (4.39MB) while connected to the WLAN with a very weak signal (1Mbps). The initial download took 129.49 seconds. With web acceleration then enabled we repeated the download and download time was reduced to 44.64 seconds – a 300% improvement approximately.

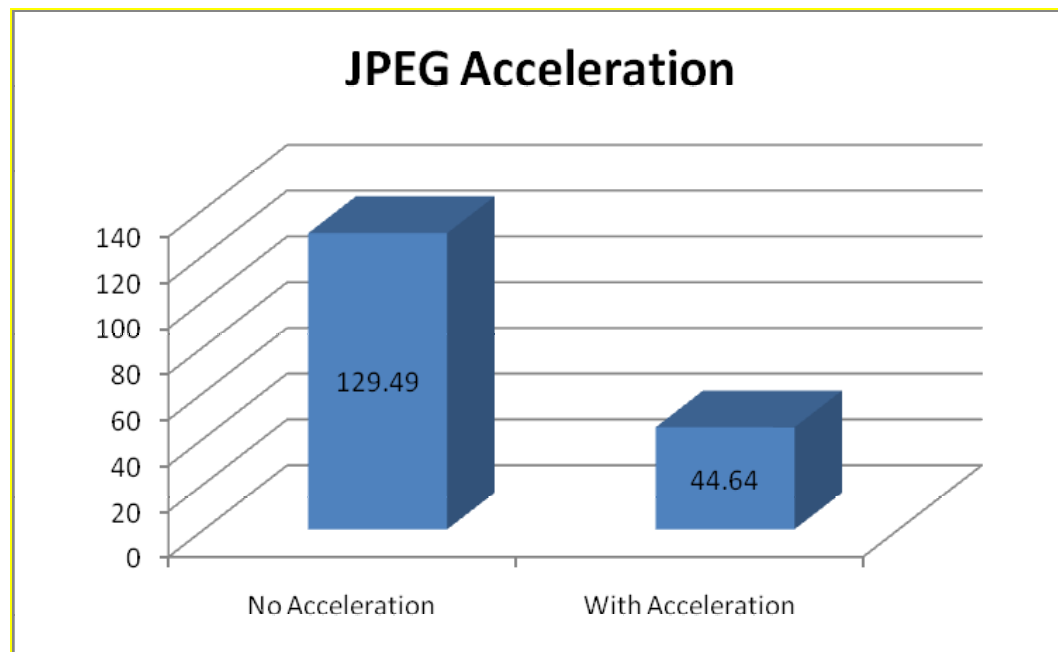


Figure 6 – JPEG Acceleration Test

We then combined compression and acceleration with a combined text/graphics download on the same WLAN link. With both compression and acceleration disabled, the download time was 29.3 seconds. With both enabled, this download time was reduced to 5.45 seconds – almost a 600% acceleration.

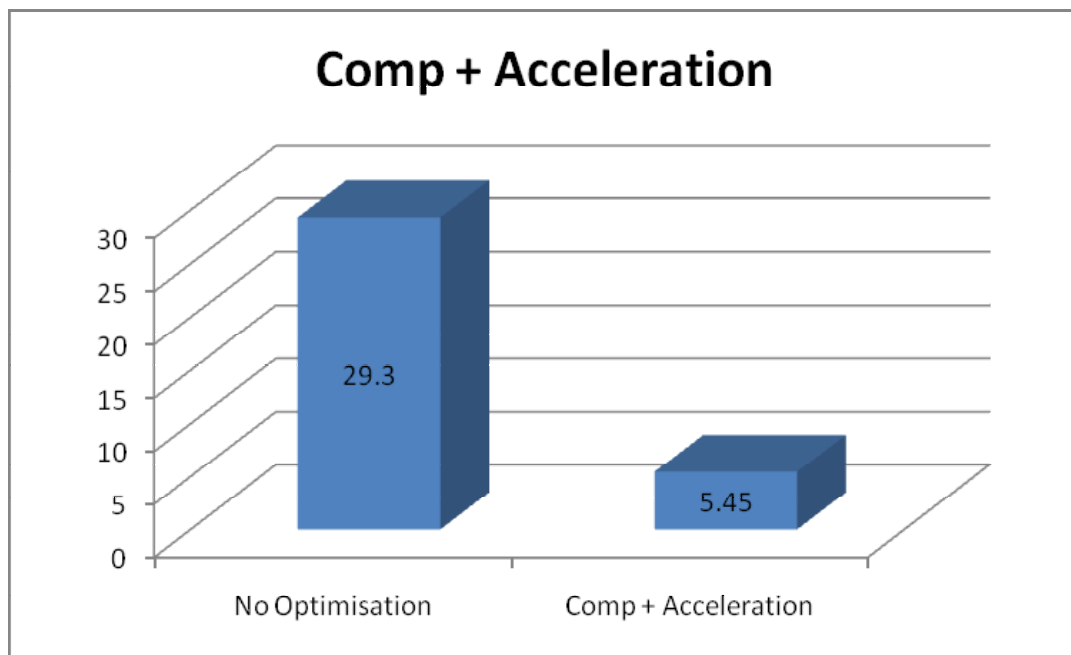


Figure 7 – Compression + Acceleration Test

Policy Management

The Mobility server maintains user policies and pushes them out to the Mobility clients in the field where they are enforced. Policy documents are stored in the Mobility warehouse. All Mobility servers in a server pool share a common set of policies. Mobility policies are pushed to the mobile device when it connects. These policies can be defined to be specific

to the device or the user. Policy updates or modifications are applied in real-time after the

administrator publishes them for distribution.

How Does It Work?

For storing, retrieving, and validating rules between client and server, Mobility uses an

XML-based rules language. A policy document contains an XML definition of the policy, as

an ordered list of rules referenced from the rules documents (a rule or set of rules that

control client network behaviour). The Mobility server parses the verbose server-side XML

into usable code objects, and validates rules against an XSD (XML schema definition)

when saving or opening the generated document. When rules are created and saved, it performs an XSL transformation from verbose server-side XML to the client-side format, creating a resource-efficient policy document containing only the information needed by each client.

The policies allow extremely flexible and fine-grained control over user and device access to network resources. While the rules are enforced at the device level, the human-readable rule sets are maintained at the Mobility server. The basic Policy

Management actions that have an impact on connections are: allow, block, disconnect, and pass-through. Mobility enforces these actions at the client. When an application attempts to send over the network, the Mobility client checks the policy list for the application, port, destination address and other parameters to see if action should be taken. The descriptions, below, assume that actions described are the base action for the policy.

Behaviour may change if the action is invoked after a session has been activated. Policy Management makes granular management of wireless bandwidth, security, and mobile productivity — even over networks the administrator neither owns nor controls — straightforward and achievable.

As Tested

We tested the software by creating a rule that blocked access to Mozilla Firefox and Skype at connection speeds below Fast Ethernet. We first connected the client via a wired Gigabit connection and access to all applications was allowed. We then disabled the wired connection, forcing the client to a wireless (802.11g) connection. While all connectivity with the Internet was maintained, we then found we could not load Skype or Firefox on that client, so the rule we created worked perfectly.

Policy Management - Rules

Cancel < Back Next > Finish **Edit rule > Conditions(s) [SBrule]**

Which condition(s) do you want to check?

Select the condition(s):

Addresses

- When the local address is address
- When the WINS server address is address
- When the DNS server address is address (Windows Vista and XP only)

Interface

- When the interface name contains keyword
- When the interface speed is less than speed Kbps
- When the DNS suffix contains keyword (Windows Vista and XP only)
- When the connection name contains keyword (Windows Vista and XP only)

Registry Key/Value

- When the Registry value is

External Key/Value

- When the External Condition is

Network Access Control (NAC) Status

Edit the Rule Description (click an underlined value):

Apply this rule

- when the interface speed is less than 54000 Kbps
- block network traffic for application(s)
skype.exe
- and all address(es)/port(s)
- with options else
- continue to the next rule

Figure 8 – Policy Rule Example

Similarly, QoS can be applied via Policy Management to provide fixed amounts of bandwidth for specific applications.

QoS

Mobility XE implements sophisticated Quality of Service (QoS) via Policy Management and integrates with DSCP [Differentiated Services Code Point]. This support can be crucial to maintaining productivity as workers move from high-speed, high-bandwidth networks, to lower capacity, high latency networks. For example, while connected to the LAN via Ethernet, performance may be just fine for the mission critical enterprise application,

running alongside e-mail, web browsing, and other applications. But on a WWAN, administrators want to prioritise use of the narrower bandwidth, and make sure that a web browser and e-mail client do not use capacity needed by the enterprise application.

Other VPNs may allow administrators to shut off non-essential applications.

Mobility XE allows administrators to specify QoS parameters which are applied between the Mobility client and server, and additionally put DSCP settings which can be used as network traffic moves beyond the Mobility Server. In particular, Mobility XE allows for different settings based on the network and its characteristics — there can be one set of QoS rules applied to an ADSL network, another to mobile data, etc. This allows an enterprise application to always have the appropriate share of network resources, but other, non-critical applications to use whatever bandwidth is available once the enterprise application has used what it needs or is assigned.

Network Access Control (NAC)

The NAC module gathers information on antivirus, antispyware, firewall software, Windows updates and registry, installed files, and processes running on the mobile device. NAC security checks enforced by the Mobility server use this information to assess the health of the Mobility client. For clients that fail a security check, NAC rules provide users with the information they need to bring the client device into compliance.

How Does It Work?

The NAC module is centrally managed using the Mobility Console. The network administrator creates a “rule“ that can check multiple device attributes. Groups of rules, known as a ”rule set“ are combined to create NAC policies suited to the organization’s needs.

Using a ‘NAC wizard,’ rules or rule sets are established to enforce security policies globally, to workgroups, by class of device, or to individual users and devices. NAC updates are automatically sent to all subscribed users and devices. The NAC rule set is

evaluated by the Mobility XE client software at start-up, and also re-evaluated every five minutes (this interval is configurable). If a client device fails a NAC policy check, based on severity or corporate policy, the administrator has a number of options for responding and remediating. If the infraction is not serious, they can configure a failure message that explains what the user must do in order to bring their mobile device into compliance. For serious security infractions, the mobile device can be disconnected from the corporate network or quarantined entirely.

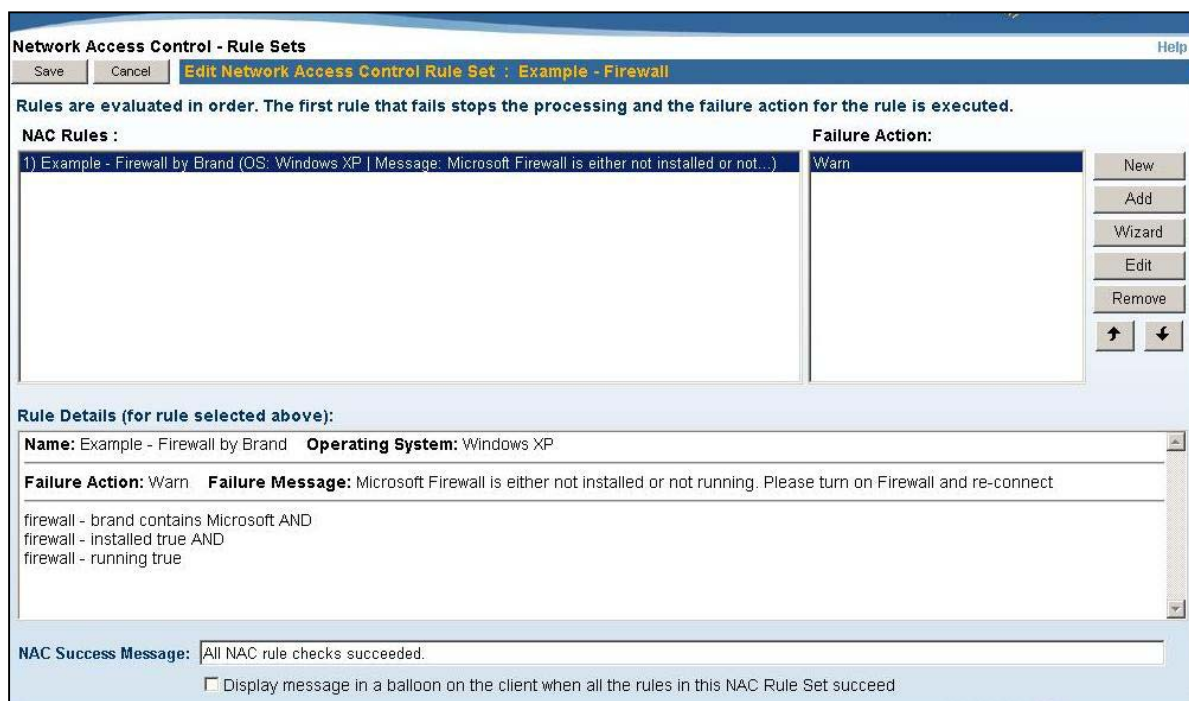


Figure 9 – NAC Rule Creation

By using the integrated Policy Management module in tandem with the NAC module, an IT administrator can also require that Mobility clients automatically download and install software updates, operating system patches, Mobility XE updates, etc., in order for the device to be considered “healthy” and thus capable of connecting to the corporate network.

Network administrators define rules that are evaluated on each client device with enforcement occurring at the Mobility server. Administrators also determine the severity of enforcement, from warnings, to remediation (using the Policy Management module), to a disconnect or quarantine.

Allow: The Mobility client device complies with NAC policy. Inbound and outbound network traffic allowed via the Mobility VPN through the Mobility server.

Warn: The client does not comply with one or more checks in a rule that causes the Mobility client to display a warning.

Remediate: The client does not comply with one or more checks in a rule that requires remediation. The action required to bring the client into compliance is determined by the system administrator.

Disconnect: The client does not comply with one or more checks in a rule that causes the device to be disconnected.

Quarantine: The client does not comply with one or more checks in a rule that causes the device to be quarantined. The system administrator must clear this state before the device can connect.

The Mobility server also collects a number of different details from the active device's session:

< Back Connection List > Session Details	
<input type="button" value="Refresh"/>	
Session Item	Value
Device Name	NEWACER
User Name	THECIRCUS\smbroadhead
Status	Connected
Device Description	
Device Class	New
Device ID	01C922290D39AA3A000475A0A62F002
Server	server
Virtual Address	192.168.1.10
POP Address	192.168.1.9:52700
Interface	Broadcom NetLink (TM) Gigabit Ethernet
Interface Speed	100 Mbps
Bytes Sent	846,812
Bytes Received	4,211,727
Registered time (UTC)	2008-09-29 13:46:46
Connection established	2008-09-29 13:46:54
Total connect time	0d 02:22:59
Client version	8.00.54422
Client operating system	x86 running Windows Vista version 6.0 Build 6001 Service Pack 1
Client docking	Undocked
Power source	External
Battery	100%
Network Access Control Level	n/a
Network Access Control Description	n/a

Figure 10 – Client Session Details

Upon initial connection, the Mobility client sends the session details to the Mobility server, which populates the information on the session details page of the Mobility Console (for each connected device).

Once the state of the connection is recorded on the details page, the information is updated when the information changes or on a periodic basis. For example, the POP address details will only change on a roam event. When the mobile device acquires a new POP address, the Mobility client will inform the Mobility server of the new session detail—no other details will be sent if they have not changed. Information that requires more frequent updating, such as the state of battery life on a device, is updated based on a configurable timer.

Because the Mobility server proxies the application traffic on behalf of each Mobility client, the bytes transmitted per application (process) can be derived from the Mobility server without having to burden the client or the network with transmitting this information.

As Tested

We put it to the test by creating a rule that checked for a firewall to be enabled on a client before it would provide a connection. We then disabled the firewall on the test client and attempted to make the connection. It failed. We then enabled the firewall and tried to connect again. This time it allowed the connection.



Figure 11 – Firewall Disabled NAC Example

By using the Policy Management module in tandem with the NAC module, an IT administrator can also require that Mobility clients automatically download and install software updates, operating system patches, Mobility XE updates, etc., in order for the device to be considered “healthy” and so capable of connecting to the corporate network.

NetMotion Mobility XE Vs a Cisco and IPSec VPN Client

As the initial part of validating a series of tests with NetMotion Wireless versus legacy/alternative technologies, we compared Mobility XE performance against a Cisco VPN client (legacy client) and and IPSec client, BT’s iNet.

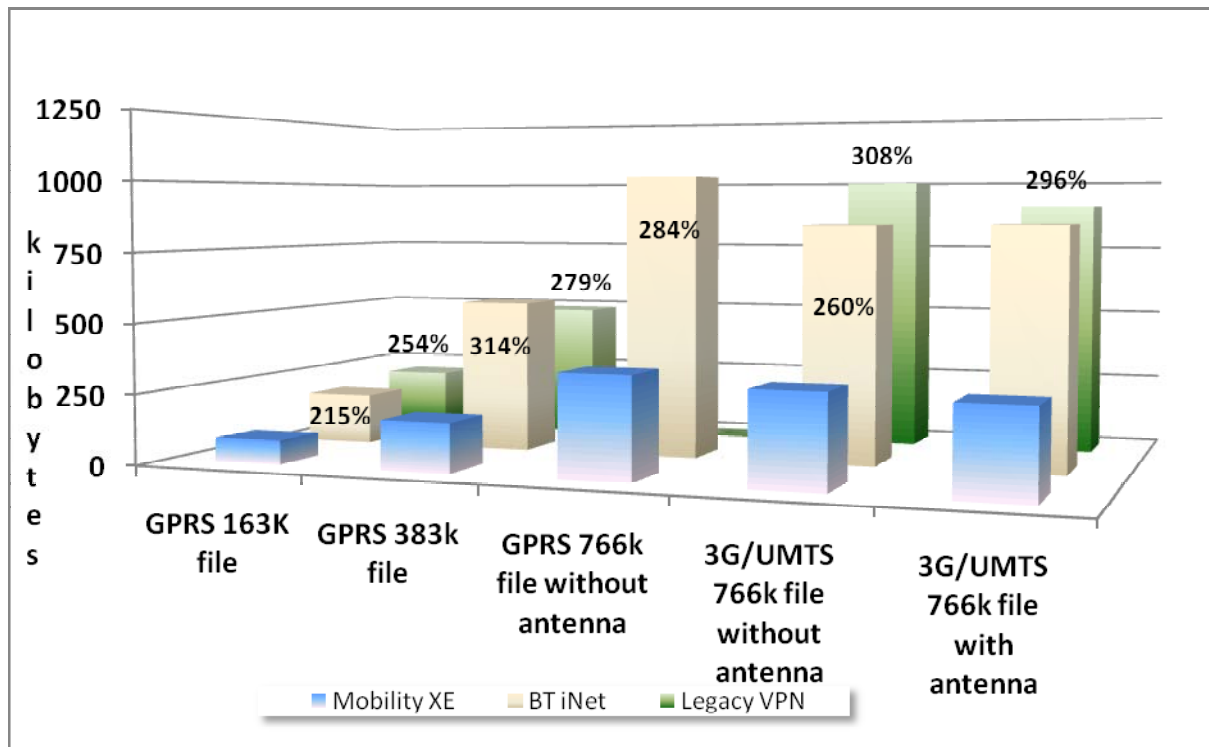


Figure 12 – Legacy VPN throughput compared to Mobility XE

Comparative throughput tests were performed to measure the effectiveness of NetMotion Mobility XE as compared to a Legacy VPN client, and to BT iNet, striving to keep environmental variables equitable. Using three different file sizes we found significant performance improvements using the Mobility XE solution on both GPRS and 3G networks.

On average, the NetMotion solution was almost 300% more effective. The Legacy VPN required 284% more cellular traffic (2.8 times as much) and BT iNet 270% more traffic to transmit the same files when compared to NetMotion Mobility XE. It is also important to note that, in some tests, the Cisco VPN could not complete the test transfer, despite repeated attempts (up to 12 repeat tries) since it simply could not sustain the connection.

Note: In our next report on Mobility XE we will be carrying out further tests to validate performance against legacy/alternative solutions.

SUMMARY & CONCLUSIONS

With Mobility XE we set out to prove the benefits of having continuous connectivity at the application layer, regardless of the data network connection.

Running over Wired, WLAN and mobile data, we showed that the NetMotion solution allows a computer or PDA/smartphone user to maximise their work opportunity by making use of any available data network connection. While complex beneath, on the surface – to the end user – the Mobility XE service is completely transparent, yet NAC and Policy Management options enable users to be controlled to a very finite degree. Using these tools, all traffic and data conditions can be optimised and maximised, users and applications prioritised. For example, the Mobility client will automatically move the user

to the highest bandwidth, most cost-effective or whichever network type has been prioritised, as and when it sees a new network option – again transparent to the user.

With a combination of NetMotion Mobility data compression and acceleration technologies we saw up to a 500% improvement in performance when downloading from and surfing the Internet. Using Policy Management we were able to deny access to certain applications depending, for example, on what the available bandwidth/network conditions were. Using NAC we were able, for example, to deny a user access to the Internet if their personal firewall was disabled.

Early testing against legacy alternatives suggests that Mobility XE is three times more efficient on both GPRS and 3G networks. More test data will follow in our, soon to be released, second report on the NetMotion product.

Overall then, NetMotion Mobility XE fully met our expectations of optimizing and controlling user activity on LAN, WAN and Internet, across all network media types. Given the ease with which we were able to both increase performance and enforce user and application control, while ensuring the user stayed connected at all possible times to their applications and services, it is easy to believe the NetMotion claims that ROI can be achieved in a matter of weeks.

