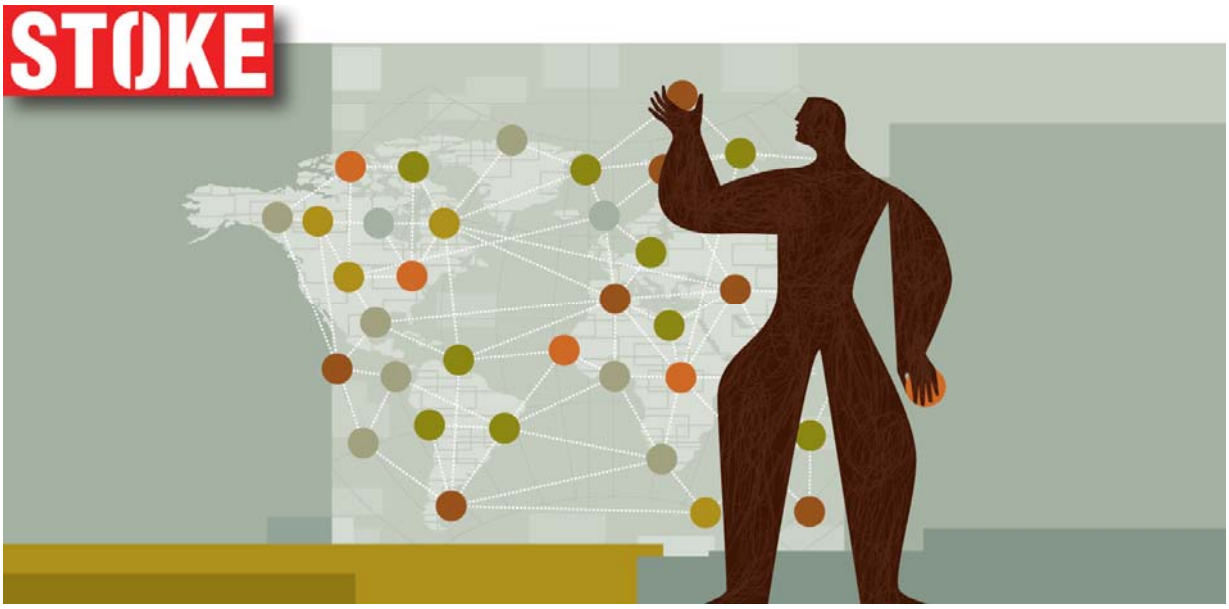

Stoke Mobile Data Offload

November, 2009

5403 Betsy Ross Drive, Santa Clara, CA 95054 USA 1.408.855.2900

Introduction

The biggest challenge for Mobile Network Operators (MNOs) today is how to handle the glut of mobile data traffic efficiently, both now and for as long as the 3G infrastructure will be in service. 'Efficient handling' means balancing customer satisfaction with extracting as much mileage as possible from the existing mobile infrastructure before triggering capacity upgrades, or even wholesale architectural changes. The overarching problem is that the vast majority of mobile data traffic is destined for the Internet rather than for operator services, and depletes capacity from the carrier infrastructure without delivering revenue. The expensive MNO data infrastructure was designed largely for low bandwidth, simple, walled garden services. Today that same infrastructure is under great pressure to manage this irretrievable shift in use towards mobile broadband Internet access with unsustainable service delivery costs. Operators have to do things differently and they have to act now.

Network operators are considering all options when it comes to relieving pressure on the mobile infrastructure. Traffic offload solution options depend on which part of the network needs relief. For example, Femtocell and WLAN Interworking solutions help to offload the RAN and backhaul systems. Stoke participates in several offload solution strategies that this paper will discuss, including a unique offering that, without any intrusive changes to the existing mobile data infrastructure, relieves a portion of the backhaul network and the core network elements from Internet bound traffic.

This paper examines the scope and scale of Internet traffic on mobile broadband services and provides an overview of different approaches that Stoke gateways achieve offload depending on their role in the network. The paper will also provide business return examples for the described options.

Managing the Explosion of Mobile Data Traffic: Where to Focus?

Mobile network operators are victims of their own success. The strategies employed to persuade subscribers to sign up for data service plans and onto 3G networks have delivered results beyond operators' wildest dreams - so much so that the resulting traffic loads are now wreaking havoc with service quality, while operators' corrective actions in some instances have created public relations challenges. Data volume increases are far outpacing revenue growth, and mobile network operators must improve service quality while reducing data services delivery costs in order to retain customers and maintain profitability. Combined with the current economic climate and its impact on rollout plans for LTE, mobile operators must significantly improve the long term profitability of their 3G networks for survival. Even as industry news outlets continue to herald the coming of LTE, there is no escaping the impact of the economy on the infrastructure supply chain's ability to deliver new mobile network solutions and solution elements, even if the MNO *could* afford to build out a new parallel network. Stoke dubbed 2009 the 'Year of Offload' because this was the year that MNOs had to re-examine their strategies and take immediate actions to increase the useful life and useful capacities of their 3G mobile data network infrastructures.

MNOs must find a way to service an increasing appetite for data-driven bandwidth while keeping capacity expansion (capital) expense in check. Approaches to this challenge can be directed towards three distinct areas of the mobile data network infrastructure (see figure 1 below) including: 1) the Radio Access Network; 2) the backhaul network; and 3) the data core network.

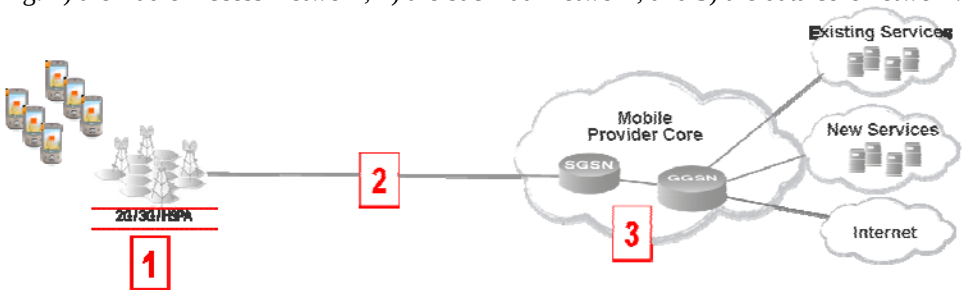

Figure 1: 3 Challenge Areas for Reducing Data Service Delivery Costs

The backhaul network is an area of intense scrutiny due to the rapidly growing and recurring expense. In North America in particular, backhaul networks are still largely based on T1 lines. Solutions for improving efficiency in the backhaul network range from IP enablement of the backhaul elements and TDM to IP/MPLS gateways to point-2-point Microwave solutions, all geared to reduce the cost-per-bit for carrying the growing traffic volume, but they do little to reduce the traffic load carried to the core.

Offloading the RAN requires employing alternative radio access such as WiMAX or Wi-Fi, but this is still a difficult sales proposition for operators as their wireless network is often their main marketplace differentiation. Adopting other networks – including ones they do not control – would certainly raise an eyebrow or two. The economic crunch and traffic glut is causing a shift in operator thinking towards alternative wireless network access.

Given the hierarchy of today's mobile networks, all traffic voice and data converge on the mobile core network and on the elements deployed there. This is most acute in the **data** core. The combination of the hierarchy and deployed network elements built for a by-gone era is driving considerable costs in the data core, and heretofore, neither equipment operators nor the standards bodies have offered true relief to this growing data plane condition. The architectural hierarchy and SGSN/GGSN systems are complex and network engineers are understandably reluctant to change. And given that large equipment providers' very livelihood is sustained by capacity expansion and upgrades of these systems, they are understandably slow to resolve the situation in the 3G network. This paper will focus on addressing the mobile data core network challenges caused by rapidly increasing mobile data traffic.

Current Options for Addressing Data Glut in the Mobile Core

MNOs do have options for addressing the spiraling data plane capacity challenge. This section will briefly describe two.

Business-as-Usual Capacity Expansion and Element Upgrades: Operators can choose to do nothing different, continuing instead to deploy more equipment to meet capacity requirements. As and when possible, MNOs can also upgrade SGSNs and GGSNs that deliver higher throughput to hopefully gain some ground on the escalating data volumes. While an improvement, this can be expensive and disruptive, particularly when deploying higher capacity network elements means introducing a new equipment supplier.

However, the real problem in this scenario is the currently deployed 3G network architecture, and the 'capacity expansion and element upgrade' approach does not deal with this fundamental issue. It might be likened to using a small adhesive bandage to dress a compound fracture.

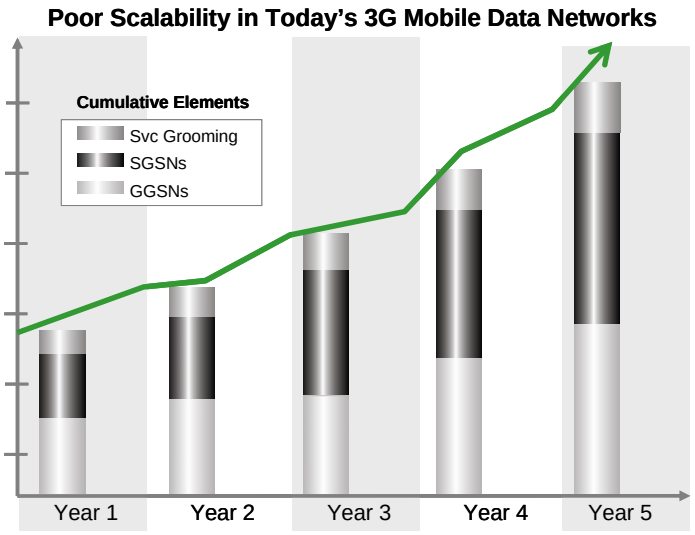

Figure 2: Core Network Element Growth in Lock Step with Traffic Growth

3GPP Direct Tunnel: Unlike the previous approach, Direct Tunnel does focus on the 3G data network architecture and it is being promoted as the answer to today's 3G data path challenges. However as a mobile network core data plan solution, Direct Tunnel falls short.

Pictured below, Direct Tunnel *does* remove the SGSN from the data path, establishing GTP data path tunnels directly between the GGSN and the RNC while keeping the SGSN in the control path. However the SGSN is only one of three elements in the data path.

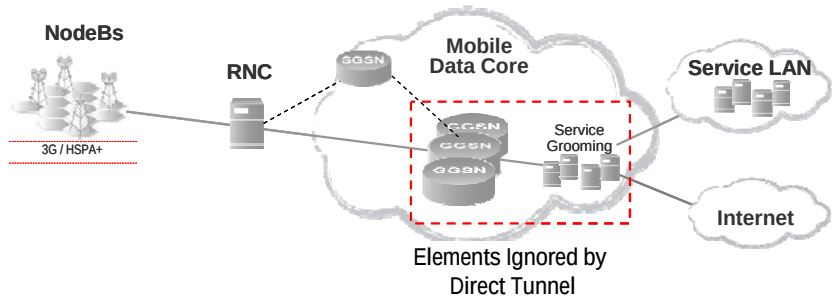

Figure 3: Direct Tunnel Only a Partial Solution

Provisioning of GGSNs and service grooming servers finds no relief with Direct Tunnel. It is very likely that MNOs choosing the Direct Tunnel approach will also have to upgrade at least the GGSNs to higher capacity systems in order to keep ahead of demand. With just these two choices for 3G network data plane solutions, more has to be done.

Taking Stock of the Problem

Web browser developer Opera reports having 20 Million users of its Opera Mini browser alone that generated 122 million megabytes of data traffic in the month of January, 2009. In his keynote address at Mobile World Congress in February, 2009, Siemens Network CEO Simon Beresford-Wylie predicted mobile data traffic on networks will outstrip revenues by a factor of 100. Even if the prediction is off by half, operators must take steps to change the status quo for growing network capacity to survive this data traffic tsunami.

Looking more closely at the mobile data traffic profile, more clues emerge as to how best to address the problem. For example, recently reported statistics indicate that the bulk of mobile data traffic is generated by a disproportionately few users and devices. Industry analyst Chetan Sharma reported that operator usage data shows that about 4% of subscribers consume nearly 70% of the bandwidth at a given tower. Those users are web surfing with 3G data cards on laptops. Ericsson estimates that 73% of traffic on wireless networks is generated by data card users, who only represent 3% of subscribers. The simple truth is that data card and 3G enabled laptop users are shaking the very foundation of mobile network operators' businesses. Any solution aiming to reduce the strain on the mobile network data plane needs to target Internet traffic, and it must include special handling for these heavy users.

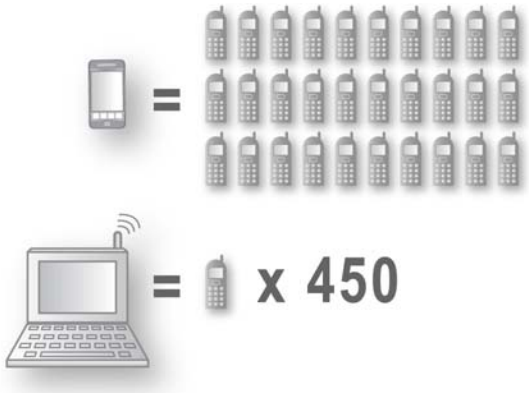

Figure 4: Data Traffic Generation Comparison
Source - Cisco Visual Network Index

Focused Attention Yields Simple, Elegant and Complete Solution

Stoke has created and patented a method for offloading mobile data traffic on 3GPP core networks by **as much as 90%.** Not only does the solution not require wholesale changes to the existing network elements, it requires no changes at all to the mobile network elements. In addition, Stoke's solution reduces the capacity requirements for the entire mobile data core network including the SGSNs, GGSNs, and service grooming resources.

Stoke has patented an approach that intercepts Internet bound traffic before it reaches the SGSN and routes it to the nearest Internet peering point. Stoke's MDO solution resides between the RNC and SGSN in the mobile data network, and has control plane intelligence to understand the type of session being established. Once it has identified a session as bound to the Internet it makes itself the serving point for that subscriber session. Critical control plane procedures like location updates to HLR or authentication against AAA are still taken care of by SGSN/GGSN and are unaffected by the presence of the data offload solution. MDO is able to intercept Internet bound sessions, bypassing the core data elements by routing them to the nearest interconnect point.

Stoke's MDO solution is fully compatible with 3GPP standards and the Iu-PS interface. The RNC and SGSNs are not impacted at all and do not require any reconfiguration when the MDO solution is deployed. MDO can interface with a large number of RNCs and SGSNs and is able to support any permutation and combination of RNC-SGSN interfaces.

Stoke Mobile Data Offload Overview

Stoke Mobile Data Offload employs an intelligent, high performance gateway to offload traffic destined for the internet – *traffic that offers no incremental value-creation opportunity for the operator* - from the 3G network as early as possible. The MDO gateway is inserted in the data path between the RNC and the SGSN (see figure 5 below). From this position, the MDO gateway transparently monitors RANAP control traffic, detecting and acting only on Internet bound sessions using APN, or specific device types like laptops with 3G modems using IMEI. In the case where the session is bound for operator services, the MDO gateway takes no action and user data flows as though the MDO is not present. For sessions bound for the Internet, MDO alters the traffic path, bypassing the SGSN and GGSN and forwarding it directly to the nearest Internet peering point. Critical control plane procedures like location updates to HLR or authentication against AAA are still taken care of by SGSN/GGSN and are unaffected by MDO's presence. Using fine-grained traffic selection policy implementations, MDO can select and direct specified Internet bound traffic to the operator's grooming services as needed.

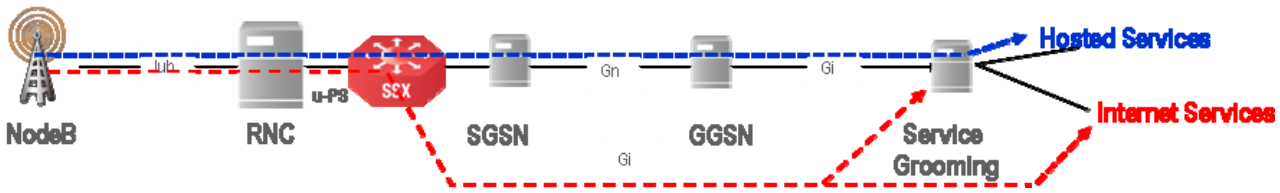

Figure 5: Stoke MDO Gateway Redirects Target Traffic Around Unnecessary Data Core Elements

Stoke Mobile Data Offload enables operators to:

- Offload their mobile data core network and service grooming resources from the bulk of the mobile broadband data traffic (i.e. Internet traffic)
- Optimize core network resources to provide value-added operator services.
- Offer a better QoS for subscribers due to a reduction in packet processing hops

Offload Traffic Selection

The MDO solution uses multiple means to identify Internet bound traffic, for example Access Point Name (APN), Client descriptor or, potentially, a user's IMEI (International Mobile Equipment Identity) to decide what traffic will be "broken out" at the MDO gateway and what will be sent to the core network. Using these selection mechanisms allows an operator to offload the 3G network without impacting the 3G network nodes themselves.

Moreover, the operator will be able to add dynamic traffic discrimination when PCRF architectures are brought on line. Traffic selection could be enhanced, for example, to use more information like tariff plan, mobile termination rate (MBR), time-of-day, application type, URL, IP 5-tuple, and other attributes.

Impact on Existing Network Nodes

Stoke is conscious that there is a clear boundary between the RAN and core networks and that they are discrete, complete entities. Each has its own management, planning, and engineering organizations, and often times each has its own incumbent equipment supplier. MDO sits squarely on the demarcation line between these two networks and can be deployed without hardware, software or configuration changes to the existing mobile network elements RNCs, SGSNs, or GGSNs. This results in a dramatically different deployment cost profile between the Stoke solution and the options provided by the incumbent equipment suppliers. Moreover, MDO is network agnostic and can be deployed in any RAN / core network combination.

The Stoke offload solution does require interaction with the transport network. While MDO decides what to breakout and where it will be sent, it is a transport node's responsibility to force Iu-PS traffic to the MDO gateway instead of the SGSN. This is accomplished in the transport router using policy based routing that sends all the SCTP traffic to the SSX.

Stoke Mobile Data Offload Financial Benefits

The location of the MDO between the RNC and SGSN means that the core network elements, SGSN and GGSN, are completely bypassed for some - perhaps most - data traffic, dramatically reducing the capacity requirements for these core network elements. To illustrate the magnitude of the potential savings, Stoke's Offload ROI model was tuned to model the effects of projected increase in data usage on the capacity demands for the GGSNs and SGSNs. The model then compares the capital and operational expense for the present mode of operations (adding SGSN/GGSN capacity as needed) against a future mode of operations (FMO) of deploying MDO. In addition, since operators have the option to upgrade their networks to higher capacity core network elements, and indeed to a Direct Tunnel architecture, Stoke's ROI tool was tuned to model these scenarios as well.

The key assumptions regarding the mobile operator model are listed in Table 1 below. They include the details of the current subscriber base and subscriber base and data plan growth factors. It also includes charateristics of the data usage per user today and forecasts for the end of the period. Actual network characteristics may vary and based on conversations with operators around the world these few tend to be on the conservative side.

Key business case example assumptions:

Beginning subscriber base	**15,000,000**
Percent of subs have data plan today	**40%**
Data plan penetration in year 5	**85%**
Data traffic destined to Internet APN and the Internet is	**85%**
Average data use volume per sub per month to start	**500Mb**
Average data use growth projected over the 5 year period	**450%**

Table 1: ROI Model Example Operator / Subscriber Characteristics

ROI Model Results

This is a comparison of the cost to provision the mobile data core network elements to handle the forecasted traffic loads using current generation SGSNs and GGSNs. MDO offers tremendous savings over today's present mode of operations. Chart 6 below illustrates model results by comparing the number of network elements that will need to be added to the network to manage the forecasted data traffic. The chart is not to scale, but the element counts tell the tale. The network that includes **Stoke MDO requires only 76** additional network elements by year 5 while continuing to deploy current generation GGSNs and SGSNs requires **310!** Not only does this signify a significant CAPEX reduction, but the reduction in network complexity is truly astounding. This simple comparison alone shows just how unsuited the current mobile network is for today's data traffic growth.

The Table 2 below present the financial view of the two methodologies. We see that over the 5 year period, **MDO reduces mobile data core CAPEX by 89% and OPEX 88%. The net addition cumulative cash flow resulting from MDO deployment is $201 million**. MDO clearly enables MNOs to scale the data capacity of the mobile network in a much more controlled fashion and ensure profitable growth.

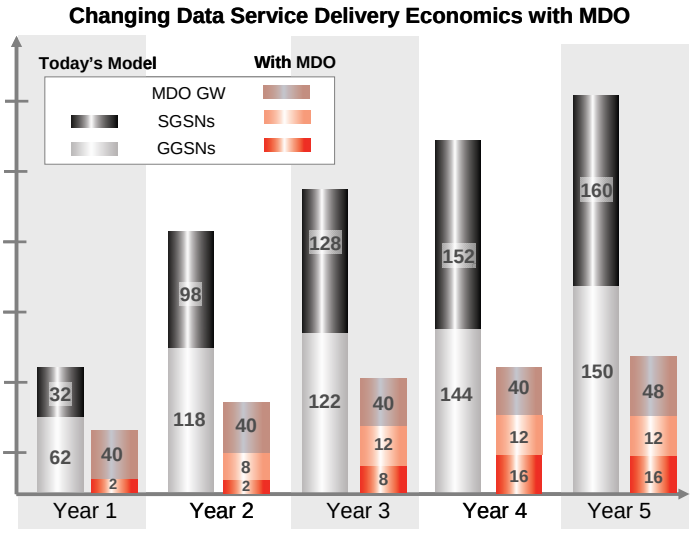

Figure: 6 Dramatically Improved Scalability with MDO

Present Mode of Operations vs. Stoke Mobile Data Offload

CAPEX Comparison Between PMO and MDO To Accommodate Forecasted Traffic Growth

		Year 1	Year 2	Year 3	Year 4	Year 5	Cumulative
Annual CAPEX - Present Mode of Operation		$ 31,122,650	$ 39,556,525	$ 19,766,950	$ 6,871,000	$ 4,312,925	$ 101,630,050
Annual CAPEX - With MDO		$ 7,756,000	$ 2,387,000	$ -	$ -	$ 1,234,800	$ 11,377,800
Annual CAPEX - Savings Due To MDO		$ 23,366,650	$ 37,169,525	$ 19,766,950	$ 6,871,000	$ 3,078,125	$ 90,252,250
Annual CAPEX Reduction (%) Due To MDO		75%	94%	100%	100%	71%	89%

OPEX Comparison Between PMO and MDO To Accommodate Forecasted Traffic Growth

		Year 1	Year 2	Year 3	Year 4	Year 5	Cumulative
Annual OPEX - Present Mode of Operation		$ 6,554,816	$ 20,020,899	$ 30,253,362	$ 34,659,190	$ 34,880,229	$ 126,368,496
Annual OPEX - With MDO		$ 1,630,865	$ 3,276,781	$ 3,405,868	$ 3,416,906	$ 3,249,616	$ 14,980,036
Annual OPEX - Savings Due To MDO		$ 4,923,951	$ 16,744,118	$ 26,847,494	$ 31,242,284	$ 31,630,613	$ 111,388,460
Annual OPEX Reduction (%) Due To MDO		75%	84%	89%	90%	91%	88%

Table 2: Stoke ROI Model Summary Table

Conclusion

With Stoke MDO operators can mitigate tens of millions of dollars of expense without disruption to their networks, solving data plane challenges and drastically reducing capacity expansion costs.

A tsunami of mobile broadband data is targeting today's 3G mobile data infrastructure. For many MNOs its affect is already being felt. The efficiencies of LTE that were just on the horizon have now slipped out of this decade due to the bleak economic climate affecting not only the operators ability to deploy, but for the entire ecosystem. With the industry unprepared for this situation, Operators are left to scramble to meet the data demands of their subscribers' on today's 3G networks for years to come. The approaching calamitous mobile broadband experiences will benefit the MNO that embraces every opportunity to reduce and manage the rising tide of data traffic. Operators that get it right will capture the hearts, minds, and contracts of the subscribers sent scurrying to higher ground. Operators that hesitate will struggle to deliver satisfactory services and initiate a downward spiral of declining reputation and customer churn with real financial implications.

Stoke has presented two of the options available to operators for managing the rising data tide, and has shown how they fall short of meeting the business needs of mobile data operators. Stoke Mobile Data Offload is the right solution at the right time. Fast and easy to deploy, MDO is non-intrusive requiring no modification or reconfiguration of existing mobile network nodes. MDO is also transparent in the network allowing all control plane traffic to pass untouched, and intercepting only operator controlled data sessions for routing more directly to the Internet peering point. Subscriber Internet sessions avoid the latency adding SGSN and GGSN network elements improving network response time. 3GPP standard-based mechanisms are used to assume control of selected sessions and all mobility events remain unaffected.

Stoke MDO also reduces CAPEX and OPEX by nearly 90% for a moderately sized mobile operator. Not only is MDO easier on the organization and the network to deploy, it dramatically improves the economics of data service delivery.

Now is the time to stem the tide, now is the time to explore every available option for managing the mobile data traffic tsunami and improve your customers' experience without breaking the bank. Now is the time to call Stoke.