quocirca
Business and IT Analysis

User experience monitoring

Ensuring IT applications are performing optimally

March 2010

The experience users receive when accessing IT applications is critical to the effectiveness of many business processes. That experience is affected by the application operating environment, the user's location and the network being used for access. To understand the overall user experience requires the ability to monitor all of these aspects and provide granular enough information to fine tune the IT infrastructure and ensure the optimal experience for all users.

This paper aims to elucidate the issues that affect the user experience in today's computing environments; how to monitor that experience, pre-empt problems and decide what actions need taking when the user experience is unacceptable. The paper should be of interest to both business and technical readers who know that delivering a good user experience is a key competitive advantage and want to be sure their organisation is benefiting from doing so.

Bob Tarzey
Quocirca Ltd
Tel : +44 7900 275517
Email: bob.tarzey@quocirca.com

Clive Longbottom
Quocirca Ltd
Tel: +44 771 1719 505
Email: clive.longbottom@quocirca.com

VISUAL NETWORK SYSTEMS

An independent report by Quocirca Ltd.

www.quocirca.com

Commissioned by Visual Network Systems

quocirca

User experience monitoring

Ensuring IT applications are performing optimally

The experience users receive when accessing IT applications is critical to the effectiveness of many business processes. That experience is affected by the application operating environment, the user's location and the network being used for access. To understand the overall user experience requires the ability to monitor all of these aspects and provide granular enough information to fine tune the IT infrastructure and ensure the optimal experience for all users.

- **Users cannot be relied upon to report poor experience**
  If they are external users, for instance using an e-commerce application, they may just abandon it and go elsewhere. If they are internal users, they may have become inured to poor performance, simply accepting it and operating unproductively with consequent impact on the business processes they participate in.
- **The issues that affect the user experience fall into three main areas; the application operating environment, the network and the location of the users themselves**
  The effective operating speed of a network and the resources available to a given application is rarely consistent and can vary through time. Understanding and controlling these issues is essential to improving the user experience wherever they happen to be located.
- **The widespread adoption of virtualisation has had particular impact on the way applications perform**
  In a non-virtualised environment the relationship between an application and the physical resources available on its host server are well understood. For an application running on a virtual server this relationship breaks down. For example, many virtual applications may compete for a small number of physical network connections.
- **The data that provides the insight into the user experience comes from a wide range of sources**
  Many already exist; network switches and routers, load balancers, firewalls, content filters etc. These can be supplemented by additional task-specific monitors. Sampling data from all of these and consolidating it centrally using a single data model allows a picture of the overall user experience to be visualised via a console. This enables business managers to direct how priorities should be changed and technical staff to make targeted changes to the way applications are managed and networks provisioned.
- **Such insight often allows problems to be solved without resorting to expensive infrastructure investments**
  If the issues that impact user experience are not well understood, the temptation is to throw money at the problem and invest in higher-performing devices and additional infrastructure. Such untargeted investment can lead to little improvement and is often unnecessary when other actions, such as modifying user behaviour, reprioritising network traffic or deploying a content delivery service, would have been cheaper and more effective.
- **To justify investment in user experience monitoring tools requires a value proposition to be built**
  Many of the issues that arise from poor user experience are intangible; lost customers, unproductive employees and ineffective business processes. Having an over-arching view of the user experience leads to an understanding of how to improve it, reduce business risk and create business value at a minimum cost. It is possible to build a compelling case, even based around a single known issue. If this can be done, the on-going ability to measure the user experience and answer questions no-one had thought to ask before is invaluable.

Conclusions

Information technology can be a wonderful thing when it works, but a miserable experience when it fails. The technology is there to enable users, not frustrate them, and ensuring the experience is more often good than bad is the only way to create a productive harmony between humans and computers.

Contents

1. INTRODUCTION – WE ARE ALL IT USERS NOW 4
2. THE IMPERATIVE OF A GOOD USER EXPERIENCE 4
3. ISSUES THAT AFFECT THE USER EXPERIENCE 6
4. PROVIDING AN AGGREGATED VIEW OF THE USER EXPERIENCE 11
5. POWER TO THE BUSINESS 12
6. TAKING ACTION 13
7. BUILDING THE VALUE PROPOSITION FOR UEM 14
8. CONCLUSIONS AND RECOMMENDATIONS 16
REFERENCES 17
ABOUT VISUAL NETWORK SYSTEMS 18
ABOUT QUOCIRCA 19

1. Introduction – we are all IT users now

In the old days it was all so simple. Mainframe computers ran applications, taking input from and serving data back to dumb terminals on the user's desk. Then things started to get complicated; client-server computing, driven by the introduction of mini-computers and PCs for users to access them, led to an increasingly hard-to-manage sprawl of IT infrastructure. Now things are coming full circle; virtualised data centres accessed via virtual desktops, web browsers or mobile devices mean applications are nice and centralised again, all back under control—phew!

This gross over-simplification of the last 50 years of the computer industry belies a truth; 1960s mainframe computing and 21st century IT infrastructure may bear some resemblance when shown diagrammatically but, of course, the underlying infrastructure and the demands placed on it could not be more different.

Behind the simplistic virtual clouds drawn on many computing diagrams lies a complex mix of contemporary and legacy IT equipment, often physically separated, yet linked by both private and public networks. However, the biggest change that has taken place over the last 50 years is in the user base. In the 1960s there were few IT users, a couple of specialists with access to green screens, everything else being run as batch jobs—if you wanted to know what the user experience was like, you just wandered along the corridor and asked them.

By contrast, today, virtually every employee is an IT user to some extent and the provision of IT resources extends beyond employees in the office to a wide range of mobile and external users; all have high expectations around how quickly they expect their applications to perform (Figure 1).

Figure 1 – response times considered acceptable for carrying out a remote transaction

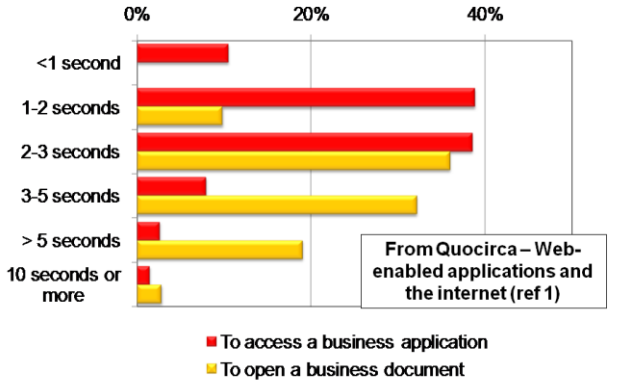

The IT department of a big company may be responsible for delivering IT services to thousands of employees and uncountable numbers of external users. Gauging the user experience and ensuring satisfaction has become a daunting task and one of the most pressing ones faced by IT managers. Deliver a poor service internally and face the wrath of the business, deliver a poor service externally and you may be responsible for failure of the business.

This paper aims to elucidate the issues that affect the user experience in today's computing environments; how to monitor that experience, pre-empt problems and decide what actions need taking when the user experience is unacceptable. The paper should be of interest to both business and technical readers who know delivering a good user experience is a key competitive advantage and want to be sure their organisation is benefiting from doing so.

2. The imperative of a good user experience

Main Findings:

- For the user, the experience is the only thing that matters, but users are fickle. They may not report a poor experience of using an application and, especially if external users, choose to go elsewhere instead.
- To be able to optimise the user experience requires monitoring of both the applications and the network equipment they rely on through time.

For the user the experience is everything. The average website visitor waiting for a response from an online booking form does not sit there pondering; *"is this delay caused by an overloaded server at the backend, a network contention issue or a problem with a load balancer?"* What most users observe is a delay and if that delay becomes unacceptable, external users, who have a choice, may go elsewhere or just go away with a sour impression of the organisation whose promised service has failed them. Internal users often have no choice; they may kick up a fuss or

find a way to work around the issues, bypassing the very process the organisation is trying to get them to use. Governance requirements and reporting capabilities may be impacted.

It may not be black and white; an application may work but its performance may be degraded for some reason. Employees who have to use such an application may continue to do so, but be working inefficiently. If the problem has arisen over time, they may not even have noticed, having become acclimatised to poorer performance. Whole departments may start to miss targets and managers may take it out on employees, when the real problem has been the slow decline in the performance of an application.

When eventually the problem gets so bad that someone notices, just like the web site visitor they will likely blame it on the application itself. This can lead to the wrong apportioning of blame—for example, a CRM or supply chain application being held responsible for a network failure—and put both internal and externals users off making best use of tools that should be making them more productive. Figure 2 shows the degree to which businesses were opening up such applications to external users, even back in 2007 when this research was conducted.

Figure 2 – internal and external users of applications

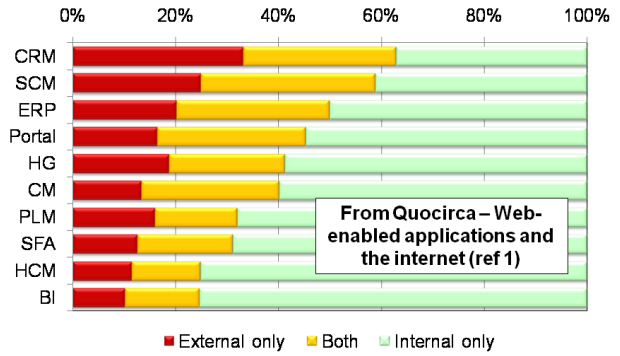

Abbreviations for application types	
HG	Home grown applications
CM	Content management
CRM	Customer relationship management
SCM	Supply chain management
SFA	Sales force automation
BI	Business intelligence
ERP	Enterprise resource planning
HCM	Human capital management
PLM	Product line management

"Whole departments may start to miss targets and managers may take it out on employees, when the real problem has been the slow decline in the performance of an application"

For some users just standing still is not good enough; expectations around performance have increased over time and a two second response time that was tolerated two years ago might be expected to be sub-second today.

Such shortcomings in performance need to be minimised. To achieve this requires the constant monitoring of the user experience through understanding the infrastructure that the experience relies on. This requires the continuous inspection of a wide range of components that link users to applications, from data centre load balancers though to branch office routers. Some problems can be pre-empted, for example recognising increasing traffic in certain areas is likely to require more bandwidth or better traffic management, others can only be fixed as they arise, for instance seeing that one network switch has suddenly become a bottleneck because another has failed.

It is also necessary to monitor the applications themselves. An application that was performing well one day may be degraded the next, even though the application itself, the infrastructure it relies on and the users accessing it have not changed in any way. Such a problem can be brought about by the introduction of a second application elsewhere. The users that rely on the first application have no knowledge of this interloper; the IT department may do, but may not realise what impact it would have elsewhere. The degree to which a new application takes up resources can be unpredictable, especially if it is one open to external users.

It is also possible that the performance degradation only happens at certain times. An internal application used by a European business may perform well all morning, but slow down in the afternoon as US workers start arriving at work accessing a service provided by the European's IT platform to their US customers. There may be no link between the two applications, other than that they contend for resources. Only by monitoring the first application and the infrastructure it relies on can the association between increased transatlantic network traffic and the afternoon slowdown be made.

So, it is not enough to take daily snapshots of the user experience, it must be measured at regular intervals during the day and historic data stored to provide baselines for comparison. To do this it is necessary to understand all the points in the delivery of IT applications at which the user experience may be impacted in some way.

3. Issues that affect the user experience

Main Findings:

- Application performance is affected by the network, application infrastructure (including virtualisation) and the user location.
- At its most basic level, a network has an inherent given capacity. Only by understanding the above factors is it possible to understand why maximum capacity is not being used.
- To understand the user experience across a broad range of users requires the ability to sample data from a range of network and application-level devices and channel it to a central location for analysis.

The issues that affect the performance of an application fall into three main areas: the network, the application and the infrastructure it relies on (servers, data centre, virtualisation etc.) and the users (their location and behaviour).

Network complexity

It is easy to think of a network just in terms of bandwidth, a fluid flowing through a pipe. The analogy with cars on a road is useful; too many cars (packets of data) down too small a road (wire or airwaves) and the traffic slows. However, as anyone who has driven a car knows, it is quite possible to spend hours moving slowly along a wide open road, because a junction at the end of it has become clogged, or because of the way that the cars are being driven. Like roads, data networks are about more than just pipes. However, ultimately, if less data is getting through the network than should, in principle, be possible, the right monitoring tools can help understand why.

Data networks are, in fact, more complex to model than vehicles on roads. A road has just one layer, whilst a computer network has many. Since the 1960s these have been described in terms of the 7 layers of the OSI (open system interconnection) stack. This is still the basic standard for describing the interfaces between network equipment, although the boundaries between many of the layers have become blurred as the equipment used to manage networks often spans multiple layers, flattening the stack and complicating the problem of understanding network performance.

A more useful standard to look at to understand how networks affect application performance is TCP/IP, the protocol used for actually transmitting data. Since the use of the internet became widespread in the 1990s, TCP/IP has replaced virtually all other network protocols. It predominates, both for local area networks (LAN) and wide area networks (WAN), where TCP/IP is either the native protocol or “spoofed” over what the underlying physical transport was initially designed for. The use of a single transmission protocol means that one possible performance-impacting issue has been removed—the need to constantly translate from one protocol to another as traffic is routed from one network to another.

As with a road, we have some idea of the maximum amount of traffic a network should be able to handle. The busier it is the more contention for resources and collisions there are likely to be. If you fill a road completely with cars, then it becomes a car park and nothing moves anywhere. The same is true of a computer network “pipe”; its maximum loading, whilst remaining efficient, is around 60% of total capacity. In practice, taking this into account, there will be a hypothetical maximum data transmission rate for a given cable or other transmission medium. However, often this rate is not achieved due to other problems that effect network performance, especially at the junctions within, and the interfaces between, networks that are not operating at “wire speed”.

The issues that impact network performance will, if they go unnoticed, have a negative impact on the user experience. Some of the worst offenders are jitter, where packets arrive at different times (a particular irritation for voice and video applications) and sawtoothing, where the network speeds up and slows down due to devices buffering data. The speed at which a signal can travel through a network is called its latency and is usually measured

by checking the time taken for a signal to make a round trip and return to the starting point. Measuring latency gives an idea of the overall performance of a network and the impact such intervening devices are having on performance.

For some applications the most critical network interface is that between the inside and outside worlds—the LAN and WAN. Most LANs today are based on Ethernet; a cabling system for linking local computers based on agreed de facto and de jure standards. LANs in larger organisations are usually broken down in to network segments connected by switches that manage the junctions within a single network. Routers allow one LAN to interface with another. If switches are akin to roundabouts allowing cars to move from one road to another with the minimum amount of fuss, routers are more like car ferry terminals, where the mode of transport may change all together, as a car leaves the road and heads out to sea on a ferry. On disembarking at the destination it may be that the rules of the road may have changed altogether.

The WAN infrastructure that connects LANs can be privately owned, but more often than not it is based, at least in part, on shared infrastructure provided by network service providers and the public internet. Therefore, most wide area networking, even for internal use, is carried out using virtual private networks (VPNs). Traffic flowing across VPNs is sharing physical infrastructure with other users, over which there is no direct control.

VPNs should be better than relying on the internet. On the open public internet, all packets are scatter-gunned across the network and will eventually turn up and be put back together, but the latency introduced can be severe. With a VPN, all packets go down a single dedicated tunnel, often with an associated SLA provided and managed by a VPN provider—so performance guarantees should be much easier to provide than over the internet, although VPN end points themselves can degrade performance to some extent.

Table 1 – Network devices that may impact the user experience

- Switches and routers: junctions within and between LANs
- Network load balancers: balance the load of traffic across other network devices such as routers
- Internet load balancers: balance the traffic load across available internet access points
- Network accelerators: speed traffic between routers through data compression and other techniques; they may in themselves be routers and include local caching and network traffic control capabilities
- IP Sec VPN routers: connect the end points of VPNs
- SSL/VPN routers: allow secure remote access into private networks by authorised individuals over the internet
- Firewalls: close down the network ensuring only certain ports are available for use and the monitoring of the network traffic that passes through them
- Intrusion prevention/detection devices: prevent unauthorised access to networks
- Content and packet inspection devices: may be used as part of an intrusion detection system, but also for data leak prevention, anti-malware and controlling user behaviour on the internet
- Proxy systems: often used to provide a consolidation point for data control, for example remote user internet access
- Encryption systems: often part of a VPN device, but also used for general data stream encryption
- Application firewalls: prevent an unauthorised access at the application level
- Mail transfer agents (MTA) that receive and direct email
- VoIP routers: that route IP voice calls, often linking them with PSTNs (public switched telephone network) network using SIP (session initiation protocol)
- Integrated routers combine many of the above into single devices, often used at the branch level

With VPNs, it is sometimes possible to manage quality of service (QoS) and prioritise one type of network traffic over another, for example voice traffic over email. Only with MPLS (multi-protocol label switching) can end-to-end QoS be guaranteed over a WAN. MPLS is a low-level, network-agnostic transmission protocol that allows traffic to be routed along pre-assigned network paths.

Of course, to make a network operate requires all sorts of other devices that can speed the network up or slow it down. These devices may all be necessary, but their good operation needs to be ensured. These devices are also

sources of information when it comes to understanding the user experience. Some of the more common ones that affect network traffic are listed in Table 1 and shown diagrammatically on Figure 3.

Application performance

To get a full understanding of application performance it is necessary to understand the impact of the operating environment of the application as well as the network that has been used to access it. Most applications do not operate independently. It is not just that they contend for resources with other applications, but often they interoperate with them. Business processes are driven by the handing off of control of a task from one application to another. In some cases this may be automated, in others user-driven, but either way, if any of the links in the chain are performing badly, the user will be impacted. For applications that are owned and run by a given organisation, monitoring application traffic at source, in the data centre, provides essential information for understanding the user experience.

However, some applications that form a critical part of a business process may not be owned by the organisation that owns the process; for example external supply chain or CRM systems may be invoked. Here, the provider can offer a service level agreement (SLA) for their own application but not for the network it is accessed via. In these cases access is usually via the public internet and resources cannot be guaranteed, although measures can be taken to maximise performance if needed.

The biggest change to the way applications are run in the last decade has been the rapid adoption of virtualisation, and to understand application performance today requires an understanding of the virtual environments many of these applications run in and how physical resources are apportioned in such circumstances.

The impact of virtualisation on user experience

Before the re-invention of virtualisation for x86-based architectures, it was relatively easy to understand the relationship between the instance of an application and the network resources available to it. The application sat on a given physical server that had a number of Ethernet ports connecting it to other equipment, switches, routers, other computers and so on. Virtualisation for modern computing environments became attractive because using servers in this way had become too inefficient and expensive, particularly in energy and space costs; utilisation was often below 10%.

"Virtualisation adds a whole new set of problems to monitor and overcome to ensure the end user experience of an application is acceptable"

Hypervisors are the basis for server virtualisation. They provide an abstracted layer between the hardware and the overlying operating systems, each of which then runs in its "space", using shared pools of resource such as CPU, memory and storage. This gets rid of the tie between the application and hardware as all hardware running in a virtualised regime becomes part of a resource pool, across which the application load can be balanced. In this way, a busy application can be allocated more resources when needed and excess capacity can be temporarily provided on-demand, sometimes from external providers. Great!

Far higher utilisation rates, dynamic resource allocation – an end to the problem of applications performing poorly due to lack of server resource then? Not quite. There are a number of issues, not least that many existing applications are not yet well adapted to take advantage of virtual environments; others that are may compete better for server resources. Even virtual images can hit resource ceilings; when this happens, it may be necessary to start a new instance of the application alongside existing images, and then balance the workloads between the images.

It is also easy to abuse the virtual environment, creating multiple virtual machines (VMs) that all need maintaining with patches, upgrades and security updates and so on. There are also complex software licensing issues, but this paper is not the place to discuss the pros and cons of virtualisation, except in so far as it has an impact on the user experience, and, if not well managed, this can be profound, especially where the underlying physical servers interface to the network.

The problem is that all those virtual machines will want to open network connections and whereas before, when a single application running on an underutilised physical server had all the connectivity it needed, hundreds of VMs on a fully utilised server may be contending for what are now very limited network connections. Virtualisation adds a whole new set of problems to monitor and overcome to ensure the end user experience of an application is acceptable.

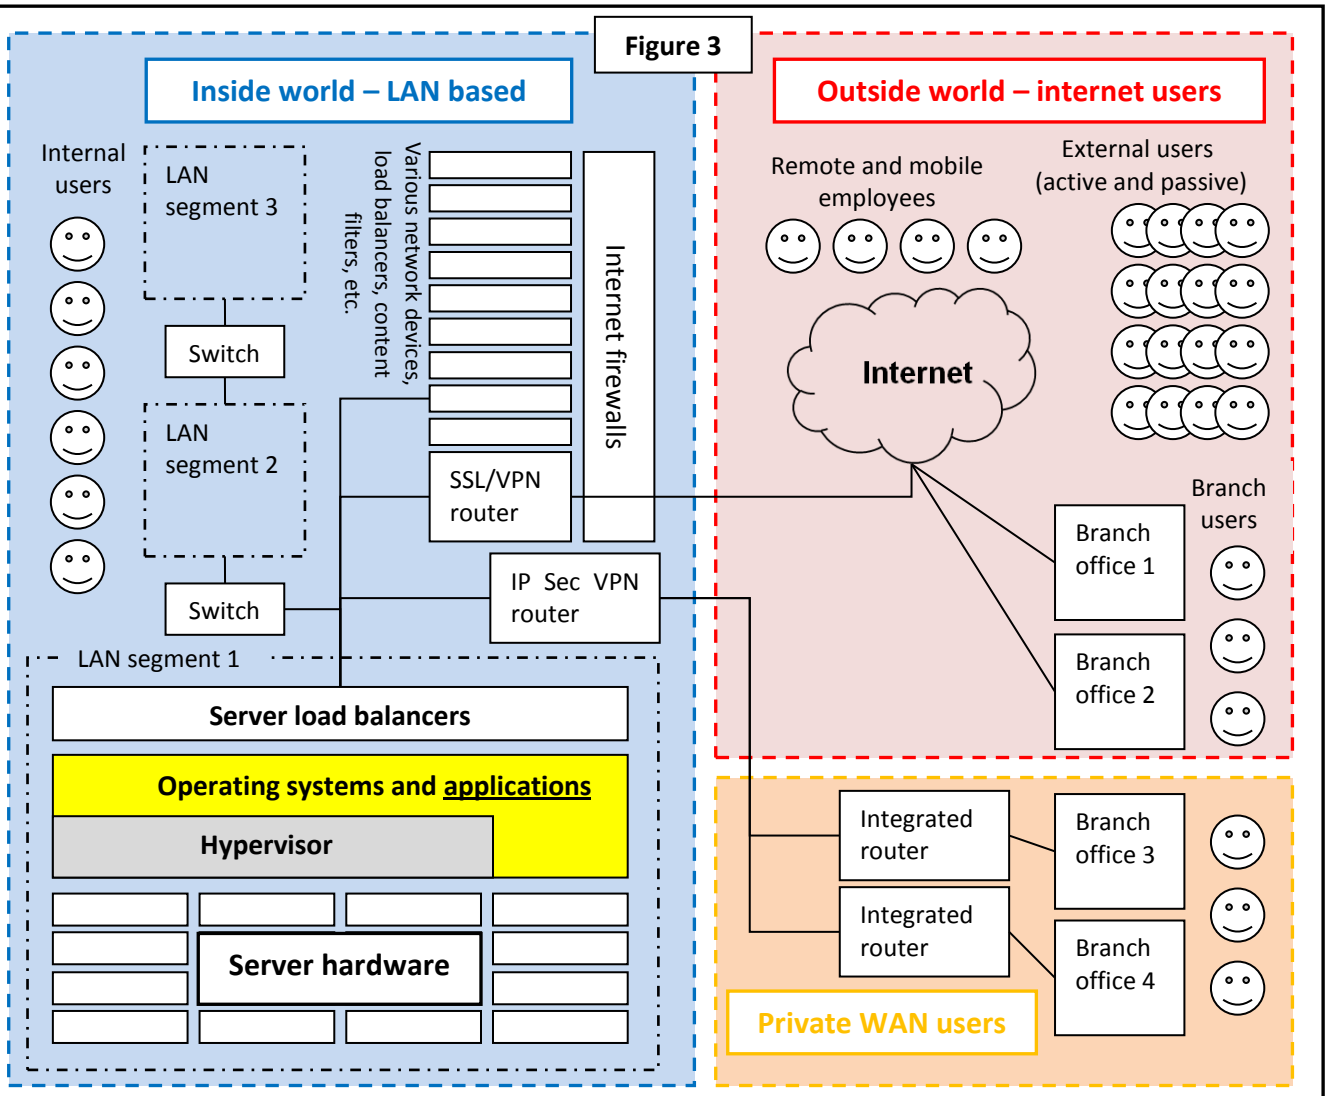

Figure 3 – The distribution of users and the barriers between them and the applications they rely on. How do you ensure they keep smiling?

User location

If an application seems to be operating well and the network seems to be running optimally, the user experience should be OK, shouldn't it? Ask a user plugging into the local LAN on the same premises as the infrastructure running the application and they may well agree. However, many users of the same application will be accessing it from a remote location and understanding and ensuring their experience is more problematic. For some applications the users that matter will be external ones.

Remote users fall into 5 categories (see Figure 3):

- Branch employees: includes point of sale workers, bank tellers as well as conventional office workers in remote branches—they may access centralised applications via a VPN or through the internet.
- Remote employees: home workers and employees working off-site may be at a customer premises accessing via the internet.
- Mobile employees: employees in the field accessing applications via mobile devices over 3G or GPRS networks.
- Active external users: traders linking to supply chain applications, internet shoppers, online bank customers etc.—usually over the internet.
- Passive external users: users of remote devices such as passenger ticketing systems, viewers of in-store displays—some over the internet or via a privately owned network of some sort.

Figure 4 shows the degree to which European businesses in various industries had opened up their networks to various external users when this research was conducted in 2008. Quocirca believes that it is likely to have increased markedly in the intervening two years.

For a given application, it is necessary to understand the experience of users in all relevant categories to understand the overall user experience. For example, for a transport ticket checking application, it is no good if operators are able to run reports on passenger numbers if the same passengers are being turned away at the entry gates because swipe card devices are running too slowly due to excessive reporting requirements.

Figure 4 – percentage saying external users are provided access to internal systems

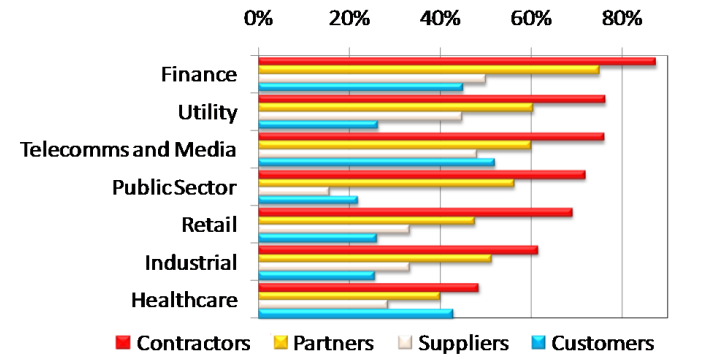

From Quocirca – The distributed business index (ref 2)

"If one branch is under performing is it because the staff have poor customer skills or because the order taking systems in consistently slow?"

There are many points in a network that may be the weak link: Figure 3 shows where some of these may lie. In the case shown, a problem with the IP Sec VPN device would only affect users in certain branches. If one branch is under-performing is it because staff members have poor customer skills or because the order taking system is consistently slow? It may become clear that all internet users are receiving a sub-optimal experience, but which of the plethora of network access and security device is causing this?

To provide that over-arching view of the experiences of all users requires a means of collecting representative data together so that it can be analysed in one place.

4. Providing an aggregated view of the user experience

Main Findings:

- A common data model is necessary to provide a single view of data collected from a wide range of sources.
- Data needs to be aggregated at a collection point; this may be a physical or virtual network appliance in its own right.
- Tools are needed to analyse the data and provide a graphical view of the user experience that can be understood by both technical and business staff.
- This view of the user experience is critical for making ongoing decisions about improving that experience.

To get an over-arching view of the user experience, it is necessary to aggregate data from all the various sources discussed so far and, where necessary, plug gaps with data collected specifically for UEM. Quocirca research shows that European businesses understand the importance of doing so (Figure 5). However, with hundreds of different devices from various vendors, carrying out a wide range of functions, all continuously generating data—how do you even start?

Figure 5 – importance of metrics for measuring the performance of web-enabled applications

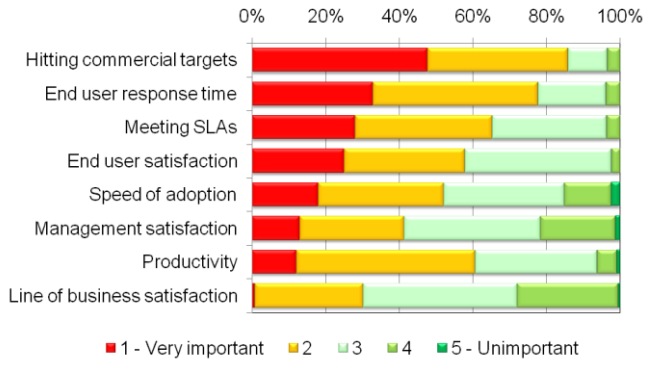

From Quocirca – Web-enabled applications and the internet (ref 1)

The first thing that is needed is a common data model. This is a schema that data can be translated into from different sources to allow like-for-like comparison and an aggregated view. Then a collection point is required; this is not for the collection of all data, but a place where point samples for given time intervals can be collected and stored. This collection point may well be a network appliance in its own right, or perhaps a virtual appliance.

To recognise that the experience of users in a branch office has deteriorated, it is necessary to have long term history of the traffic handled by the local branch router; to do this it may make sense to take regular daily snapshots of the router. On the other hand, to understand what impact a server load balancer is having on an application, when its performance is seen to deteriorate in the afternoon, may require taking minute-by-minute records of its activity. With more frequent sampling it becomes less practical to store long term historic data.

The actual data collected will vary from device to device and even once it has been translated into a common data model, it requires intelligent tools to formulate a meaningful overview. Fortunately, there is some level of standardisation that makes the job easier than it might otherwise be.

One of the most useful reporting formats at the network level is NetFlow, a protocol developed by Cisco for collecting IP traffic information. NetFlow has been adopted by other network equipment manufacturers, albeit under different names (e.g. Jflow or cflowd from Juniper Networks; NetStream from 3Com/H3C (now HP) and Huawei Technology; Cflowd for Alcatel-Lucent). NetFlow measures the output of routers and switches. Collecting such information can itself be a burden on a device so “sampling” is a normal practice where every nth packet is inspected.

Another useful source of network information is SNMP (Simple Network Management Protocol), which is overseen by the Internet Engineering Task Force (IETF) and is used for reporting on the status of network attached devices. It can report on issues such as "free memory", "system name", "number of running processes" and "default route”.
However, information provided by NetFlow and SNMP is at a low level in the network and provides little information about the performance of individual applications and the traffic they generate. For applications owned by a given organisation, this requires data centre based performance monitors that can provide such specific information.

Application performance monitors (APM) can measure the overall response times of specific applications and understand where time is being spent, specifically what proportion of overall time taken to complete a transaction is

network transfer time and how much is taken by application processing. To do this requires collecting data packets that are specific to a given user transaction for an individual application. So, this is much more than just counting bits and bytes but an understanding of the application conversation. For example, if QoS is being applied, it may be policy that some users get a slower response than others; therefore application performance can only be judged in the context of an individual user transaction.

APMs also enable the identification of anomalous use; for example inappropriate or unanticipated use of applications that generates unexpected network traffic (e.g. large file downloads, streaming video). Understanding when demand is exceeding expectations allows actions to be taken; this may just be reinforcement of policy or a change to the application itself. In the latter case it is necessary to provide the application owners a trace of the anomalous behaviour. This can be done by constructing a trace file from captured data that reproduces the behaviour of a given application for a given time period. The application owners can use this to understand where the problem lies, if necessary down to individual lines of code.

Understanding the experience of users who rely to a greater or lesser extent on the internet is more complex than monitoring the experience of purely internal users. There is, however, plenty of useful information available from branch office routers, internet load balancers and the devices on which the users themselves rely.

Other gaps can, when necessary, be filled by the use of purpose-placed network probes, although these should be a last resort when it is certain there is a gap in the understanding of the user experience that cannot be filled in any other way. Such probes may be on the wire itself, looking at network traffic or acting as application user proxies, ascertaining what the given user in a given location would experience at given times. There are external service providers who can provide such information for internet users at a wide range of global locations through the use of synthetic users.

5. Power to the business

Main Findings:

- Tools that visualise the user experience are required to provide insight for both business and technical staff.
- The view provided should be from the user perspective rather than a technology one.
- Flexible reporting is required to ensure newly arising problems can be quickly analysed.

At the end of the day, resources will always be limited. Even after a network upgrade it might not be long before peak capacity is reached, so ensuring a good user experience is always going to be a compromise of some sort. Some users and some applications will need to be prioritised over others. For example, in a retail outlet, the point-of-sale (PoS) system that drives a customer's transaction should probably be given a higher priority than the branch manager's incoming email.

Such decisions about who and what should have precedence should be taken by the business and not IT managers; the latter should advise what is possible. This means that tools to understand the user experience should have visual outputs that make sense to non-technical staff and enable them to make decisions about how user experience should be prioritised.

"Tools to understand the user experience should have visual outputs that make sense to non-technical staff and enable them to make decisions about how user experience should be prioritised"

To do this it helps to identify constituencies that are based on the distribution of users and not centred on devices. For example a large site may have multiple routers and a small site just one; but all the same they are both sites, so it makes sense when understanding the user experience to consider the multiple routers of the large sites together. Beyond sites, broader user groups can be defined and monitored, for example those from different departments, countries or time zones.

The tools that visualise user experience data should allow the easy sorting of data by these different classifications, such as allowing the rapid identification of a few sites—out of what may be thousands—that are receiving unacceptable response times for a given application or searching for a particular site that is generating an unexpected amount of inbound or outbound traffic. In some cases these will be regularly run reports but in other cases it will be necessary to construct ad hoc reports to provide insight in the newly occurring problems.

Such powerful reporting capabilities can provide positive insights as well as negative ones. For example, if a given network sector is running at 80% capacity but the overall response times experienced by users that rely on the network is acceptable, then IT managers can prove they are meeting given service levels whilst maximising use of IT infrastructure.

Some businesses will also be faced with the task of providing a user experience monitoring service to multiple "internal customers". Most obviously this includes service providers who provide virtual networking services to many organisations. However, it can also include businesses where, perhaps through a previous merger, a unified IT infrastructure is now serving multiple lines of businesses. Reducing IT and other infrastructure costs may have been the aim of such a merger in the first place, but the unified IT department will need to report on the service being provided to two or more different sets of business users. For these reasons it will often be a requirement that user experience monitoring tools can collect and report on multiple sets of data.

Once the business has a view of the experience its users have, it can work with IT to decide the most effective way to address problems that have been identified.

6. Taking action

Main Findings:

- Data collected by user experience monitoring tools allows the most cost effective actions to be taken to improve the user experience.
- In many cases, this will not be investment in new equipment but may be as simple as asking for, or forcing, a change in user behaviour or subscribing to a new network service of some sort.
- The prudent targeting of investment plus the improved customer experience and reduced risk to the business are the basis for justifying investments in user experience monitoring.

Knowing that there is a problem with user experience is one thing; doing something about it is another. With the fine detail provided by monitoring tools, the actions taken to improve user experience can be targeted and precise. Without such information the temptation is to throw capacity at the problem through investing in extra servers, more bandwidth, bigger switches and routers and so on. As well as being costly, this can be frustrating when the problem persists. More subtle solutions are often going to be cheaper and more effective.

For example, if it is known that the problem faced by internal users in a European organisation on weekday afternoons is down to increased traffic from external users getting to work in the USA, a better and cheaper resolution to the problem may be to invest in a content delivery service (where more content is cached locally by an external service provider), rather than faster servers.

Other problems can be resolved simply by voluntary or enforced modification of user behaviour. If internet access times to externally-hosted applications have become degraded and monitoring shows this is due to other low priority activity by some users on the internet, for instance viewing bandwidth-intensive video content, then software can be put in place to limit or block such activity. These capabilities may even exist within existing routers or can be added to them or purchased as an external service.

Some problems can be solved at no cost, simply by changing user behaviour and/or expectations, for example telling employees that running complex reports on customer databases at peak trading time is likely to be slow—do it early in the morning or later in the afternoon, or submit it as a batch job where the system can make the decision to run a

report based on real-time resource availability. In this way, both the employees and customers will benefit. Or it may be the case that changing priorities solves the problem, making sure the internal report always runs at a lower priority than a customer transaction.

It is not always easy to control priority, especially when relying on the public internet, but even here things can be done. Content delivery providers can select routes for traffic that will be optimised and faster than the random routes taken by most internet traffic. It may be justified in some cases to invest in virtual private network (VPN) infrastructure for users in some locations to provide guaranteed bandwidth allocation across public networks.

However, there will be some situations where the only choice is to invest in new infrastructure, but at least with the UEM tools in place a level of confidence can be provided to those that approve spending that the investment is justified. Such investments may be to provide better WAN performance by migrating to MPLS or introducing network acceleration devices, to undertake a server virtualisation project or to engage with a third party, for example to have access to additional computer power from a "cloud based" platform-as-a-service provider when necessary.

It is not that improving user experience is always going to be cheap, but if the problem is well understood the best way to approach it can be taken, for the lowest possible cost. Take these two things together—the benefits of improved user experience, plus the cost saved in doing so and add in the downside risk of not taking any action—and you can build the case for investment in user experience monitoring tools that will be required before getting such an investment agreed to by those that sign the cheques.

7. Building the value proposition for UEM

Monitoring so many data sources to build a complete picture of the experience of a wide range of users is a complex task and tools that have the capability to do so come at a price—an investment that needs to be justified. The problem is that quantifying the return on that investment can be hard when many of the benefits are intangible. For instance, how do you quantify the value of a performance improvement that was made to an area of a network, due to a problem being noticed before it had any chance to impact the business processes that relied on it?

Once such tools are in place, their repeated use can help ensure the occurrence of such problems is minimised and their availability becomes taken for granted—"*how did we manage before*"?—but getting agreement to the initial investment can be a hard hurdle to jump when the business benefit is indirect. All too often the investment happens because an emergency has arisen and solving a given problem has become so critical that barriers to investment are more easily broken down. But even if it is not an emergency situation, considering the impact that user experience monitoring tools may have on a single issue can be the way to build a compelling case.

Often such cases are built purely around return on investment—if we spend X we will save Y, and Y is greater than X. This is a reasonable expectation, but for softer investments a broader case often needs to be built that can take into account immediate benefits, possible future ones, and looks beyond pure cost control to increasing business value and reducing business risk. Quocirca calls this a total value proposition (TVP) and it is built around those three elements; cost, value and risk. The TVP will vary for each individual customer and application but it is possible to list the generic factors that should be taken into account; in a given case some will have more impact than others.

Cost to the business

Cost is about tangible figures, those which can actually be credibly measured. One cost that is very tangible is the amount a provider of user experience monitoring tools will charge for their products, the services required for putting them in place and training business and IT managers to make best use of them. For a distributed enterprise with thousands of employees, and maybe countless external users, this is likely to be a six figure sum, be it €, $ or £.

Another set of costs for the business will be any investment in new infrastructure that the UEM process identifies as necessary (see section 6). These costs will range from zero upwards depending on the problem identified and the remedy suggested. However, whatever the amount, the confidence provided that the money spent will address the problem will be higher with the insight provided by UEM.

Finally, costs to the business need offsetting over a period of time that is commensurate with the duration of the benefit being delivered. In many cases that benefit will be on-going so the cost should be spread over multiple years.

Cost savings for the business

The more tangible the cost savings identified to offset these investments, the stronger the case that can be made to those signing off the procurement. The following list identifies some of the main directly measurable cost savings that might be expected, the importance of each varying on a case by case basis:

- Eliminating outages of applications leads to more orders: if a given problem has led to regular application down time that directly impacts the ability to take customer orders then it is possible to calculate the value of orders that would be taken if that down time was eliminated. This may be an application used by internal users to process orders or an e-commerce website accessed directly by customers.
- Speeding up transactions: if application delays are constraining employee output, then it is possible to calculate what revenue increase or cost reduction a percentage increase in that output would lead to if the delay is reduced by a given factor.
- Reducing the help disk burden: if the number of trouble tickets and support calls being logged by a help disk can be reduced, it is possible to calculate the savings in help desk time. Some staff members are then freed up for other activities or total staffing numbers can be reduced over time.
- Shorter repair times: when a real problem is identified, having a rapid and detailed insight into where the problem lies means quicker fixes and lower use of engineering resources.
- Customers that never were: well-performing applications keep customers engaged; for example improving the number of visitors to an e-commerce site that actually buy or of those to a self-service support that actually solve their own problems.
- Ensuring cost savings of given IT investments are realised: for example if a voice over IP (VoIP) system has been put in place is it performing as well as expected now a shared network is relied on and are the promised reductions in telephony costs being delivered?
- Eliminating redundant equipment: improving performance in one area may mean out-dated equipment in another area is no longer needed. In some cases it may already be redundant but the fact has gone unnoticed. Energy, maintenance and service charges can be saved.

Creation of business value

Other benefits are softer and cannot be so easily quantified; however, the business value they create should not be overlooked:

- More motivated employees leading to lower churn: having user experience monitoring in place will mean that, in the long term, employees are better able to perform, getting more job satisfaction and perhaps higher earnings.
- Work/life balance: guidelines can be issued for personal use of the internet and those employees that abuse them taken to task. Many employees have come to expect such liberties, which often amount to the online equivalent of going to the shops at lunch time.
- Dynamic customer interaction: the ability to run front line activities such as customer sales at the same time as back office reporting means customer behaviour can be monitored in real time and the products and services available to them modified accordingly.

Reduction of business risk

Factors around reduction of business risk are also intangible but should be taken into account:

- Brand reputation: poor service may lead to short term loss of orders but it can also lead to long term reputational damage which may keep customers away even after a problem has been fixed.
- Less money spent on upgraded equipment: without the insight provided by UEM, IT departments are prone to throw resources at a given problem often with too little effect. UEM helps eliminate such wastage. This is particularly true of network bandwidth; there is a temptation to over-invest, even when it may not be the source of a particular problem.
- Ensuring low priority activities are not hogging resources that could be allocated to higher priority ones.

Not all of the above apply in all situations, but selecting those that do and identifying others that are specific to a given problem allows a TVP to be put forward that will be much more likely to offset the combined costs of the tools themselves and the untargeted investment in infrastructure that would be necessary without the level of insight provided by the tools.

8. Conclusions and Recommendations

Anyone who has read this paper may well recognise that their organisation faces an existing problem with user experience or that it may do so in the near future. However, it may just be that you do not know—you might be asking yourself questions like:

"What is it really like for my customers serving themselves online?"

"Could the problems in the call centre be an application performance issue rather than an employee motivation one?"

"Why did we spend all that money last year on new network switching equipment for no apparent improvement?"

Questions like these can be answered if the tools to gather and analyse all the available data are in place. The motivation to purchase them may be down to a single initial problem, but once the benefit of the insight provided by such tools have been realised, they are available to be used over and over again; perhaps answering questions you had never thought to ask in the first place.

Information technology can be a wonderful thing when it works, but a miserable experience when it fails. The technology is there to enable users, not frustrate them, and ensuring the experience is more often good than bad is the only way to create a productive harmony between humans and computers.

References

1. Quocirca: Web enabled applications and the internet – October 2007
   http://www.quocirca.com/reports/144/web-enabled-applications-and-the-internet
2. Quocirca: The distributed business index – March 2008
   http://www.quocirca.com/reports/137/the-distributed-business-index-dbi

About Visual Network Systems

Visual Network Systems provides innovative solutions that support IT professionals responsible for enterprise-wide application, network and VoIP performance. Our flagship solution, Visual Performance Manager, is a unified system that provides enterprise service intelligence to help organisations effectively deliver these services. Our solutions are available directly to enterprises or as part of a managed services offering.

Visual Network Systems is part of the Danaher Corporation, a Fortune 500 company and leading manufacturer of professional instrumentation, industrial technologies and components with revenues of more than $12 billion (USD) annually. Visual Network Systems is headquartered in Colorado Springs, Colorado, USA, and distributes its products in more than 50 countries. More information can be found by visiting Visual Network Systems web site at www.visualnetworksystems.com or by calling (888) 293-5853 (US & Canada) or 425 446-6400 (other locations).

REPORT NOTE:
This report has been written independently by Quocirca Ltd to provide an overview of the issues facing organisations seeking to maximise the effectiveness of today's dynamic workforce.

The report draws on Quocirca's extensive knowledge of the technology and business arenas, and provides advice on the approach that organisations should take to create a more effective and efficient environment for future growth.

Quocirca would like to thank Visual Network Systems for its sponsorship of this report.

About Quocirca

Quocirca is a primary research and analysis company specialising in the business impact of information technology and communications (ITC). With world-wide, native language reach, Quocirca provides in-depth insights into the views of buyers and influencers in large, mid-sized and small organisations. Its analyst team is made up of real-world practitioners with firsthand experience of ITC delivery who continuously research and track the industry and its real usage in the markets.

Through researching perceptions, Quocirca uncovers the real hurdles to technology adoption—the personal and political aspects of an organisation's environment and the pressures of the need for demonstrable business value in any implementation. This capability to uncover and report back on the end-user perceptions in the market enables Quocirca to advise on the realities of technology adoption, not the promises.

Quocirca research is always pragmatic, business orientated and conducted in the context of the bigger picture. ITC has the ability to transform businesses and the processes that drive them, but often fails to do so. Quocirca's mission is to help organisations improve their success rate in process enablement through better levels of understanding and the adoption of the correct technologies at the correct time.

Quocirca has a pro-active primary research programme, regularly surveying users, purchasers and resellers of ITC products and services on emerging, evolving and maturing technologies. Over time, Quocirca has built a picture of long term investment trends, providing invaluable information for the whole of the ITC community.

Quocirca works with global and local providers of ITC products and services to help them deliver on the promise that ITC holds for business. Quocirca's clients include Oracle, Microsoft, IBM, O2, T-Mobile, HP, Xerox, EMC, Symantec and Cisco, along with other large and medium sized vendors, service providers and more specialist firms.

Details of Quocirca's work and the services it offers can be found at http://www.quocirca.com

quocirca