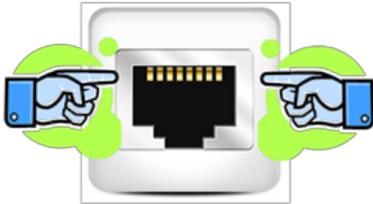


Desktop VDI May Be Ready for Prime Time but Is the Network?



Lori MacVittie, 2012-06-02

#VDI #quasar #mobile The proliferation of mobile devices is pushing VDI closer to being “the solution” of the year to resolve the increasing complexity – and costs – associated with consumerization.



Considering the innate differences between just the two most popular mobile operating systems – Android and iOS – gives rise to understanding how costly and complex an infrastructure might need to be to support both. It’s not at all unlike the issues with server virtualization. Management and delivery architectures require different solutions depending on the platform, so despite potentially costly investments to scale, organizations are often staying single-vendor with respect to

its virtualization platform strategy.

Organizations had taken that approach – standardized on a single mobile platform – only to discover that employees blatantly ignored such mandates and began using whatever they brought from home. Worse, they expected support when applications didn’t work quite right.

Thus IT is stuck trying to figure out how to efficiently deliver, secure, and manage applications to multiple operating systems right now. Not tomorrow, not next week. Today.

VDI is thus rearing its head as a viable solution; one that promises consistency regardless of platform, without worry about Bob wanting to access corporate resources via his Internet-enabled HDTV. For the most part, experts and implementers deem VDI ready to meet the challenge. But what they haven’t asked, nor considered, is whether the *network* is ready for VDI.



Desktop Virtualization Ready for Prime Time

The appeal of VDI remains the same: it improves flexibility, simplifies administration and boosts security. What has changed are ongoing price drops and a growing need to seamlessly manage an [IT infrastructure](#) that includes desktops running Windows, Mac laptops using Apple OS X and mobiles devices using iOS and Android. In many cases, VDI streamlines data exchange and accessibility in an increasingly bring-your-own device (BYOD) IT world.

VDI OFFLOADS the PROBLEMS to the INFRASTRUCTURE

Interestingly enough, the problems with delivering applications to multi-endpoint environments do not actually disappear with the introduction of VDI. Oh, the problem of supporting every device under the sun is neatly resolved, but other problems quickly arise, and these are not necessarily easy problems to solve.

1. Roaming

The issue with roaming isn't just that of a device roaming across service boundaries or WiFi networks, it's roaming geographically. VDI deployments carry with them some strict and often constraining infrastructure requirements that are not easily overcome without the assistance of infrastructure. Typical network environments are ill-prepared to deal with not just the basic constraints but the resolution to those constraints. They lack the flexibility of an application delivery tier to mediate between roaming users and virtual desktop infrastructure.

A user that roams between two completely different network types may in fact appear to be two different network users from an IP perspective. While we know we must one day eliminate our dependency on IP addresses, today it remains a factor that must be addressed. Users who suddenly move from one network to another may cause undue stress along the entire infrastructure – but especially on VDI servers that maintain their understanding of users based on connections, which base their identification on IP addresses. A mediating connectivity layer such as an application delivery tier eliminates not the dependency, but the impact on the actual VDI servers and applications by becoming the endpoint and handling the possible volatility in device identification on behalf of the services, mitigating the impact by absorbing and managing it itself.

2. Network Impact on Performance

What, exactly, is the network over which the mobile device is connecting? Is it WiFi? Is it the mobile network? The network over which a device is connecting has a significant impact on performance, particularly [from the perspective of the end-user](#). It isn't so much a question of whether or not the network is fat enough, it's whether or not the external (read: out of IT control) network is fat enough, or fast enough.



Latency is the biggest concern among networking pros considering a VDI deployment, according to an informal survey of 1,197 VMworld 2010 attendees conducted by storage vendor Xiotech and WAN optimization vendor Silver Peak. The vendors say 62% of respondents named latency as their top VDI network consideration.

A minority named other WAN-related issues as concerns, such as the ability to shape or prioritize traffic (7%) and [bandwidth](#) (6%).

[-- VDI over the WAN: How latency affects virtual desktop performance](#)

While WAN Optimization and similar technologies can certainly address issues when a WAN is involved, it won't necessarily be of assistance when mobile devices experience issues over WiFi or mobile networks or any configuration in which there is no control over both endpoints. Yet the same network problems will plague these devices, and likely with more frequency than remote desktops over IT controlled WAN channels.

3. Scale of Dependent and Primary Services

Likely the most overlooked of all is the scalability of dependent network services. Simple things like NAT, like application access control, like network security infrastructure that must deal with the possibility that users will be logged in from several places at the same time, trying to access different resources. It's the scalability of network security devices that suddenly must contend with connections coming from a wide variety of networks and locations, and must decide – quickly – which of those connections will be allowed, and which should – and must – be denied.

It's also about the ability of applications themselves to scale when faced with suddenly very different network profiles that significantly impact the capacity and load on existing services. Applications that have performed well with capacity X suddenly perform poorly even though concurrent user counts have not changed. This is because the network characteristics may have changed in such a way as to change the way in which the applications are served. Users connecting over the LAN are able to receive content quickly and thus reduce the overall burden on server infrastructure by clearing queues and releasing connections that can be used by other users. Users connecting over mobile networks are not able to receive content as quickly and thus increase the burden on server infrastructure by receiving content more slowly and taking more time to complete a session. This reduces the capacity of server infrastructure and may require additional scaling to ensure consistent, acceptable performance levels across all device types and users.

Thus, while VDI may be ready for prime-time, and is certainly a valid solution to the problem of consumerization with respect to mobile device proliferation in the enterprise, the network may not be ready for VDI – regardless of endpoint form factor.

VDI, like server virtualization and [cloud computing](#), will necessarily change the way in which we architect and ultimately view the network because of the very characteristics that make these technologies appealing – abstraction, elasticity, dynamism. These characteristics make it more difficult to deliver applications and services like VDI because of the volatility and diversity they introduce into the data center and impose on the network.

New architectural and technological solutions will be required in the network to manage such issues as they arise.

F5 Networks, Inc. | 401 Elliot Avenue West, Seattle, WA 98119 | 888-882-4447 | f5.com

F5 Networks, Inc.
Corporate Headquarters
info@f5.com

F5 Networks
Asia-Pacific
apacinfo@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.com