



J-

Thanks for the opportunity to review the VDI initiative proposal for your client.

This is a very professional and well-thought out proposal and project. Since this is a pilot project and since VDI is not as 'mainstream' as server virtualization is, I wanted to make several observations and comments on the proposal/project.

- ❑ I have reviewed several configs for VDI and monitored a few VDI projects over the past year, and by far the most difficult issue for this architecture is application compatibility and (closely related) application support from the software vendor. The general rule of thumb is that the more 'generic' an application is (think Microsoft Office or Web browsers), the more likely it will run well in VDI. The corollary is that the more specialized (in hardware or OS dependency) it is, the *less* likely it will work well (if at all) in a VDI environment. In this case, the VDI is to be used in laboratory settings, and whatever software is used in those settings will likely be the issue. We know that any kind of device interface dependencies are especially troublesome, such as lab equipment requiring serial-port connections or specialty USB-devices such as signature pads (which sometimes get recognized as mouse devices by VDI). (Special attention should be given to any high-end lab equipment that might have been procured with special Grant monies which would require integration into this infrastructure.) I would assume that XYZ has already tested those apps in a 'test lab' VDI environment, utilizing already-installed systems (without a \$200K price tag). [This is recognized in the vendor proposal by excluding Applications from the SOW, other than those mentioned in the SOW.]
- ❑ Closely related to the above--for those applications and devices--would be the issue of software support by the software vendors/ISVs for those elements in a VDI environment. We have seen several cases in med-tech and manufacturing where the software vendor would NOT support their specialty software in ANY virtualized environment. XYZ should also be sure to make explicit any limitations of that support (e.g. is a 'vanilla' hosted environment supported, but not ThinApp or not PCoIP?).
- ❑ This last point raises the issue of patch management for those apps. One of the main benefits of VDI is the centralization of images and easier patch management. Automated updates/patching processes are offered as benefits, but this assumes either a homogeneous footprint (i.e. everybody can/wants to be upgraded to the next version or patch level) or a technology that sandboxes applications allowing multiple versions to run on an OS simultaneous (e.g. ThinApp). Since homogeneity cannot be assumed for the future (especially if student-owned devices will be incorporated in the future), the software vendor support for ThinApp becomes even more critical--and this must be determined beforehand.

- ❑ The pricing comes out to about \$2K per user, which is about right or a tad high. The capex for VDI is significantly higher than for traditional desktop users, and the gains (cost wise) are supposed to come from life cycle opex savings and/or cost avoidance from slower refresh. Deployment costs (including capital costs) are estimated to be from 1.3x to 1.5x that of a PC [Gartner].
- ❑ Capex costs are not expected to decline much for VDI, largely because the market for it is limited. VDI penetration of enterprise desktop footprints is not expected to rise above 15-18% (IDC) or 10-12% (Gartner). [Current estimates for 2011 are 3%.]
- ❑ The proposal is on-target in arguing for a move to 10GbE connections, since that will be required for many VDI situations. As VDI was being piloted in large organizations over the last, say, 3 years, the bottlenecks that occurred were (in this order) RAM constraints, then NIC bottlenecks, and then storage access (IOPS).
- ❑ But VDI itself does not require any sort of 'converged architecture' *per se*. There are no real advantages (for the desktop user, at least) of converged network adapters on the servers nor the convergence capabilities of a Nexus box. These are not bad technologies--of course--but they should be recognized as peripheral to the VDI pilot. "Non-converged" servers, switches, and storage would do just as well in a VDI test (as opposed to in a large data center move to virtualized infrastructure). XYZ should be sure that they are not paying *more* for these 'Cisco-preferred' technologies than they would be for non-converged gear. If the pricing is at parity with non-converged options, then XYZ should exploit this (recognizing, however, that these approaches are 'propriety' or at least 'vendor-specific' and could conceivably prevent later shifts to other/competing technologies).
- ❑ In fact, these additional 'novel' elements in the design--on top of the VDI/PCoIP innovations--may create some ambiguity in the results. If the VDI pilot successfully hits all the performance goals, it will not be a simple matter to isolate the 'success' to the VDI, the converged servers, the converged fabric, or the new SAN. On one hand that is okay--we want to know a *complete solution* that works. But on the other hand, we might not be able to determine whether we 'overbought' on the SAN or on the converged technologies. XYZ obviously would not want to have to replicate this large investment for every 100 VDI users(!), but sizing future systems might be complicated by these additional variables. [The proposal has an upgrade path specified, for deploying this configuration for server virtualization, but not for scaling the VDI footprint. This sizing issue is raised explicitly in the proposal and the vendor has explained the uncertainty clearly.]
- ❑ The above point might not be relevant, if XYZ already has some VDI in place. There is a reference in the proposal to "the current DVI environment", so perhaps this issue (and perhaps many of these issues) have already been addressed.
- ❑ Sizing is the next major issue in most VDI implementations, and is the subject of considerable debate. There are the 'validated' limits of VMware software, and there are client experiences. VMware validated numbers include 16 virtual desktops per

core in Nehalem systems (for vSphere 4.1), but current experience places the 'safe' number around 12 VMs per core. Given those numbers, the configuration of 24 cores (four 6-core Xeons) would have a 'validated' max of 384 (24x16) or a 'safe' max of 288 (24x12), apart from back-end performance issues (e.g NIC, SAN).

- ❑ The RAM issue in sizing is dependent largely on the mix of client desktop images supported. The more variety in those images, the less memory optimization we will get from the various 'page sharing' technologies; and the more Windows 7 we have to support (versus XP) the larger the RAM/VM requirements will be (memory sharing with XP got up to 80% in some cases--we are barely getting 40% with Win7).
- ❑ From a pilot standpoint, one of the major challenges is predicting future system behavior 'under load'. To do performance testing for 110 systems on a configuration that could support (by 'spec') 384 users, may give a possibly misleading result (unless performance is 'bad', of course). Response times for regular 'office-class' applications may be abnormally high, and performance on graphics-intensive engineering or math applications likewise. The vendor proposal rightly specifies a close look at performance testing, and this should be attended to by XYZ with all due diligence. [This is, of course, a problem with ALL pilots, and not something specific to VDI or the technologies under test.]
- ❑ One minor point on proposal terminology: I would assume that the vendor would support *embedded* versions of Windows (in the thin-client devices) since those devices are in their proposal. However, this is not made explicit in the list of supported OS's in the document and should be added to the wording--just for contractual clarity.
- ❑ On life cycle timing and lease term issues, there are two points that should be mentioned, one about display technology rate of change and the other about pilot-type investment risk. The proposal speaks of "six years plus" as a possible refresh cycle for thin clients, and while this is conceivable for the Wyse units (slow rate of change in the technologies inside them), it is not necessarily so for the Samsung displays. Of all the technologies undergoing rapid innovation today, surely the user interface designs are some of the greatest. The use of touch-screens, multi-touch, gestures, face recognition via screen-web, and even screen-based digital pen/ink technologies will be a major part of student experiences (and expectations!) over the next several years. The Samsung units are excellent devices, but they will simply not be able to 'stretch' into these other interaction modalities. XYZ should consider putting the displays under the shortest lease term possible, to facilitate replacement without undue change expenses.
- ❑ The second issue has to do with lease term on 'pilot' equipment. The central point of a pilot is to determine IF some architecture should be invested in. Although this design is excellent and would certainly achieve high levels of performance, there is still an element of risk associated with it. If any of the issues raised above become 'show stoppers' after the equipment has been put in deployment, the cost to 'send it back' can be considerable--and this is simple business investment risk. The servers,

switch, and storage could all be redeployed elsewhere within XYZ, but this is not the optimal way to do architecture upgrades ("we are stuck with this gear, so let's use it *here*, instead of getting the exact equipment we *really* wanted to use here..."), and the thin-clients and displays may not be useable anywhere (at least not at the level of what they cost). So XYZ should consider its 'fallback plan' in case the pilot outcome is not what was hoped, and should look at shorter lease terms (to reduce costs of project abandonment). Once the pilot proves that the architecture will work, then leases can be renegotiated as part of the procurement cycle for the large-scale VDI deployment.

I think XYZ--as a leader in its engineering education field--is wise to consider this technology, and the benefits articulated so well in the proposal by Generic Solutions would be of great benefit to achievement of their organizational mandates. But the issues that others have faced along this road (e.g., application issues, vendor support, capacity planning) should be worked through carefully with the vendors, the internal IT experts, and Macquarie.

Thanks again and let me know what I can do to help--Glenn, May 2011

Note: I have attached two articles on the subject that might be of use to XYZ, one on sizing on one on successful VDI deployments.