

## Fermilab comments on Draft IG report on "Management of Data Centers at Contractor Sites"

Victoria White, Head Computing Division – 6/27/08

We believe that this report may be misleading in how it defines a "data center". It appears to be claiming that consolidating all servers and scientific computing in one physical place is clearly the most cost effective. At Fermilab we have almost all application hosting, servers and data serving consolidated into a data center that is administered and managed as one entity – but in fact spans 3 physical locations. We gave the auditors lots of documentation about how we had done careful cost-benefit analyses when taking this approach. In fact we might have hoped to be called out for special mention in the way in which we have cost effectively re-used building space for computer centers and also built one space at a time to optimize each space for maximal utilization of space, power and cooling together. The three physical locations are linked by many hundreds of fibers allowing us to treat them effectively as one space, but in fact to gain better efficiency in power and cooling as we use different spaces for systems with different physical infrastructure needs. We believe this is a "best practice" that should be called out for special mention. Many factors were taken into consideration in this – looking at the entire picture of what it would cost to upgrade existing space in one center, to build completely new space. Issues of availability of power feeders, condition of the buildings all have to be considered and a superficial look at "how many centers" does not adequately take these factors into account.

In addition although some labs were singled out for criticism of allowing multiple organizational units to have "data centers" and run core services such as email we feel that it is only fair to name those organizations, such as Fermilab, who essentially have consolidated all such services (with one small exception) as doing things well.

There is a statement on Page 1 "The Department had not always taken advantage of opportunities to improve the operational and energy efficiency of its data centers at contractor sites.". The issue of cost of this is not adequately considered. Perhaps a deeper comment in the report might look at the money that DOE has as a whole to invest in infrastructure that would allow contractors to improve operational and energy efficiency. There is no doubt that upfront capital investment is needed in order to, over a longer period of time, save money on operational costs. Why is this factor not brought out in the report?

Another best practice that we feel Fermilab should take credit for, but which was not even raised by the auditors in their interview with us, is the way in which we buy the large amounts of scientific computing servers. The procurement process does take account of energy efficiency as one of the evaluation factors in the procurement.

A large part of the report is claiming that virtualization is the key to cost savings. Fermilab DOES employ virtual hosting extensively. There is extensive use of VMWare for both business systems (HR, Accounting, etc) and for core services such as web servers. There is also increasing use of Zen for virtualization of many scientific computing services – across all of our scientific programs at the lab. No-

one can recall that specifically being on the list of questions asked by the auditors and so this should be corrected and Fermilab should be added to the list of three institutions that are called out as effectively employing virtual hosts.

However, from Fermilab's viewpoint the report is naïve in its assertions about the potential cost savings involved in virtualization and the potential cost savings involved in energy efficient equipment.

Firstly the statement that "Without improvements in Department guidance and contractor monitoring of server utilization, contractors and field sites will continue to spend more than necessary on data center operations." is highly questionable. Do the auditors really think that the reason sites might not use servers as efficiently as possible is because they lack department guidance and that would make all the difference? At Fermilab we are constantly striving to improve efficiency. Our folks attend conferences to look at data center efficiency – such as Uptime, Afcom, 7X24 and strive continuously to make long term plans that replace old equipment, consolidate and use our people, spaces and equipment most effectively. The report is very negative and does not look at the slope of change – which is continually improving. Furthermore it seems to be based on an old snapshot now. We have provided documentation 3 separate times to 3 different sets of auditors – and some of this dates back to 2006. In the lifetime of technology and computer centers 2006 is a long time ago. The limitations on moving more quickly to implement change are in manpower, capital equipment dollars and cultural change. It is clearly the job of management, such as myself, to move forward as quickly as possible, given these constraints. Some acknowledgement of the work and struggle that does go on to do this, in a report such as this, would be welcome.

Secondly the assertion that at Sandia having over 100 virtual machines led to an annual savings of more than \$700,000 is highly questionable – and clearly cannot be simply extrapolated to the \$7.3M savings figure cited. The savings yielded through virtualization is complex and not the whole story. The amount of effort it takes to run systems depends not on the number but on the uniformity of the systems. At Fermilab we run many thousands of compute servers for scientific use with very few people – because they are all the same and all used in a certain limited way. I strongly object to the simplistic model for extrapolating potential savings for the department – from one or two questionable datapoints and without a better economic model underlying this for the cost of running servers, services, physical infrastructure, power, cooling – based on reliability and redundancy of services required. All of these factors must be considered in a costing model – and fed into a model of potential savings. There is no doubt that all of us could do better and are striving to do so – but this report is plain misleading.