

AI

Hi Bob,

I know AI is a growing interest with your team, so passing this along for your assessment.

We're working with a vendor named Liquid who builds really, really dense GPU servers (think 20-30 GPUs per server).

The resulting denser servers = faster systems and less infrastructure/power needed for platforms like AI, HPC, VDI, etc.

Further, they leverage components from Dell, Nvidia, AMD, etc (brands already in your environment) and, if helpful, the flexibility allows GPUs assigned to your AI systems by day to be reassigned to other environments during off-peak hours.

Are you open to a quick session to see if there is value in digging deeper on this?

Thanks,
Bob