


AI Decision Tree: How to Match AI Readiness to the Right AI Solution

The maturity of AI in your organization has a huge impact on the type of AI solution that will work best for you and what sort of hardware vendor would make the best infrastructure partner.

Understand your AI deployment readiness and discover the best type of AI system and system provider for you:

We know we need AI to extract greater value from our data but we want to move carefully into AI.



- Consider a workstation (if you will only have one or two people conducting the analysis) or a small, project-scale cluster.
- For AI workstations, ensure you buy from an organization that specializes in GPU-accelerated computing and not just general hardware. This will ensure they have the relationships with the right component suppliers and have access to engineers who can ensure strong ROI for AI workloads, based on the understanding that AI puts a significant burden on hardware.
- For small clusters, ensure the vendor you are working with is up to speed on the latest technology and has strong partner relationships. This means that you will have access to the components you need as well as be able to take advantage of new developments coming on scene.

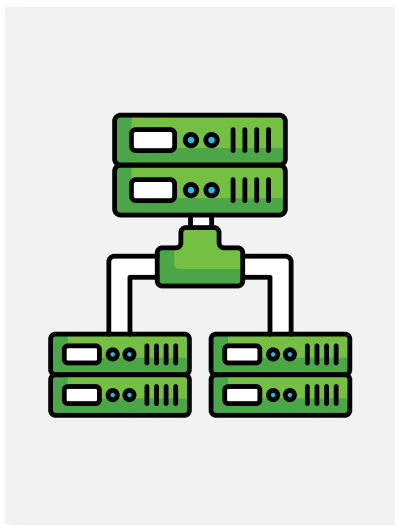
We have an existing AI system but its old technology that needs to be replaced and we're sure we can find something more up-to-date at a reasonable cost that gives us more value.

- Consider a custom-tailored cluster built using a modular methodology, as this allows customization to your key concerns (i.e., processing speed, processing power, etc.) but that can grow with you or just change as new technology arises.
- For cluster upgrades or mid-sized systems, remember that your unique needs are much more likely to be met with a vendor that specializes in cluster design, rather than one that also needs to devote scarce resources to designing, marketing, and selling laptops, cloud-related technology, or other non-core offerings. The cost of supporting these



huge businesses means overhead costs are more likely to be passed on to clients and there is little ROI for these large firms to spend the engineering or manufacturing resources to customize solutions for each client. Also consider whether you might need some professional services help (i.e., training or just rack and stack) to help you make the most of your cluster and ensure your chosen cluster partner can provide that.

We've been using AI for a while and feel that we could handle a bigger, more robust system and get more value from AI.



- Consider a cluster that you can build upon over time but don't get caught up in having technologies with buzz. Instead, focus on technologies you really need and allow your infrastructure design team to use more traditional technology in other areas so that you can spread your budget out further, for more nodes (which are key to training the system).
- For larger clusters, you might consider a vendor who can offer a "supercomputer-in-a-box" if cost is not an issue, you are not concerned about vendor lock in, and you have no need to customize your cluster. Otherwise, choose a vendor that uses a building block approach to building clusters, so you can avoid those issues. Ideally, your vendor should also be able to customize hardware on-site, so that you can be assured of quality in this custom infrastructure. Here too, consider whether your team has the time to do all the necessary set up tasks or if you need professional services help (i.e., rack and stack) and whether your chosen cluster partner can provide that.

Interested in finding the right AI solution for your specific needs?
Contact Silicon Mechanics

📞 1-866-352-1173

💻 SiliconMechanics.com

✉ info@SiliconMechanics.com