Panavisor Compared to a Public Cloud

How is Panavisor different from a public cloud solution? Control. The differences between Panavisor and other computing architectures reflect the allocation of control.

At one end of the control spectrum are local servers at the user's own premises. The on-premises data center puts everything is under the control of the user – the user selects the desired applications procures the platforms required to host the applications. The systems are accessed over the local LAN, which can be segmented using LAN segmentation tools under the user's control. This LAN segmentation implements security through isolation by placing different uses on different LAN segments, also called subnets. Through proper segmentation, internal processes can be separated from external processes – the "intranet" and the "internet." The users are 100% responsible for their own up-time and business continuity solution. The users have to purchase and maintain the hardware and software, house it, power it, cool it, back it up and administer it. The user can delegate control to a solution provider, who acts as an out-sourced IT department. The independence of each user's systems is assured by physical isolation.

The difference between a LAN and a WAN can be illustrated by the simple every day task of printing. What is the span of control of the print job? In public cloud environments, the print job reaches the internet enabled printer over the internet – the print job leaves the control of the user. In a LAN environment the printer is a local device which receives the print job over the LAN – the print job is under the control of the user for its entire life cycle.

At the other end of the spectrum is the public cloud. Users of public clouds have no control over their data center – the control of all IT functions is based on the cloud provider's infrastructure and choices. The user must accommodate to the provider's API requirements for configuring IP addresses, subnets, firewalls and data service functions. The computing requirements of different users reside on shared hardware. While sharing hardware offers significant cost advantages, the data is under the control of the cloud provider. Privacy concerns abound in cloud computing because the service provider can access the data that is on the cloud at any time. It could accidentally or deliberately alter or even delete information. The data is exposed to government data requests which may be made without the knowledge of the user. Security lapses of the cloud provider or other tenants on the same hardware place the user at risk. The user's compute processes and data may be moved from server to server and country to country at the whim of the provider seeking to minimize its operating costs. The user uses web-based applications which are accessed through web-browsers. For established businesses, moving to web applications which are supported by the public cloud will require data migration and user retraining. Business processes which require multiple web applications may be difficult or impossible to implement.

In an effort to provide hosted solutions that more closely mimic the on-premises LAN, cloud providers provide two types of upgrades. The first upgrade is providing dedicated computing hardware instead of shared hardware that is allocated according to the needs of the hosting provider. This addresses many of the security and privacy concerns that arise from using shared hardware. The increased isolation provided by controlling one's own hardware comes at a significant increase in cost and complexity.

The second upgrade is to provide more control over the remote data center environment by allowing the user to control network functions such as internal and external firewalls and load balancers in order to allow network segmentation. These segmentation services can include vLAN services in which virtual LANs are created to segment the network. The cost and complexity of management tools is considerable.

Moreover, this architecture does not solve the problem that users are still required to use web apps accessed through browser sessions.

An interesting blending of shared resources and increased control is the Savvis Symphony service which offers users network controls (internal and external firewalls, and load balancers) on shared hardware.

At the apex of the cloud solution, both in terms of cost and complexity, the user controls a private cloud in which he is managing dedicated compute resources, configuring firewalls, load balancers and VPN's. This environment is more complex to manage than a LAN because of the added key management tasks imposed by VPN's. In this environment, the remote resources provide the same compute services as on-premises servers. In this environment, the user can establish LAN services and does not have to migrate to web applications.

Panavisor provides private cloud services using shared public cloud resources. This is accomplished through our end-to-end architecture which:

- isolates the multi-tenant processes at the hosting provider
- allows the user to establish LAN subnets for enhanced security
- assigns the user's processes to the data center of his choosing
- provides LAN access to the data center

The complexities usually associated with a private cloud are performed by the management functions of *Panavisor* appliances. The dedicated equipment costs of a private cloud are replaced with the cost *Panavisor CloudPort* at the user's site, the user can seamlessly replace its on-premises data center with managed hosted resources. The *Panavisor* architecture allows the user to authorize a solution provider to manage the compute resources on its behalf.