

第 16 回 IT ソフトウェア翻訳士認定試験

<1次試験> 4月14日（日）10：00～15：00

問題 1・2 の両方について解答のこと。選択ではありません。

<問題 1> 全文を訳して提出してください。

Exploiting the business potential of BYOD (bring your own device)

Just as the Internet has changed the way we do business, mobile devices are quickly becoming a similar catalyst. Once deemed off-limits by businesses, smartphones are now being used by hundreds of millions of employees to access corporate data. Tablets are also increasingly being used for both work and leisure activities, with mobile functionality a key driver.

In fact, a recent survey showed we have reached a tipping point in the business use of mobile devices as mainstream business activity. The report found companies were embracing this change: nearly three quarters of the organizations surveyed were looking to develop their own customer mobile applications, well over half were running line-of-business applications, and two-thirds were discussing implementing a corporate ‘app store’.

So why has the use of personal devices as business tools changed from a largely forbidden activity to an acceptable practice? The answer lies in today’s emphasis on business agility. Mobile devices are helping companies to increase workforce effectiveness, increase efficiency, and generally get work done more quickly. Adding something around potential for new business as customers and partners are equally becoming more mobile.

Enterprise mobility delivers on productivity because employees know how to use, and enjoy using, their own devices. BYOD can also reduce capital expenditure as businesses can leverage devices employees may already be paying for. In addition, employees often take better care of devices they have selected and purchased.

However, the rise of BYOD culture also brings its own problems for

businesses; it can increase productivity but can leave an organization with a risk assessment nightmare. A 2011 Gartner report stated that: “Although BYOD is attractive to some enterprises, it poses a range of new challenges for the support organization. No organization can afford to fix an unlimited range of issues on a large portfolio of devices, many of which the organization doesn't own and therefore can't control in conventional ways.”

Organizations are aware of the potential dangers mobility can pose, rating it highest among IT initiatives in risk. They're worried about losing devices, data loss and malware infecting the corporate network through smartphones and tablets.

<問題 2> 下線部のみを訳して提出してください。

Deploying Database Appliances in the Cloud

2 Deployment and Tuning Challenges

Our focus is on deploying and tuning virtual machines running database systems (i.e., database appliances) on large clusters of physical machines (i.e., computing clouds). This raises deployment and computing challenges, which we describe next.

2.1 Deployment Challenges

Creating a database appliance that can easily be deployed in a cloud, and obtaining an accessible, usable database instance from this appliance require addressing many issues related to deployment. These issues are not the research focus of our work, but we present them here since these seemingly simple and mundane tasks can be very tricky and time consuming. These issues include:

Localization:

When we start a VM(Virtual Machine) from a copy of a database appliance, we need to give this new VM and the database system running on it a distinct “identity.” We refer to this process as localization. For example, we need to give the VM a MAC address, an IP address, and a host name. We also need to adapt (or localize) the database instance running on this VM to the VM’s new identity. For example, some database systems require every database instance to have a unique name, which is sometimes based on the host name or IP address. The VMM(Virtual Machine Monitor) and the underlying operating system and networking infrastructure may help with issues such as assigning IP addresses, but there is typically little support for localizing the database instance. The specific localization required varies from database system to database system, which increases the effort required for creating database appliances.

Routing:

In addition to giving every VM and database instance a distinct identity, we must be able to route application requests to the VM and database instance.

This includes the IP-level routing of packets to the VM, but it also includes making sure that database requests are routed to the correct port and not blocked by any firewall, that the display is routed back to the client console if needed, that I/O requests are routed to the correct virtual storage device if the “compute” machines of the IaaS cloud are different from the storage machines, and so on.

Authentication:

The VM must be aware of the credentials of all clients that need to connect to it, independent of where it is run in the cloud.

2.2 Tuning Challenges

Next, we turn our attention to the challenges related to tuning the parameters of the virtualization environment and the database appliance to achieve the desired performance objectives. These are the primary focus of our research work, and they include:

Placement:

Virtualization allows the cloud provider to run a user’s VM on any available physical machine. The mapping of virtual machines to physical machines can have a significant impact on performance. One simple problem is to decide how many virtual machines to run on each physical machine. The cloud provider would like to minimize the number of physical machines used, but running more VMs on a physical machine degrades the performance of these VMs. It is important to balance these conflicting objectives: minimizing the number of physical machines used while maintaining acceptable performance for users.

A more sophisticated mapping of virtual machines to physical machines could consider not only the number of VMs per physical machine, but also the resource requirements of these VMs. The placement algorithm could, for example, avoid mapping multiple I/O intensive VMs to the same physical machine to minimize I/O interference between these VMs. This type of mapping requires understanding the resource usage characteristics of the application running in the VM, which may be easier to do for database systems than for other types of applications since database systems have a highly stylized and often predictable resource usage pattern.

Resource Partitioning:

Another tuning challenge is to decide how to partition the resources of each physical machine among the virtual machines that are running on it. Most VMMs provide tools or APIs for controlling the way that physical resources are allocated. For example VMM scheduling parameters can be used to apportion the total physical CPU capacity among the VMs, or to control how virtual CPUs are mapped to physical CPUs. Other VMM parameters can be used to control the amount of physical memory that is available to each VM. To obtain the best performance, it is useful to take into account the characteristics of the application running in the VM so that we can allocate resources where they will provide the maximum benefit. Database systems can benefit from this application-informed resource partitioning.

以上