Chapter 12
Effective Infrastructure Protection through Virtualization

Dennis C. Guster
St. Cloud State University, USA

Olivia F. Lee
St. Cloud State University, USA

ABSTRACT

In this current digital era, organizations are increasingly aware of the need to protect their computer infrastructure to maintain continuity of operations. This process involves a number of different concerns including: managing natural disasters, equipment failure, security breaches, poor data management, inadequate design, and complex/impractical design. The purpose of this chapter is to delineate how virtualization of hosts can be used to address the concerns above resulting in improved computer infrastructure that can easily be restored following a natural disaster and which features fault tolerant hosts/components, isolates applications security attacks, is simpler in design, and is easier to manage.

INTRODUCTION

Numerous types of disasters, both natural and manmade, can be catastrophic to businesses. Without a well thought out disaster recover (DR) plan, such events can seriously disrupt routine business operations. Often times, it is difficult to comprehend the devastation of an unknown future event, let alone create a comprehensive approach to meet and survive it. The most critical challenges are related to understanding the scope and complexity of DR requirements and the risk of inadequate deployment of recovery efforts. This lack of understanding is especially applicable to small or medium size businesses (Hill, 2008) due to the limited IT resources they have available. Larger firms are able to apply economy of scale to develop an information technology (IT) depart-

DOI: 10.4018/978-1-60960-573-5.ch012

ment that equips them with the basic infrastructure to support the addition of DR mechanisms. Small and medium sized businesses, on the other hand, often do not have adequate infrastructure and since they operate on smaller profit margins, devising and supporting a DR plan can be a huge burden. Recently, Search Security.com reported that disaster recovery often accounts for as much as 25% of the IT budget. Hence, sound disaster recovery planning is a very important undertaking not only due to what might be lost, but also from a budget perspective. Hence, devising a strong DR infrastructure is further justified by the "information intensity" structure of many companies in the 21st Century.

Specifically, for many companies in the 21st century, information resources are their livelihood. The loss or unexpected long term disruption of information or data could have a detrimental effect on business operations. Phillippi (2008) reports that 92% of small businesses that experience significant data loss due to a major disaster go out of business within five years. Indeed, due to the high level of internet connectivity required by most operation functions today, the risks are high and warrant a well thought-out plan with appropriate risk assessment (e.g., Hiles, 1992; Jones & Keyes, 2001; Stephens, 2003). Although security risks of the internet increase the need for an effective disaster recovery mechanism, the internet connectivity is nevertheless advantageous as it can be effectively used in the data replication process. An efficient and cost-effective disaster recovery strategy is to utilize the geographic distribution of the critical components model (Adam, 2002). The connectivity can be inexpensively provided by the internet provided secure transmission methods, such as virtual private networks or VPNs (a way of isolating and double encrypting data sent across the internet), are used, and using the internet can minimize the huge cost of leasing dedicated lines (such as T1 a non-switched digital phone line). The data replicas should be at least 150 miles from the data center headquarters (Phillippi, 2008).

Information resource or data recovery can take many forms. In the past, pools of computers, on which a few members worked together and shared resources, could be used to house backup systems. While this approach still has some merits and fits reasonably within a service oriented architecture approach (SOA), for small businesses with remote sites and existing corporate partners, there are major trust issues to resolve in regard to the partners that make up the pool. An alternative is the use of virtualization which can potentially minimize costs and server density (Safigan, 2008). Organizations can logically partition one high-end computer and place each of its production servers on it in separate zones, thereby reducing management overhead. Another concept to consider is automation as it may significantly reduce the recovery time during an unexpected disaster. Whatever the chosen method, it is important to consider the expensive and on-going personnel costs which could be significantly higher than the additional hardware required.

Another data recovery aspect organizations must consider is the effect of distributed processing in disaster recovery. The advent of distributed processing and cluster computing vastly altered the manner in which data is stored and how access is granted to resources in an enterprise computing environment. No longer are projects simply done on a stand-alone computer. In fact, any given project may share and retrieve resources from several computers. While such processes often improve performance and lead to some degree of fault tolerance, organizations are required to have a resource profile on each required host. As a result, data spread out among hosts, in separate login accounts, can rapidly become an end-user's nightmare. However, by using a global file (such as NFS) and authentication system (such as LDAP) that allows single sign-in capability, as well as a file system that is attached to that sign-in no matter which host is being accessed, the above problems can be rectified.