INTERNAL AUDIT REPORT

FROM:	Audit and Risk Manager	**SUBJECT:**	Management of the Virtualised Server Environment
TO:	Chief Executive (in absence of Head of Corporate & Community Services) ICT Services Manager	**DATE:**	16th December 2013
C.C.	Infrastructure Manager Head of Finance		

---

1 INTRODUCTION

1.1. In accordance with the Audit Plan for 2013/14, an examination of the above subject area has recently been completed. This report is intended to present the findings and conclusions for information and action where appropriate.

1.2. Wherever possible, results obtained have been discussed with the staff involved in the various procedures examined and their views are incorporated as appropriate in any recommendations made. My thanks are extended to all concerned for the help and co-operation received during the audit.

2 SCOPE AND OBJECTIVES OF AUDIT

2.1 Approximately 90% of the server estate within the data centre at WDC is hosted in a virtual environment. This utilises the VMWare vSphere 5.0 product. A virtual environment provides for multiple virtual servers to be created from physical server hardware, with the environment at WDC hosting multiple servers from five physical ESXi servers within a class C chassis. Virtual environments can provide significant benefits but must be designed and configured to meet specified requirements. A review has therefore been undertaken to assess capacity management, security, resilience, performance monitoring and backup arrangements.

2.2 The evaluation was based on the following control objectives:

- Capacity: Projected resource requirements are monitored against current provision on an ongoing basis;
- Security: Administrative access is adequately controlled at the VMWare level;
- Resilience: Resilience is managed to minimise the risk of downtime for virtual servers;
- Performance Monitoring and Resource Management: Usage and performance is monitored and managed across the virtual environment;
- Training and third party support: In-house training and access to third party support is provided to minimise the risk of downtime;
- Backup and recovery: Strategies for backup and recovery in the virtualised environment have been formulated to meet requirements

2.3 The audit approach aimed to assess the approach towards management of the virtual environment through:

- Consultation and discussion with key officers associated with the administration, management and monitoring of the virtual environment; and
- Inspection of configuration settings and examination of documentary evidence recorded at the audit meetings.

3 FINDINGS

3.1 Capacity Planning

3.1.1 It is understood that capacity planning was undertaken by the suppliers when the virtual environment was first scoped around five years ago. This assessed demand from users, anticipated growth and the resilience requirements that would apply. The latter was relevant in terms of the spare capacity that would be required to utilise some VMware resilience facilities.

3.1.2 The two key components for monitoring of capacity are CPU (processing) and RAM (memory). The ICT team state that current utilisation of capacity generally runs at 30% CPU and 50% memory at WDC. This is monitored on an ongoing basis. Inspection of utilisation reports confirmed this estimate to be broadly accurate across the Production cluster of five ESXi physical boxes.

3.1.3 In terms of resilience in processing, the current capacity would allow two of the five physical ESXi boxes to be lost without preventing any live applications from running. The loss of one or two ESXi boxes would result in the High Availability functionality redistributing the relevant virtual machines to the remaining boxes. The capacity available for use by High Availability is maintained by minimising the number of test virtual machines in the production environment.

3.2 Security over the VSphere Environment

3.2.1 Virtual machines are the containers in which guest operating systems and applications run. By design, all VMware virtual machines are isolated from each other. This isolation enables multiple virtual machines to run securely while sharing hardware and ensures both their ability to access hardware and their uninterrupted performance.

3.2.2 The audit reviewed administrative access at the vSphere level i.e. the provision of privileges that enable users to configure objects within the virtual environment. For example, administrative privileges:

- Allow virtual servers to be created and existing servers to be deleted;
- Permit services such as Distributed Resource Scheduler and High Availability to be enabled, disabled or configured; and
- Allow administrative privileges to be assigned to other users

3.2.3 The installation of vSphere involves the creation of the Root account that can be used to access each ESXi box directly or via the vClient. This is designed to ensure that access to the environment can be gained in the event of the vCenter server administration tool being unavailable, see 3.2.4 below. Discussions confirmed that access to the Root account is restricted to the four members of the ICT Infrastructure team with administrative responsibilities.

3.2.4 Administrative tasks for the virtual environment are undertaken at WDC using the vCenter server application. The built-in Administrator role is the most powerful role for vCenter Server, as those assigned to the role have full administrative privileges over all objects in the environment. This includes control over the virtualization layer that controls the hardware and allocates the hardware resources among the virtual machines; the virtual machines that host WDC applications; and the virtual networking layer. It therefore also allows users to configure services such as Dynamic Resource Scheduler and High Availability and to manage all privileges and permissions within the vSphere environment. Access to the Administrator role should therefore be strictly controlled and subject to periodic review. Inspection of the membership of the role confirmed, however, that the following are assigned:

- 2 local Windows server users (for the virtual server that hosts vCenter server);
- 4 Active Directory (AD) users; and
- 3 AD groups (Enterprise Admins, Domain Admins, and Schema Admins). These groups contained 14, 2 and 28 user accounts respectively and included a combination of individual users and service accounts.

Risk
Unauthorised access to VCenter/EXSi privileges

Recommendation
The current assignment of user accounts to the Administrator vCenter server/ESXi role should be reviewed and revised as necessary to ensure that only those accounts with a genuine operational need are assigned with full privileges for the vSphere environment.

3.2.5 The Administrator role can also be assigned on a more restricted basis by assigning the role to users or groups but for selected objects only. At WDC, this is applied by folder whereby one folder contains one or more virtual machines. Three folders have been created for the Production cluster. These are:

- App Supp Full Access;
- VMWare View; and
- PC Support Access

Discussions confirmed that the access to the above has been honed over time to ensure that access is commensurate to need. Inspection of the user list of users with Administrative access to the App Supp Full Access folder showed that only the Head of ICT plus members of Application Support had access.

3.3 Resilience

3.3.1 Discussions confirmed that resilience measures have been built into the environment. The C class chassis has dual power supplies, two onboard controllers and all key components are duplicated. Connections to the SAN and LAN are also dual-homed. In addition, vSphere facilities have been utilised to provide resilience. These include vSphere Distributed Resource Scheduler, a feature that helps improve resource allocation and power consumption across all hosts and resource pools; and High Availability, a feature that ensures that all virtual machines that were running on a failed ESXi host box are automatically restarted on different hosts in the same cluster.

3.3.2 Inspection of settings confirmed that DRS is enabled. The 'Fully Automated' option has been configured for DRS, with an aggression threshold of 3 out of 5 to limit unnecessary activity. This approach reflects the availability of the spare capacity. It is understood that advanced settings have also been configured for DRS to specify that two virtual machines running Exchange cannot be moved to the same ESXi box. It has also been specified that the SQL Server virtual machine cannot be moved to another ESXi box by DRS. Inspection of settings also confirmed that High Availability is enabled.

3.4 Performance Monitoring and Resource Management

3.4.1 Reports of CPU and memory usage can be accessed in real time via vCenter Server. These can be run to show this data at:

- Overall physical ESXi box level and;
- At virtual machine level within each of the ESXi boxes

These reports are reviewed on an ongoing basis as part of the daily administration of the environment.

3.4.2 In order to minimise the risk of downtime and to maximise performance, the physical hosts within the environment are replaced with upgraded hardware every 3-4 years. Additional memory is also added from time to time where the vCenter monitoring reports indicate a need over an extended period.

3.4.3 The vCenter Server administration tool provides the means to monitor resource usage and issues within the virtual environment. The 'Alarms' vCenter function allows alerts to be configured to record triggered events, highlight these events within the vCenter server screen and notify specified individuals via email. Discussions have confirmed that these facilities have not been specifically configured, however. As the vCenter tool tends to be in use each working day, reliance is placed upon the default triggers, the onscreen highlighting of alarms within the vCenter server screen and administration staff to identify any issues outside of the defaults.

Risk

Delays may occur in identifying and addressing issues within the virtual environment.

Recommendation

The current configuration of the Alarms facility within vCenter server should be reviewed to ensure that all key events that may adversely affect users will result in a trigger being invoked and notifications being issued by email to the VMWare administration staff.

3.5 Training and Third Party Support

3.5.1 The primary administrator within ICT for the VSphere environment is Andy Walsh, with Andy Watts being the main deputy. Richard Bates, John Molloy and Lee Millest are also skilled in administering the product. These ICT team members have all been on formal VMware training.

3.5.2 External VMware support is provided by Organised Computer Services Ltd (OCSL). The contract includes hardware maintenance plus support and also provides support for backup issues. The services of Organised Computer Services Ltd are rarely used as the four trained members of ICT are able to deal with most issues. The company will be used to assist in the migration from VSphere 5.0 to vSphere 5.5, however. This is planned for late in quarter 2 2014/15.

3.6 Backup and Recovery Strategy

3.6.1 There is no standby environment at WDC for the virtual environment. A disaster recovery contract is in place with Phoenix, however, to ensure that replacement hardware can be acquired in a timely manner if a significant event affects the virtual environment hardware and storage.

3.6.2 The recovery strategy is therefore based upon:

- The resilience measures in place for the processing environment e.g. the use of High Availability for use where up to two ESXi boxes within the chassis fail or suffer physical damage;
- Resilience measures within the storage provision e.g. RAID for disks and dual power supplies;
- The use of backups to enable restoration using server hardware and storage hardware acquired by WDC following a physical incident. This hardware would currently be provided by Phoenix under disaster recovery contract provision. Backups would also be used to restore in the event of a logical issue, such as data loss or corruption.

3.6.3 HP Data Protector is used to manage backups in the virtual environment. Multiple backup jobs are configured for the environment, with each job covering one or more virtual machines. The content covered for each virtual machine under each job depends upon the recovery strategy for the application concerned. Where an application involves a very large database, an incremental backup approach is taken to record a full backup once a week and then incremental backups for the remaining days. Applications with less data are generally subject to daily full backup. Weekly backup tapes are taken off site each Monday and are stored in a safe at the Town Hall. Daily backups remain on site.

Note: Audits of the overall approach to disaster recovery and of the approach to backup and recovery have been undertaken in 2011/12 and 2012/13 respectively. Please refer to the audit reports for details of recommendations raised.

4. CONCLUSIONS

4.1 The provision of the virtual environment is designed to minimise the physical hardware required in the data centre whilst providing significant benefits in terms of resilience, capacity planning, performance monitoring and recovery capability.

4.2 Two issues have been raised in the report. One is rated as medium priority and one as low priority.

4.5 In view of the above issues, the findings are considered as giving SUBSTANTIAL assurance that appropriate controls are in place to effectively manage the risks relating to remote access and the use of portable devices.

6. MANAGEMENT ACTION

6.1 Recommendations to address the issues raised are reproduced in the appended Action Plan for management response.

Richard Barr
Audit and Risk Manager

ACTION PLAN

Management of the Virtualised Server Environment Devices – 16 December 2013

Rec. No.	*Report Ref.*	*Recommendation*	*Risk Rating* [1]	*Responsible Officer*	*Management Action*	*Target Implementation Date*
1	3.2.4	The current assignment of user accounts to the Administrator vCenter server/ESXi role should be reviewed and revised as necessary to ensure that only those accounts with a genuine operational need are assigned with full administrative privileges for the vSphere environment.	Medium	Andy Walsh, Database Administrator	Undertake review as per recommendation	31st March 2014
2	3.4.3	The current configuration of the Alarms facility within vCenter server should be reviewed to ensure that all key events that may adversely affect users result in a trigger being invoked and notifications being issued by email to the VMWare administration staff.	Low	Andy Walsh, Database Administrator	Undertake review as per recommendation	31st March 2014

1 Risk Ratings are defined as follows:

Low - Minimal adverse impact on achievement of the Authority's objectives if not adequately addressed.
Medium - Moderate adverse impact on achievement of the Authority's objectives if not adequately addressed.
High - Requires urgent attention with major adverse impact on achievement of Authority's objectives if not adequately addressed.