Technical Report

IBM FileNet P8 Hot Backup Solution Blueprint

Nathan Walker, NetApp
Month Year | TR-4023-0212

IBM FILENET P8 HOT BACKUP SOLUTION OVERVIEW

Many enterprises today rely on enterprise content management (ECM) solutions such as those built on IBM FileNet to help them managed structured data, unstructured data, and business process management. NetApp® storage solutions enhance IBM FileNet by delivering time-efficient hot backups, outstanding data protection, and business-ready flexibility to help you go further, faster.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1 INTRODUCTION

The amount of unstructured data managed by today's enterprises is growing at an overwhelming rate. New corporate regulations and compliance requirements are making data management even more difficult. To tackle these challenges, many organizations have chosen an enterprise content management (ECM) solution to address these complex requirements.

The IBM® FileNet® P8 Platform is a next-generation ECM solution with comprehensive business process management and compliance capabilities. FileNet P8 addresses the most demanding compliance, content, and process management needs for your entire organization.

While enterprise records management, advanced case management, and information lifecycle problems are solved by FileNet P8, many organizations fail to create a storage foundation to support ECM backup and disaster recovery. This paper describes a solution blueprint that can be used to provide a storage foundation to lower data management costs while simultaneously addressing hot backups and disaster recovery for FileNet P8.

2 EXECUTIVE SUMMARY/PROBLEM STATEMENT

NetApp provides the necessary enterprise storage platform and services to store, manage, and protect the numerous data types underlying the entire FileNet P8 data infrastructure. By taking advantage of the tight integration of FileNet P8 with NetApp storage systems, you can create a scalable and cost-effective ECM platform that is highly available and secure. In addition, NetApp storage platforms provide unique capabilities to enhance your FileNet implementation in ways that reduce cost, risk, and cycle time. These capabilities include application-wide, application-consistent "hot backup"; storage savings up to 50% for document versioning; and the ability to easily leverage production data for fast and secure development and testing.

FileNet P8 has numerous server roles, datasets, and I/O access patterns with associated storage and management requirements. Because there is no "central brain" in FileNet P8, an application-wide, application-consistent hot backup is only possible with the NetApp storage platform. A typical FileNet P8 environment uses several Java® EE application servers, databases, and standalone applications such as IBM Legacy Content Search Engine (Verity) or Content Search Services. Most organizations find it nearly impossible to back up and replicate all of these components without negatively affecting the business services with a full shutdown of the application. By contrast, NetApp provides a cost-efficient and scalable storage foundation to back up and replicate these IT processes *without* affecting user access.

NetApp will help you lower FileNet P8 costs through improved automation, storage efficiency, and data protection. By modernizing your IBM ECM infrastructure on NetApp storage solutions, you reduce risk and increase competitive advantages. NetApp helps you go further, faster.

3 AUDIENCE/SCOPE

Traditionally application backup has been approached as a separate concept from disaster recovery, as if each existed in a vacuum. Both have the central focus of managing data to protect business processes. A recent study [(1)] has noted that the NetApp Snapshot™ and replication technologies simultaneously address backups and disaster recovery. This document introduces the enabling NetApp technologies that will help you to plan and deploy a hot backup solution for your IBM FileNet P8 environment. These same application hot backups can be replicated to one or more disaster recovery facilities for maximum data protection. This paper is written for backup administrators, storage administrators, and ECM architects who make recommendations and investigate technologies that address their ECM needs. It is not intended as a replacement for FileNet manuals or product training. For an introduction to IBM FileNet P8 on NetApp storage solutions, refer to TR-3927. [(2)]

The primary goal of this paper is to help you create a hot backup infrastructure for FileNet P8 4.5.x. The concepts, architectural considerations, and NetApp components described herein are general enough to allow you to extend this to other distributed enterprise applications such as IBM Content Manager on Demand, IBM FileNet Image Services, or indeed non-IBM applications. The purpose of this approach is to build “factories,” not “castles,” in that a castle represents a single, custom monolithic design that is not flexible to change and cannot be extended to other enterprise applications. A factory, by contrast, can be replicated across geographies and organizations with only minor adjustments for each application. In this way you will make some minor application-specific changes in your scripts as you upgrade and extend your ECM landscape, but the basic architecture and overall solution remain the same. As an example, refer to TR-3709 for a demonstration of hot backups with Symantec™ Enterprise Vault™ [3]. This paper follows the same design philosophy.

The initial lab proof of concept testing for this paper was based on Content Engine 4.5.1, Process Engine 4.5.1, IBM WebSphere 7.0, IBM DB2 9.7, and Verity K2 6.5.1, all on several virtualized Microsoft® Windows® servers. Therefore, a certain level of modification might be required for your FileNet P8 landscape. For example, instead of using the Snap Creator™ module for DB2, you might be required to use the SnapManager® module for Microsoft SQL Server® or Oracle®.

Two of the main structural changes with the advent of FileNet P8 5.0 are the introduction of IBM Content Search Services and the change in Process Engine from a C++ application to a standalone Java application. Nevertheless, there is WebService and REST API backward compatibility, as well as feature parity and transparency to existing applications, which are written to the public APIs supported in the ported FileNet Process Engine. This compatibility will ease the application of the methods described in this paper to an existing FileNet P8 5.0 or to a soon to be upgraded FileNet landscape.

This reader of this paper is assumed to have a basic understanding of ECM concepts and NetApp storage technologies.

4 FILENET P8 REFERENCE ARCHITECTURE

When speaking of hot backups for FileNet P8, it is helpful to put this in the context of an enterprise reference architecture (ERA). An ERA is a blueprint that defines a set of building blocks and shows how they fit together, providing a common vocabulary and conceptual framework for information technology environments. A FileNet P8 ERA includes a number of core services, namely, image, presentation, and output services that are provided by Workplace, WorkplaceXT, and Application Engine; enterprise content management and business process management provided by Filenet Content Engine, Process Engine, and their respective databases; and storage services provided by the intelligent disk subsystem.

Figure 1 illustrates this reference model.

Figure 1) Reference architecture for IBM FileNet P8 consistency group on NetApp storage solutions.

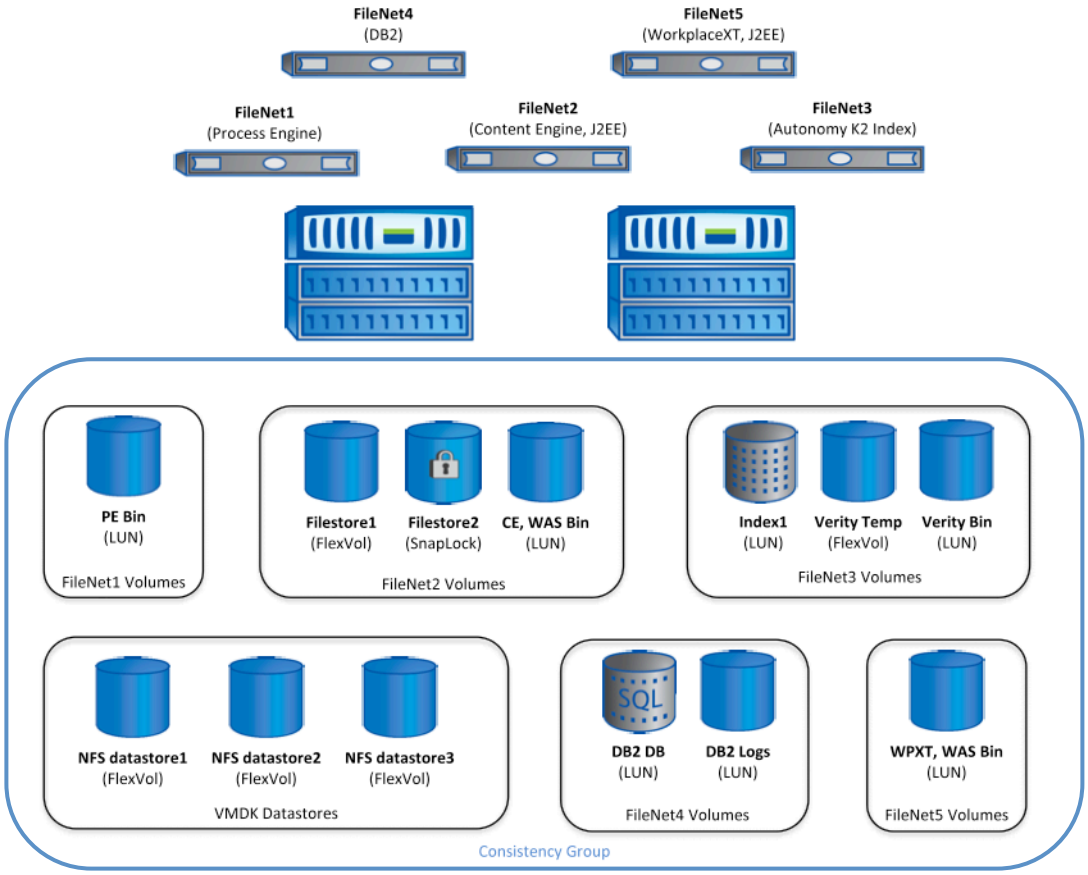

Your production FileNet P8 deployment might have additional applications and servers. For example, you might have FileNet Process Simulator, FileNet eForms, FileNet Rendition Engine, IBM Enterprise Records, or even third-party application integration. In each case you might be able to extend the consistency group to those servers and applications. The details will obviously vary by application and how the datasets are stored. For example, FileNet Rendition Engine employs a database to store persistent data. Because this database can be captured in the same workflow sequence as the core Content Engine and Process Engine databases, the consistency group is extended to the Rendition Engine. Note that the Rendition Engine was outside the scope of testing for this paper. As with any IT infrastructure project, you are advised to consult with a qualified systems integrator to match your particular business needs and with the appropriate technology solution.

5 ENVIRONMENT REQUIREMENTS

Before embarking upon a backup and disaster recovery project, it is important to first assess the state of the current environment. The first steps are generally to review the operating procedures and IBM FileNet architecture diagrams to verify alignment with vendor best practices. Then return to the original solution objectives in designing such an application infrastructure. Review the value of FileNet for the business in order to allocate an appropriate amount of resources into the backup and disaster recovery project. A thorough understanding of the recovery point objective and the recovery time objective will help to justify the additional expense needed for this project to succeed. After completing the business impact analysis and infrastructure assessment, you can begin aligning project resources and adjusting budget forecasts to accommodate the enhanced business value that end-to-end data protection for FileNet will bring to the enterprise.

Because operational staff costs make up the largest percentage of the data protection budget, it is extremely important to automate backups and disaster recovery as much as possible [4,5]. Reducing manual intervention will not only reduce costs, but also help to eliminate operator error. The right solution is the union of people, process, and technology.

For more information about consistency checking, review the P8 Content Consistency Checker information center [5].

6 SNAPCREATOR INTRODUCTION

The NetApp Snap Creator framework addresses the needs and challenges administrators and developers face by providing a centralized and consistent solution for backing up critical information. The NetApp Snap Creator framework seamlessly integrates with existing application environments to reduce costs and enable a faster return on investment, with no additional cost for acquisition.

Figure 2) Snap Creator management GUI.

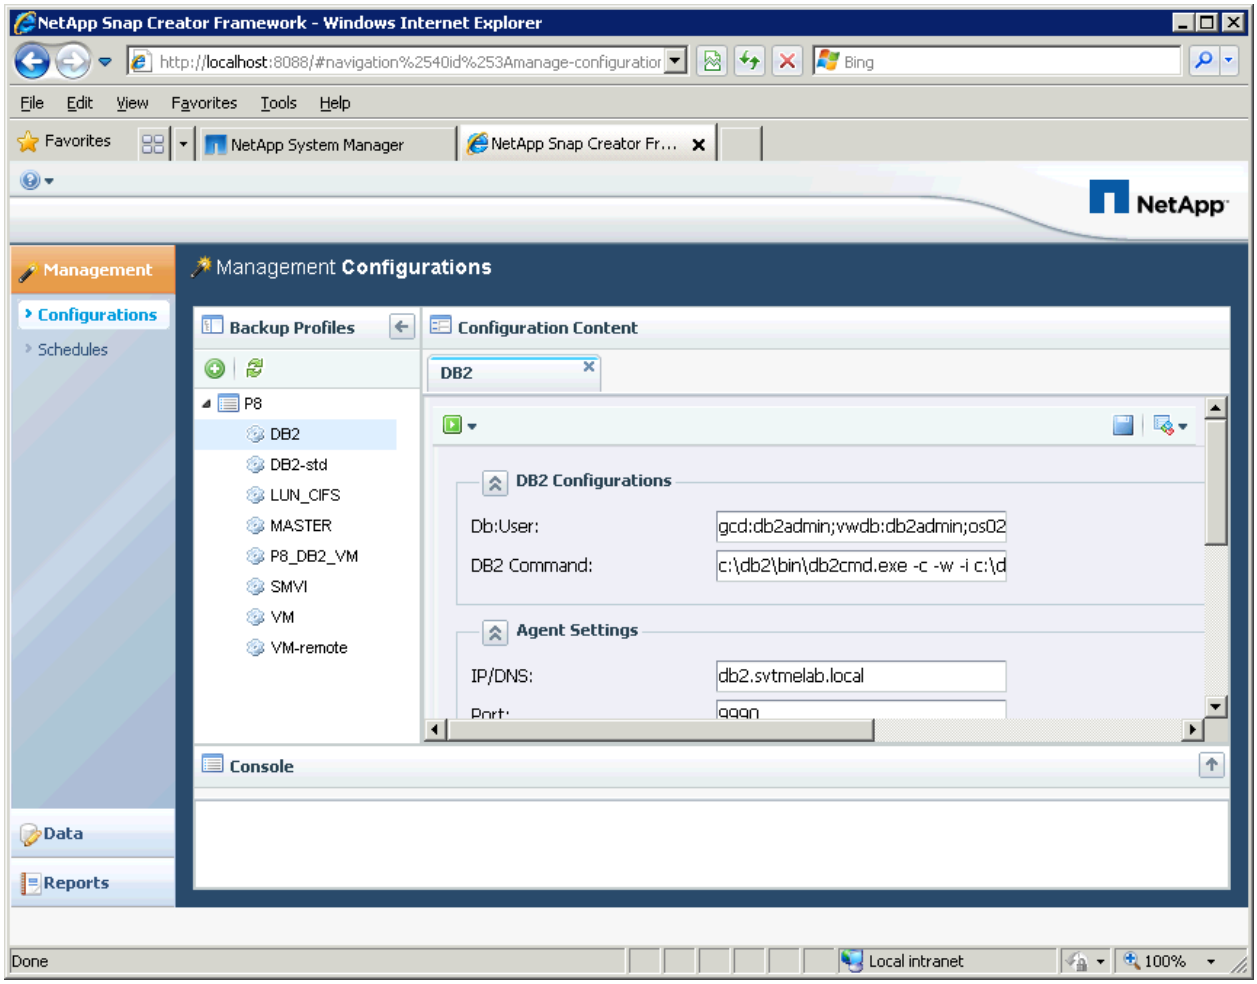

Snap Creator is a modular framework with a convenient command line interface as well as a Web browser graphical user interface based on the NetApp Web framework (NWF) to provide consistency across NetApp Web-enabled applications.

The modular Snap Creator includes plug-ins for Oracle, DB2, Lotus Domino, MySQL, Sybase ASE, SnapManager for Exchange, SnapManager for SQL Server, and VMware® packaged with the framework and supported by NetApp Global Support with a focus on a plug-in architecture and a modular design. Its integration with other NetApp technologies allows you to easily manage complex data relationships across applications, storage controllers, and distributed data centers. You can create and manage one-to-one, fan-in, and fan-out SnapMirror® or SnapVault® relationships. Learn more about Snap Creator in the NetApp Community or from your NetApp sales representative [(6)].

7 BACKUP OVERVIEW

IBM FileNet P8 is a mature and robust ECM platform that has grown over the years since its initial release. It is designed to integrate business processes with numerous IBM and third-party business solutions. This flexibility, unfortunately, creates a problem for backup and application administrators who are entrusted to manage and protect these critical business assets. Furthermore, backup and/or disaster recovery are often afterthoughts and are not considered in the initial design and implementation of the ECM deployment. As the object stores grow over time, it becomes increasingly difficult to create consistent backups while minimizing the impact to the business community. Even more critical, the wrong choice of storage can make it nearly impossible to implement disaster recovery in a way that maintains the internal consistency of the application.

This backup design does not assume any particular arrangement of server roles, network infrastructure, or volume layout. Therefore, many of the requirements for FileNet P8 hot backups find their foundation in NetApp storage efficiency and data protection technologies. A key design consideration is to preserve the IT investments by applying them in a reusable manner to other enterprise applications. Indeed, this design is intended to help you build repeatable factories, not one-off castles. You can get back to business faster by relying on NetApp storage technologies for your most critical business applications.

DISTRIBUTED APPLICATION MANAGEMENT

It's not always clear why we focus so rigidly on consistency in a FileNet P8 environment. Since datasets are potentially spread across numerous servers, storage systems, and applications, it is fundamentally important to manage consistency throughout the backup and restore processes. Otherwise, you run the risk of data inconsistency or unacceptable data loss if something goes wrong and you need to recover from your backup or fail over to your DR site. When managing this application consistency on NetApp storage solutions in a distributed environment, the backup window is minimized, but not instantaneous.

The key to capturing consistency across the complete P8 domain or site lies in the databases and the transaction nature of the Content Engine. The FileNet P8 transaction processing is implemented using the standard J2EE transaction mechanisms through the Java transaction API (JTA). Should the transaction fail (for example, one of the operation steps fails or is not committed to the database), the JTA should handle the rollback of the entire transaction, whether distributed or on a single host. This means that the hot backup will be consistent within itself. But it remains possible that pending transactions made during the short backup window might not have been committed to the database and are therefore not available for recovery.

Ideally the consistency group is mapped to a complete P8 site or to the complete P8 domain in order to protect distributed resources and offer a logical unit of restore. The backup window is further minimized by placing the constituents of the consistency group as close as possible. While there is nothing in Snap Creator that forbids targets from being widely distributed across a corporate WAN, for backup and disaster recovery performance the workflow will complete faster if the latency across the elements of the consistency group is minimized.

Within the context of application-wide, application-consistent hot backups, either you can protect both the operating system (OS) and all FileNet P8 pieces or you can protect only the FileNet P8 components. By including the OS in the FileNet P8 consistency group, you remove the need to manually synchronize patches, hot fixes, and kernel-tuning parameters between production and disaster recovery servers. This

paper discusses using VMware to manage the OS. Using SAN boot and then replicating those volumes is another option. You might define one consistency group to include only the FileNet P8 components, which is then backed up on an hourly basis. Another consistency group composed of the operating system, supporting application binaries, and FileNet P8 can be captured on a nightly or weekly basis. This paper will focus on hot backups and disaster recovery of a virtual environment based on VMware.

It is important to not omit the operating system, operating system patches, tuning parameters, and configuration files in the production consistency group that must be replicated to the disaster recovery data center. When a catastrophe hits, the last thing an IT staffer needs to deal with is inconsistency between the two data centers or to navigate complex, manual operations required to restore service availability. It is not uncommon to have inconsistency between production and disaster recovery servers if patches, hot fixes, or kernel-tuning parameters are manually applied to both sites. A recent ESG Lab paper has discussed the business and technical benefits of disaster recovery synchronization across all levels of mission-critical applications through intelligent automation and virtualization [(7)].

Before continuing with the hot backup architecture, it is advisable to first complete a health check to verify your FileNet P8 and storage environments are in good running order and are following best practices. Enterprise architecture diagrams should be checked for monitoring guidelines and infrastructure readiness. Next consult with the business constituents to begin mapping business requirements to technology capabilities. This exercise will aid in defining consistency groups, setting retention levels, and establishing service-level objectives. See the IBM FileNet P8 Information Center page on enterprise-wide backup for a more complete discussion about which FileNet P8 components must be included in the backup consistency group (8).

The goal of this phase is to outline the essential foundation requirements for your distributed enterprise application backups. The architecture must provide enterprise-level performance, yet remain flexible enough to accommodate new applications and allow extension to other applications. Enterprise application backup solutions based on NetApp storage technologies will help you balance cost with features and flexibility with simplicity.

8 BACKUP SEQUENCE

As illustrated in Figure 4, where the databases were not quiesced before initiating the Snapshot copies, the correct sequence of activities is necessary to create the consistency group backup. Figure 3 provides an overview of the correct sequence of actions.

Figure 3) FileNet B8 backup sequence.

The details of your infrastructure will determine the amount of time for the entire sequence to complete. In our tests, the time required for the NetApp functions in the backup process was less than a minute. The inclusion of the virtual machines or other supporting applications will extend the backup window. Since you are probably applying patches and application upgrades during nightly or weekly maintenance windows, you might choose to back up the VMDKs at the same time. This particular arrangement minimizes end-user impact during the application backups and provides outstanding restore granularity.

The sequence of activities and proper state management is important to provide a consistent FileNet P8 backup. For example, if the database is not suspended during the brief backup window, the environment is vulnerable to lost data.

Figure 4) Consistency check of an incorrect backup sequence.

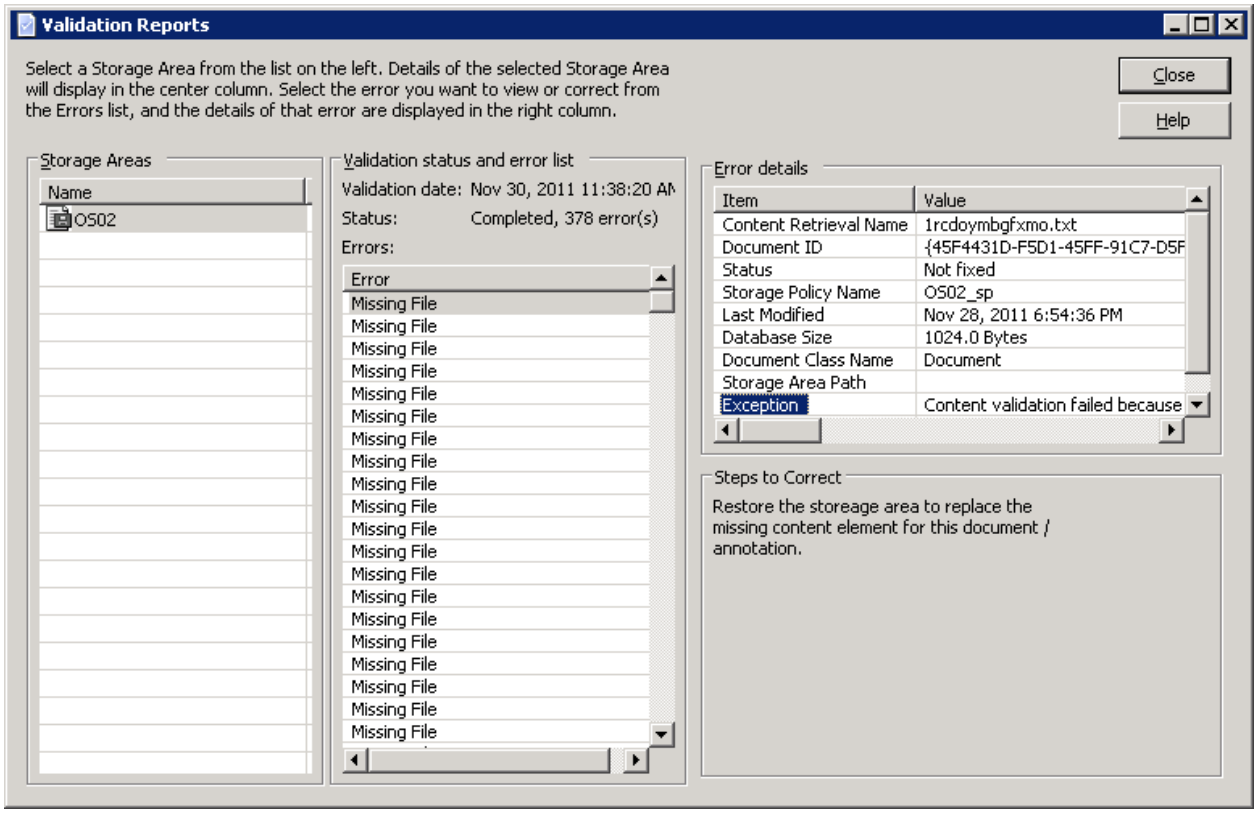

The screenshot in Figure 4 was captured after running the consistency checker following a backup where the database was not quiesced to freeze changes. The inconsistency resulted from changes being applied to the database during the backup window. The metadata in the database was not consistent with the content on disk due to the small time delta between the backup of the two. Snap Creator will manage the overall state changes that you have defined, but it is unable to guarantee those transitions will occur with a particular timing. Therefore it is imperative to follow recommended sequence shown in Figure 3.

8.1 SNAP CREATOR INSTALL

Begin by installing Snap Creator server on a non-P8 server. The installation is driven with a user-friendly wizard. Then install the Snap Creator agent on the vSphere® server, the DB2 server, and any servers with LUNs, such as the Verity server. The Snap Creator agent does not need to be installed on any servers with content in only VMDKs or CIFS network shares.

Following the install and configuration, you can verify the client is listening by telnetting to the Snap Creator pot or by selecting the test agent selection button in the Snap Creator GUI.

```
mbp:~ nwalker$ telnet p8.svtmelab.local 9999
Trying 10.10.46.57...
Connected to p8.
Escape character is '^]'.
```

```
.
HTTP/1.1 400 Bad Request
Date: Sat, 17 Sep 2011 18:31:42 GMT
Server: libwww-perl-daemon/5.827
Content-Type: text/html
Content-Length: 57

<title>400 Bad Request</title>
<h1>400 Bad Request</h1>
```

8.2 WRAPPER SCRIPT

Next create a wrapper script in the language of your choice. This script serves as a simple method to manage the state transition within the FileNet P8 backups and exercise granular control over the process. It should follow the workflow as described in the backup sequence section of this document. The verbose and debug switches are useful for troubleshooting purposes, but optional.

```
Snapcreator.exe --profile P8 --config DB2 --action quiesce --verbose --debug
sc \\verity stop "Verity K2 6.2.1 Administration Server"
Snapcreator.exe --profile P8 --config VM --action snap --policy hourly --verbose --
debug

:while1
sc \\verity query "Verity K2 6.2.1 Administration Server" |findstr STOPPED
    if [%errorlevel%]==[1] (
        echo still running!
        sleep 1
        goto :while1
    )

Snapcreator.exe --profile P8 --config LUN_CIFS --action snap --policy hourly --verbose
--debug
sc \\verity start "Verity K2 6.2.1 Administration Server"
Snapcreator.exe --profile P8 --config DB2 --action unquiesce --verbose --debug
```

In the preceding example, a Windows batch script calls the Snap Creator binaries directly, instead of using the scheduling capabilities. Used in this manner, the Snap Creator server service does not need to be running. Note the sleep.exe binary is not installed by default in Windows. It is available in the Microsoft Windows Resource Kit.

If you run Snap Creator twice within same shell environment, the parameters of the second instance will overwrite those of the first. This is because all of the variables and parameters in the Snap Creator config file are exported to ENV. For this reason it is easier to use a wrapper shell script to manage granular control over the sequence of P8 states and backup activities. This shell script calls Snap Creator in a sequential manner to transition from one backup state to the next. You might want to augment this script with controls for enhanced error handling and advanced service detection.

General settings should be consistent across Snap Creator config files, such as `SCNAME`, `NTAP_SNAPSHOT_RETENTIONS`, `NTAP_PWD_PROTECTION`, and `NTAP_SNAPSHOT_RETENTION_AGE`.

8.3 DB2

Place the databases and logs on a LUN, preferably not a VMDK. By configuring DB2 to place the database and log files on a LUN, you have the flexibility of including them in a consistency group with all the P8 components and another consistency group with everything but the virtual machines. Verify the arrangement of the databases and data integrity before continuing with the implementation.

```
C:\DB2\BIN>db2 create db os02 automatic storage yes on y: pagesize 32768
DB20000I  The CREATE DATABASE command completed successfully.
C:\DB2\BIN>db2 connect to os02

   Database Connection Information

 Database server        = DB2/NT 9.7.3
 SQL authorization ID   = DB2ADMIN
 Local database alias   = OS02


C:\DB2\BIN>db2 update db cfg using NEWLOGPATH Z:/os02_log
DB20000I  The UPDATE DATABASE CONFIGURATION command completed successfully.
SQL1363W  One or more of the parameters submitted for immediate modification
were not changed dynamically. For these configuration parameters, all
applications must disconnect from this database before the changes become
effective.
```

Next, create a DB2 configuration using the Snap Creator server GUI. This configuration type will be plug-in, as it supports the quiesce and unquiesce actions. Verify the Snap Creator agent is listening on the correct port on the DB2 server. Pay special attention to the DB2-specific settings within the DB2 configuration file.

```
DB2_DATABASES=gcd:db2admin;vwdb:db2admin;os02:db2admin

DB2_CMD=c:\db2\bin\db2cmd.exe -c -w -i c:\db2\bin\db2.exe
```

The account used to run the Snap Creator agent needs to have the security context that allows it to write suspend and write resume the DB2 databases for backup operations. For restore operations the account must have DB2 recovery operation permissions. It is recommended that you not use a local administrator or root user for this purpose. You can verify the account permissions from the command line. Run this under the context of the Snap Creator agent service account.

```
C:\db2\bin>db2 connect to os02

   Database Connection Information

 Database server        = DB2/NT 9.7.3
 SQL authorization ID   = DB2ADMIN
 Local database alias   = OS02

C:\db2\bin>db2 set write suspend for database
DB20000I  The SET WRITE command completed successfully.

C:\db2\bin>db2 set write resume for database
DB20000I  The SET WRITE command completed successfully.

C:\db2\bin>db2 connect reset
DB20000I  The SQL command completed successfully.
```

The NetApp deployment and best practices guide for IBM DB2 is found in TR-3272 [9,10]. Consult the guidelines in this document before continuing with your hot backup planning.

8.4 VMWARE

When preparing to implement the P8 hot backups in your production site, you are advised to first validate the virtual machine backup and restore process, independent of FileNet P8. This could be done by first shutting down all relevant ECM services or by conducting a test according to your site operating procedures on lab or development servers.

Create a VMware configuration that will use the NetApp Virtual Infrastructure Backup Engine (VIBE) plug-in to communicate with the vSphere server. Verify the Snap Creator agent is listening on the correct port of the vSphere server. Verify the VIBE_VSPHERE_NAMES is correctly defined in the configuration file. The configuration type should be standard, as it supports the snap action.

An alternative is to use SnapManager for Virtual Infrastructure (SMVI) to capture the virtual machines in the FileNet P8 consistency group. This might be chosen in scenarios where SMVI is already the method of choice for the backup and restore functions. Be aware that when the Snap Creator VIBE plug-in is used to back up the virtual machines, the SMVI metadata is not updated to reflect the backup. Therefore the `smvi backup list` command would not list any Snap Creator VIBE backups. In either case all Snapshot copies of the volume are available from the NetApp storage controller, through either the command line or a GUI. Both SMVI and VIBE were verified in our lab to provide virtual machine consistency in the P8 hot backup sequence. The following example shows one way that you could use SMVI for the wrapper script.

```
"C:\Program Files\NetApp\Virtual Storage Console\smvi\cli\bin\smvi.bat" backup create
-id P8_datastore  -verbose
exit /b 0
```

Take note that in a virtualized environment there are special considerations and best practices for file system alignment. Before continuing with the hot backup planning, make sure you have followed the virtual file system alignment guidelines described in TR-3747 [(10)]. A review of NetApp best practices for VMware vSphere is found in TR-3749 [(10)]. Finally, advanced deployment scenarios, beyond the scope of this document, are discussed in TR-3933 [(12)].

8.5 FLEXVOL VOLUMES AND LUNS

Create another Snap Creator configuration for the LUNs and volumes where the databases, database logs, Verity indexes, Verity temp directories, file stores, and fixed file stores are located. As noted earlier, there is no constraint for any of the FileNet P8 data components to be located on a single NetApp storage controller. Indeed, for scaling purposes these SAN and NAS volumes are often spread across a number of storage controllers for scaling and resiliency.

You can configure a single config file to snap both NAS and SAN targets. Define the NAS targets with the VOLUMES parameter in your Snap Creator configuration file.

```
VOLUMES=controller1:p8_vol1,p8_vol2;controller2:p8_vol3
```

You should use SnapDrive® for Windows or SnapDrive for UNIX® to define the SAN targets. SnapDrive will flush cached writes to disk to preserve consistency. The Snap Creator account will need to have permission to execute SnapDrive on the remote servers. Depending on the desired sequence, you can define them in one of the CMD statements. Refer to the Snap Creator documentation for further details on these parameters.

```
POST_APP_UNQUIESCE_CMD01="c:\Program Files\NetApp\SnapDrive\SDCLI.exe" snap create -m
db2.svtmelab.local -s %SNAME-%SNAP_TYPE-%SNAP_TIME -D y -D z
POST_APP_UNQUIESCE_CMD02="c:\Program Files\NetApp\SnapDrive\SDCLI.exe" snap create -m
verity.svtmelab.local -s %SNAME-%SNAP_TYPE-%SNAP_TIME -D I –D j
```

The Snapshot treatment is identical for both FlexVol® volumes and immutable SnapLock® volumes. So if you have a fixed storage area on a SnapLock volume, there is no special procedure required to create the Snapshot copies within the consistency group. For a detailed introduction to SnapLock, see TR-3263 [(13)]. You can learn more about how FileNet uses SnapLock for fixed content in the IBM FileNet P8 information center.

9 RESTORE SEQUENCE

Just as the hot backups created synchronized Snapshot copies of all the P8 data repositories, the restores also need to be coordinated according to server roles and which particular Snapshot copies are restored. Because the restore is returning all servers and data stores to an earlier moment in time, the process will take the applications offline until revert operations are complete. There is a designated sequence to the restore process. Likewise IBM has published a detailed integrity process that should be followed after the restore is complete.

Begin by shutting down all virtual machines and then restore all volumes corresponding to the target consistency group.

The restore Snapshot actions can be managed from the command line or a GUI, such as NetApp System Manager or SnapDrive. Note that SnapDrive is a host-based application and cannot be used if the server is offline. Therefore you would only use SnapDrive in a use case when you have chosen to restore only the application data, not the complete environment including virtual machines and P8 data.

Figure 5) NetApp System Manager showing FileNet P8 data components.

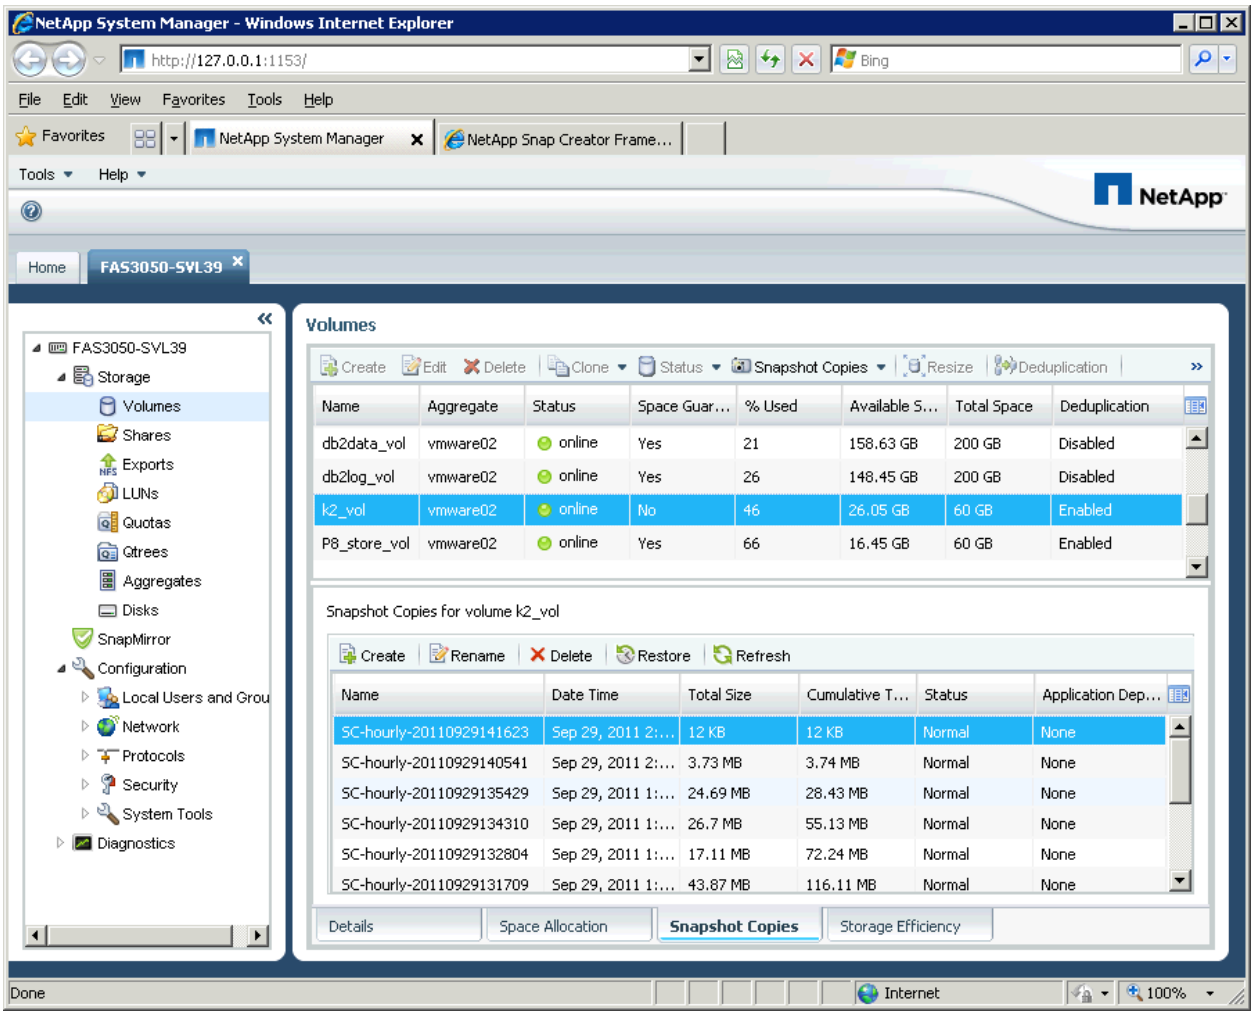

Figure 5 shows all the Snapshot copies for the k2_vol volume that were created as a result of the wrapper script executing while the P8 environment was placed under a simulated load. Each Snapshot copy belongs to a consistency group. Therefore it is important to restore the corresponding Snapshot copy for that particular consistency group.

Figure 6) SnapDrive for Windows showing Verity LUNs.

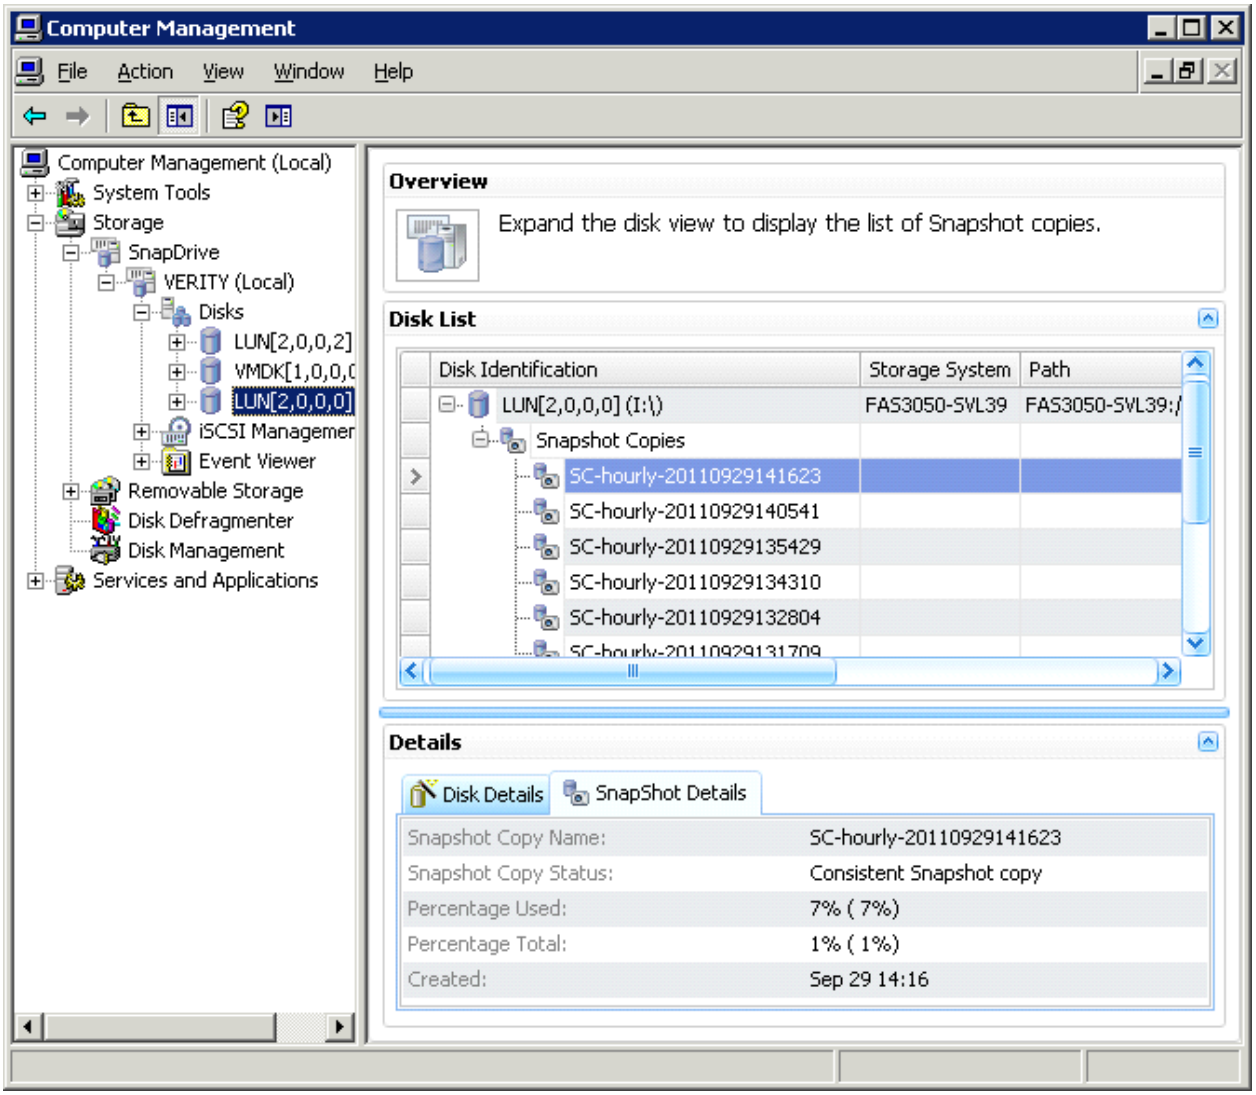

Both SnapDrive for Windows and SnapDrive for UNIX can be used to restore a Snapshot copy of a LUN. SnapDrive can help you to manage host operations of your NetApp SAN from the server. In addition to restoring a Snapshot copy, from the Windows host you can resize storage, create or destroy a LUN, initiate storage reclaimer, and more. It is tightly integrated with the NTFS file system and provides a layer of abstraction between application data and physical storage associated with that data.

This console block shows how to list and restore a Snapshot copy from the command line.

```
FAS3050-SVL39> snap list k2_vol
Volume os01_vol
working...

%/used       %/total  date          name
----------  ----------  ------------  --------
  9% ( 9%)    1% ( 1%)  Sep 29 14:16  SC-hourly-20110929141623
  9% ( 0%)    1% ( 0%)  Sep 29 14:05  SC-hourly-20110929140541
  9% ( 0%)    1% ( 0%)  Sep 29 13:54  SC-hourly-20110929135429
 10% ( 0%)    1% ( 0%)  Sep 29 13:43  SC-hourly-20110929134310
 10% ( 0%)    1% ( 0%)  Sep 29 13:28  SC-hourly-20110929132804
 10% ( 1%)    1% ( 0%)  Sep 29 13:17  SC-hourly-20110929131709

FAS3050-SVL39> snap restore -s SC-hourly-20110929141623 k2_vol

WARNING! This will revert the volume to a previous snapshot.
All modifications to the volume after the snapshot will be
irrevocably lost.

Volume k2_vol will be made restricted briefly before coming back online.

Are you sure you want to do this? y

You have selected volume k2_vol, snapshot SC-hourly-20110929141623

Proceed with revert? y
Volume k2_vol: revert successful.
```

Begin the restore process by shutting down all servers. Then restore the volumes for all the NAS and SAN components in P8. Next, restore the virtual machines, if they were included in the hot backup consistency group. If you used Snap Creator and the VIBE plug-in, then you can use the same config file to direct the virtual machine restore from either the command line or the Snap Creator GUI.

```
snapcreator.exe --profile P8 --config VM --action restore --verbose --debug --policy
hourly --snap_name SC-hourly_20110826195219
```

If you used SMVI to manage the virtual machine backups, then you can restore from the command line or the SMVI GUI.

```
smvi backup restore –id netfs://fas3050-svl39/vol/vmware_p8_vol –backup-name
backup_40dd515659dc623078ee235b7fb6e8423
```

Now you are ready to begin the recovery process of the application. Start the DB2 server. When the server is online and connected to the storage, execute a restart resume command on the DB2 databases.

```
db2 restart database gcd user db2admin using password write resume
db2 restart database os01 user db2admin using password write resume
db2 restart database os02 user db2admin using password write resume
db2 restart database vwdb user db2admin using password write resume
```

From this point in the recovery workflow, the core startup order is Content Engine, Process Engine, WorkplaceXT, and finally Verify. Carefully review the startup, service availability, and data integrity of

each restored component according to section 10 of the Disaster Recovery and Backup Solutions for IBM FileNet P8 (14,7).

10 TESTING METHODOLOGY

In order to confirm the hot backup and restore procedures and data protection, we created a FileNet P8 landscape in our partner compatibility lab. Figure 7 shows the relevant components involved in the hot backup testing.

Figure 7) Proof of concept reference architecture.

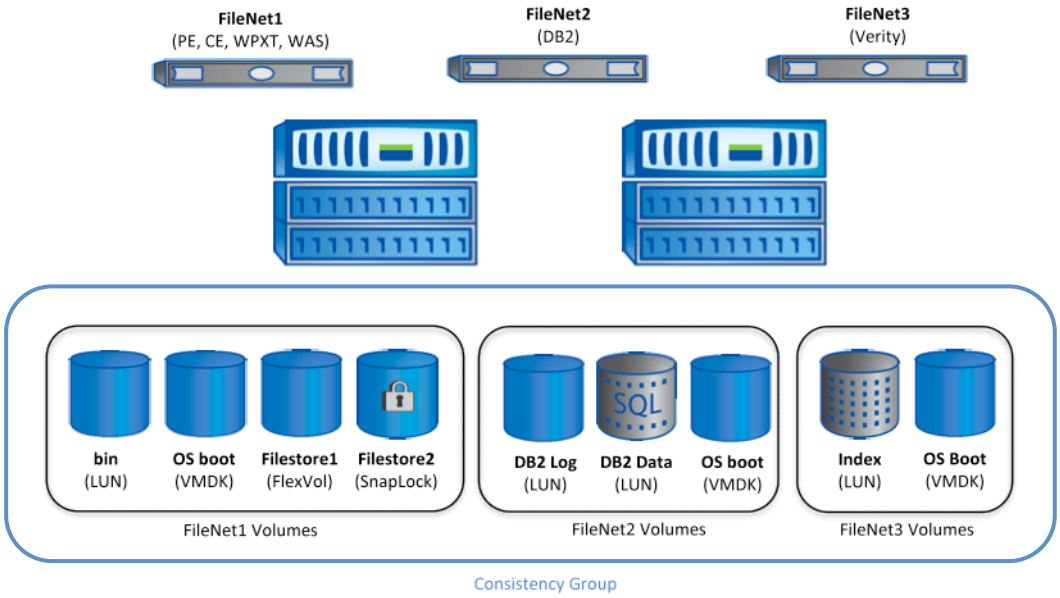

Not shown are the directory servers, the vSphere server, the Snap Creator server, and other common infrastructure elements. All servers were virtualized in VMware. In the case of the Process Engine and Content Engine server, the application binaries were installed on a VMDK. The Verity server, by contrast, had the indexes and index binaries installed on iSCSI LUNs. The DB2 server was configured to hold the databases and database logs on separate RDM LUNs. This mix of disk configurations demonstrates the versatility of NetApp storage technologies in managing complex, distributed applications.

A separate server, also not shown, was configured with an IBM engineering toolset that continually generated files and pushed them into the object store by the Content Engine. This in turn triggered index requests that were processed by the Verity server. The load server had another IBM tool that ran content through the Process Engine in a longevity test. This tool was originally designed to test the robustness of FileNet advanced case management. The J2EE applications never needed to be taken offline or paused during the backups. However, the Verity indexing functions required the corresponding services to be stopped and restarted during the hot backup window to provide data integrity with the indexes. The service stop and restart steps were recommended by IBM FileNet engineering. During the load and timed hot backups, the FileNet P8 servers were continuously close to 100% CPU utilization and had higher than average disk I/O.

When all the test backups were complete, the environment was restored at each hot backup consistency point, starting from the newest and moving toward the oldest consistency group. The application integrity was confirmed using the tools and methodology discussed in the IBM Redbook on FileNet P8 backup and disaster recovery.

In summary, all of the application restores resulted in a recovery point where all of the P8 components were determined to be completely consistent with one another. Each backup in the test suites was found to be consistent, with no data loss.

10.1 BACKUP TEST SUITE 1

The first backup test suite contained all servers, with the exception of the load generator, in the consistency group. The inclusion of the virtual machine images increased the backup window since the virtual machine phase included a VMware snapshot of the servers in the VMware datastore before continuing with the NetApp Snapshot copy of the storage volume.

Table 1 summarizes the results of each backup. The total backup window column was the time between the first database write suspend command and the last write resume command. This time period represents the complete backup window from the client perspective.

The VM Snapshot wait time column was the time used by vSphere to create the virtual machine Snapshot copies in the VMware datastore and to create the NetApp Snapshot copies of the storage volume. The following debug command output shows exactly how Snap Creator effectively manages these operations. First a signal is sent to the vSphere server to create a Snapshot copy of the virtual machines. Next a ZAPI request is sent to the NetApp storage controller to create a Snapshot copy of the volume where the VMware datastore is located. Finally, the virtual machine Snapshot copy is removed. This approach creates a consistent virtual machine Snapshot copy within the P8 consistency group. Although our lab environment had all virtual machine images placed in a single VMware datastore, there is no inherent requirement for this. The Snap Creator profile can be directed to manage multiple virtual machine datastores on multiple storage controllers.

An extract of the debug output for the virtual machine Snapshot sequence is shown in the following console text.

```
[Fri Aug 26 19:39:25 2011] DEBUG: Creating snapshot for VM P8 ...
[Fri Aug 26 19:39:25 2011] DEBUG: Checking power state of VM P8 ...
[Fri Aug 26 19:39:25 2011] DEBUG: Checking snapshot capability of VM P8 ...
[Fri Aug 26 19:39:25 2011] DEBUG: Removing leftover snapshots for VM P8 ...
[Fri Aug 26 19:39:27 2011] DEBUG: Creating snapshot of VM P8 (attempt #1) ...
[Fri Aug 26 19:39:45 2011] DEBUG: Creation of snapshot for VM P8 successful.

(output truncated)

[Fri Aug 26 19:41:35 2011] INFO: Creating NetApp Snapshot for vmware_p8_vol on
FAS3050-SVL39.svtmelab.local
[Fri Aug 26 19:41:35 2011] DEBUG: ZAPI REQUEST
<snapshot-create>
        <snapshot>SC-hourly_20110826193916</snapshot>
        <volume>vmware_p8_vol</volume>
</snapshot-create>

(output truncated)

[Fri Aug 26 19:41:36 2011] DEBUG: Checking power state of VM P8 ...
[Fri Aug 26 19:41:36 2011] DEBUG: Checking snapshot capability of VM P8 ...
[Fri Aug 26 19:41:36 2011] DEBUG: Removing snapshot for VM P8 ..
```

The value in the effective backup window column in Table 1 is simply the difference between the second and third columns. These values represent the time Snap Creator consumes to process the backup job and manage the database states as well as the time to wait for the database and Verity servers to flush their file system cache in order to create the NetApp Snapshot copies of the SAN file systems.

The DB2 db2admin.indexrequests values were read after restoring the databases, but before starting the Content Engine J2EE services. This demonstrated the high rate of document creation by the load

generator, which increased the pending index requests as the test suite continued to run. The final two right columns, Object Store Elements and Verity Collection Count, were determined after completing the consistency check and allowing Verity to process all pending index requests.

Table 1) Backup test suite 1 results summary.

Backup Number	Total Backup Window (Seconds)	VM Snapshot Wait Time (Seconds)	Effective Backup Window (Seconds)	DB2 db2admin.indexrequests	Object Store Elements	Verity Collection Count
1	94	42	52	22,543	359,139	232,316
2	99	48	51	43,126	379,734	252,899
3	97	46	51	61,089	400,697	273,862
4	94	48	46	82,513	422,121	295,286
5	117	50	67	100,641	443,249	316,614

For brevity, only two consistency check screenshots are included. There were no errors in any tests, except when a step was deliberately skipped or transposed, such as the database quiesce at the beginning of the backup workflow noted in Figure 4.

Figure 8) Backup 1, test suite 1 consistency check summary screenshot.

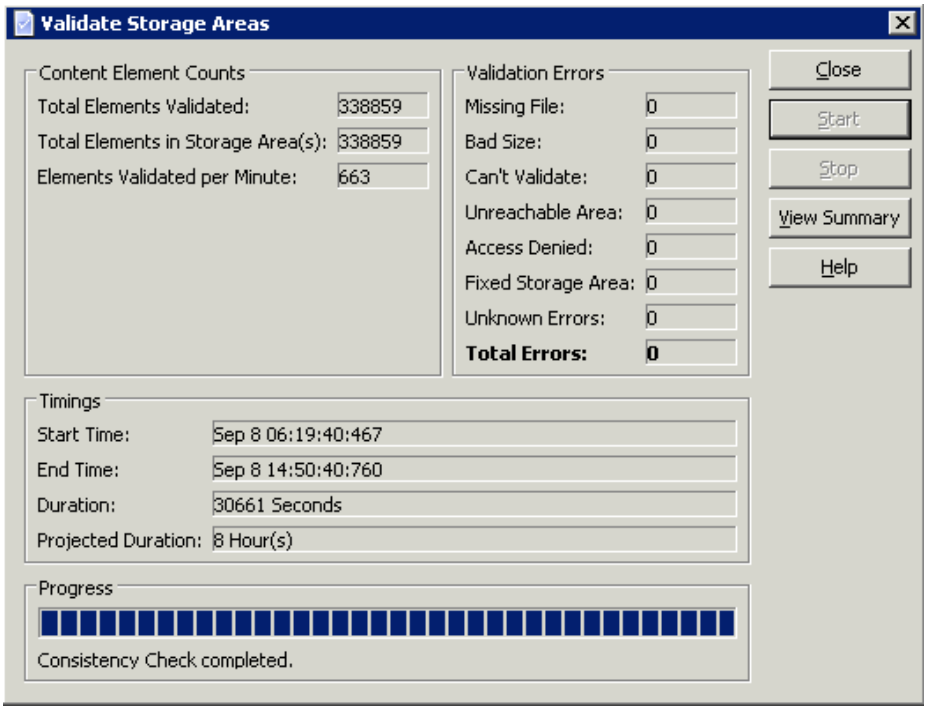

Figure 9) Backup 5, test suite 1 consistency check summary screenshot.

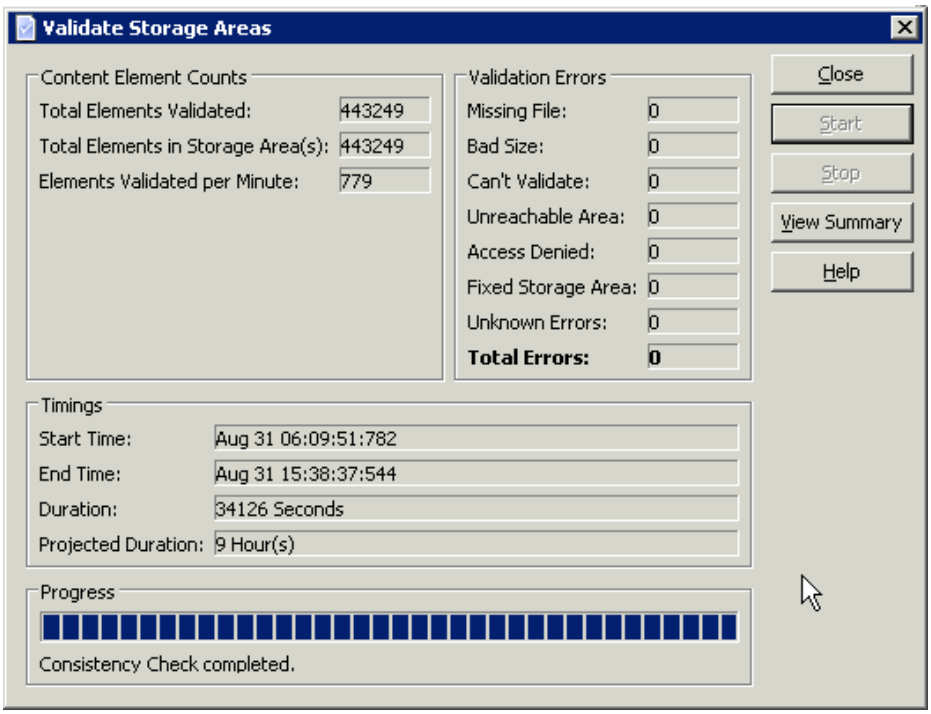

10.2 BACKUP TEST SUITE 2

Table 2, the results table for the second backup test suite, follows the same format as Table 1. The notable exception is the exclusion of the virtual machine backup phase. The goal of this sequence was to verify the ability to create hot backups and then restore the backups without the inclusion of the virtual machines. In all instances the backups were faster. You might choose to execute this type of backup on an hourly basis, during production hours, while the full virtual machine and FileNet P8 backups could be run once a week. There is a measure of flexibility in how you choose to create your consistency groups and when to execute the backups as each backup represents a consistency point of all its constituent parts.

This test suite included measuring the time required to wait for the Verity services to stop. This check was included in all test suites. In the first test suite this time was always shorter than the virtual machine Snapshot phase. The Verity services had stopped by the time the NetApp Snapshot creation was about to begin. In this suite there was no such step, so the process had to pause while the Verity services were being stopped. The physical ESX® hosts, active client load, indexing level, and storage multi-tenancy will all affect the amount of time needed to complete the hot backups. Your results might vary from those of this paper.

Table 2 summarizes the backup details for the partial backups. As previously stated, the virtual machines were excluded. The FileNet P8 file store, DB2 databases, Verity binaries, collections, and temp areas were included.

Table 2) Backup test suite 2 results summary.

Backup Number	Total Backup Window (Seconds)	Verity Wait Time (Seconds)	Effective Backup Window (Seconds)	DB2 db2admin.indexrequests	Object Store Elements	Verity Collection Count
1	73	36	37	28,115	540,659	413,824
2	78	35	43	45,722	562,266	435,431
3	78	35	43	61,122	584,666	457,831
4	92	34	58	83,286	607,830	480,995
5	89	34	55	104,187	630,731	503,896

10.2 PROCESS ENGINE CONSISTENCY

Another tool used to verify the consistency of the backups is the vwverify command line tool for the Process Engine workflow database. Using vwverify, you can scan for inconsistencies in the queues and rosters within the database and make certain repairs.

Following is some vwverify sample output following the restore of a hot backup.

```
P:\FNSW\BIN>vwverify -r 408 -Y p8_bind_user+password
2011/09/29 14:32:37.029  <p8admin> VW/Process  (2476.3272.93 0x9ac.cc8) ... [INFO]
VW (vwverify): VWideCharAuxAttach(): Default system locale is English_United Sta
tes.1252, MB_CUR_MAX=1

2011/09/29 14:32:38.326  <p8admin> VW/Process  (2476.3272.93 0x9ac.cc8) ... [INFO]
VW (vwverify): Tried 0 times: PEDirectory HAS registered with IOR.

Connecting to VWService0:ProcessEngine:FileNet
Isolated Region = 408


Queues:
-------
Delay(0)
        queue work objects:     2378, total:     2378
Tracker(0)
        queue work objects:     1351, total:     3729
Conductor
        queue work objects:        0, total:     3729
CompleteWF
        queue work objects:        3, total:     3732
CE_Operations
        queue work objects:        0, total:     3732
WorkQueue
        queue work objects:      293, total:     4025
InstructionSheetInterpreter(0)
        queue work objects:        0, total:     4025
Inbox(0)
        queue work objects:     2222, total:     6247
WSRequest(0)
        queue work objects:        0, total:     6247

Rosters:
--------
```

```
LoadTest
        roster work objects:       15, total:       15
DefaultRoster
        roster work objects:     6232, total:     6247
Summary:
--------
        Total work objects counted: 6247
sort < vwvfytmp.txt > vwvfyrec.txt
del vwvfytmp.txt
No fixup required
```

A comprehensive FileNet P8 integrity discussion is beyond the scope of this document. Review the IBM information center for additional information.

11 PROFESSIONAL SERVICES

This hot backup methodology has been applied in data centers where hot backups for FileNet P8 were required. The solution blueprint described herein does not imply compatibility or extended support for your particular P8 landscape.

The backup reengineering of a distributed enterprise application can be complicated. A hot backup solution will help you improve business efficiency and data availability while reducing the complexity of your storage architecture. Correctly configuring and integrating hot backups into your existing FileNet P8 environment can be complicated and time consuming.

NetApp and IBM professional services have experience in these types of situations. We apply industry best practices and our experience in hot backup implementations, giving you the full benefit of our wealth of technical knowledge. Our design and implementation consultants will architect, test, and implement a solution that will work for you. Our comprehensive service handoff includes complete verification testing, operating documentation, and knowledge transfer. Go further, faster with NetApp.

12 CONCLUSION

Innovations at NetApp allow application-consistent backups for ECM and archive solutions such as IBM FileNet without downtime even when the application runs across multiple servers and storage systems. This hot backup capability allows backups to be created more frequently without affecting ongoing business operations, improving availability and reducing risk. By automating your complex backup operations on NetApp, you save money and free up precious resources to devote to your core business model. Customers around the world choose us for our “go beyond” approach and broad portfolio of products and services. Our solutions can provide nonstop availability of critical business data and speed product development, so you can deploy new capabilities with confidence and get to revenue faster than ever before. Discover our passion for helping companies around the world go further, faster at www.netapp.com.

13 BIBLIOGRAPHY

1. ESG Briefs: NetApp: Blurring the Lines Between DR and Backup
   www.enterprisestrategygroup.com/2011/06/netapp-blurring-the-lines-between-dr-and-backup/
2. An Introduction to IBM FileNet on NetApp Storage Solutions
   www.netapp.com/us/library/technical-reports/tr-3927.html
3. Enterprise Vault 8 Cascading Snapshot Copies on NetApp Storage Systems
   www.netapp.com/us/library/technical-reports/tr-3709.html
4. 2010 Data Protection Trends
   www.enterprisestrategygroup.com/2010/04/2010-data-protection-trends
5. Administering IBM FileNet P8: Storage Area Consistency
   http://ibm.co/IbxTn9
6. NetApp Community
   http://communities.netapp.com
7. ESG Lab Validation Report: Automated, Affordable Disaster Recovery Solutions
   http://bit.ly/HqP2b8
8. Administering IBM FileNet P8: Enterprise-wide backup
   http://ibm.co/HwbJwN
9. IBM DB2 on NetApp Storage: Deployment and Best Practices
   www.netapp.com/us/library/technical-reports/tr-3272.html
10. NetApp and VMware vSphere Storage Best Practices
    www.netapp.com/us/library/technical-reports/tr-3749.html
11. Best Practices for File System Alignment in Virtual Environments
    www.netapp.com/us/library/technical-reports/tr-3747.html
12. NetApp Snap Creator Framework—Advanced Deployment Scenarios
    www.netapp.com/us/library/technical-reports/tr-3933.html
13. WORM Storage on Magnetic Disks Using SnapLock® Compliance and SnapLock Enterprise
    www.netapp.com/us/library/technical-reports/tr-3263.html
14. Disaster Recovery and Backup Solutions for IBM FileNet P8
    www.redbooks.ibm.com/abstracts/sg247744.html

Go further, faster®

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

© 2012 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, FlexVol, Snap Creator, SnapDrive, SnapLock, SnapManager, SnapMirror, Snapshot, and SnapVault are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. FileNet, IBM, the IBM logo, and ibm.com are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Microsoft, SQL Server, and Windows are registered trademarks of Microsoft Corporation. Oracle and Java are registered trademarks of Oracle Corporation. Enterprise Vault and Symantec are trademarks of Symantec Corporation. UNIX is a registered trademark of The Open Group. ESX, VMware, and vSphere are registered trademarks of VMware, Inc. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. TR-4023-0112

NetApp®
www.netapp.com