



White Paper

SAP HANA Data Protection Made Better

NetApp IT Use Case for SnapCenter Setup and HANA Backup

Rajan Vaidheeswaran, NetApp IT
July 2020 | WP-7331-0820

Abstract

When NetApp IT upgraded their SAP CRM on-premises landscape, a top concern was minimizing downtime, ensuring speedy backups, safeguarding against disasters, and provisioning test and production systems quickly with production data. The team chose NetApp SnapCenter®, a user-friendly data protection solution that encompasses NetApp Snapshot™, NetApp SnapMirror®, and NetApp SnapVault® technologies.

TABLE OF CONTENTS

1	The Challenge	3
2	SAP CRM Inside NetApp	3
2.1	Breadth of SAP CRM System	3
3	Backup and Recovery	4
3.1	Possible Solution Options	4
3.2	Our Selection	5
4	SnapCenter: NetApp Data Protection Tool	5
5	Backup Architecture and Success Stories	6
5.1	Backup Architecture	6
5.2	Results of Disaster Recovery Test	7
5.3	Provision a New System	8
6	About the Authors and Contributors	8

LIST OF TABLES

Table 1)	CRM systems and components	3
Table 2)	Options evaluated.	4

LIST OF FIGURES

Figure 1)	Architecture highlights	6
Figure 2)	SAP system refresh (existing system)	7
Figure 3)	SAP database recovery savings.	7
Figure 4)	SAP system copy (new system)	8

1 The Challenge

All IT organizations understand the downside of downtime. Users expect availability and consistent performance levels 24/7, even in the face of ever-increasing volumes of data. System backups cannot paralyze operations or enterprise applications. As a result, backup windows are shrinking, while the amount of data to be backed up is increasing. It is equally important to provision a test or development system with production data quickly and to perform disaster recovery testing.

Minimizing downtime was a top concern when NetApp IT upgraded our SAP CRM on-premises landscape in late 2018. We chose NetApp SnapCenter®, a data protection solution that encompasses NetApp Snapshot™, NetApp SnapMirror®, and NetApp SnapVault® technologies.

NetApp SnapCenter integrates with SAP HANA databases through plug-ins to deliver data protection and flexible provisioning of SAP environments. It provides critical functions such as scheduling and managing backups, cloning a sub-production environment with production data, and testing disaster recovery without affecting system availability and performance.

2 SAP CRM Inside NetApp

SAP customer relationship management (CRM) software is a critical business application for tracking and managing NetApp customer support cases. It is used 24/7 by approximately 7,000 users ranging from engineers to support teams to external users across geographical locations. It is tightly integrated with other critical business systems such as NetApp® AutoSupport®, Oracle E-Business Suite, the NetApp Support Site, the data warehouse, mobile device management, and Oracle Fusion Middleware.

Our SAP CRM on-premises landscape went through a major upgrade in December 2018 from AIX to Red Hat Enterprise Linux, Oracle to SAP HANA, and CRM from 5.2 to 7.0 EHP4. The current SAP CRM landscape is running on SAP NetWeaver 7.5.

2.1 Breadth of SAP CRM System

The various systems and components that were to be provisioned in our production and sub-production environments are shown in Table 1 below.

Table 1) CRM systems and components.

Type	System	Application	Components	Number of Hosts
Production	Production	CRM	EHP4, NW 7.5, HANA 2.0	4 app + 2 DB
		Process Orchestration (PO)	NW 7.5 & Oracle 12c	3 app + 2 DB
		Gateway	NW & Oracle12C	1 (app & DB)
		Solution Manager (SOLMAN)	NW & Oracle12C	2 app + 1 DB
Sub-production	Patching	CRM	EHP4, NW, HANA 2.0	1 app + 1 DB
		Process Orchestration	NW & Oracle12C	1 app + 1 DB
		Gateway	NW & Oracle12C	1 (app & DB)
		Solution Manager	NW & Oracle12C	1 app + 1 DB
	Development	CRM	EHP4, NW, HANA 2.0	1 app + 1 DB
		Process Orchestration	NW & Oracle12C	1 app + 1 DB
		Gateway	NW & Oracle12C	1 (app & DB)
	QA	CRM	EHP4, NW, HANA 2.0	1 app + 1 DB
		Process Orchestration	NW & Oracle12C	1 app + 1 DB
		Gateway	NW & Oracle12C	1 (app & DB)
	Staging*	CRM	EHP4, NW, HANA 2.0	4 app + 2 DB
		Process Orchestration	NW & Oracle12C	3 app + 2 DB
		Gateway	NW & Oracle12C	1 (app & DB)
		CRM	EHP4, NW, HANA 2.0	1 app + 1 DB

Type	System	Application	Components	Number of Hosts
	Training and Support	Process Orchestration	NW & Oracle12C	1 app + 1 DB
* Will be repurposed as disaster recovery system as necessary				

Service managers coordinated with the business units to gather requests to enhance the CRM system. They shared prioritized lists with the development team and coordinated delivery to predefined release dates. As in the past, we planned two release dates per quarter. In general, such releases do not require any downtime.

The platform team planned proactive, minimally invasive upgrades and patches to the application and Red Hat Enterprise Linux OS to maintain vendor support, mitigate any security vulnerabilities, and leverage any updated features. However, these upgrades required downtime. The platform team planned these upgrades at the start of the fiscal year and published a plan that aligned with the predefined downtime window for each quarter.

3 Backup and Recovery

Like all our critical applications, the objective for SAP CRM backup and recovery was a solution that would not affect performance or availability and addressed these four use cases.

1. Back up production and sub-production systems—both data and non-data volumes
2. Refresh sub-production systems, e.g. patch, development, test, staging, and training, as necessary with production data.
3. Build a new system, as necessary.
4. Demonstrate recoverability of production data in the disaster recovery (DR) systems as part of DR testing.

3.1 Possible Solution Options

Table 3 below describes three solution options that were considered for a 3TB SAP HANA database. Each option had to reduce backup and recovery times as well as be cost-effective, scalable, reliable, and easy to use, learn, and support.

Table 2) Options evaluated.

#	Option	Key Points	Metrics	Considerations
1	Native backup and recovery method as recommended by vendor	It addressed all use cases and incurs no additional licensing cost. It did not require a dedicated host. Yet, there was a performance impact with backup runs.	5 hours for backup and 6 hours for recovery.	<p>Needed to wait for the backup to be available at the target site; it would take approximately 12–16 hours for the backup to be available.</p> <p>Required scripting and expertise to schedule and manage things like backups in Red Hat Enterprise Linux systems.</p>

#	Option	Key Points	Metrics	Considerations
2	Implement database replication for HANA	It addressed use cases 1 and 4 listed above. Yet, might have required another SAP license. It required a dedicated host, which meant more hardware, licensing, and maintenance costs. Synchronization did not affect performance.	0 minutes (synchronous) 15 minutes (asynchronous)	Depending on whether the method was synchronous or asynchronous, data might be available in real time or near real time at the target site.
3	NetApp SnapCenter data protection tool	It addressed all use cases and incurred no additional licensing cost. It did not require a dedicated host and synchronization did not affect performance.	5 minutes for backup and 1 hour for recovery.	Depending on whether the method is synchronous or asynchronous, data might be available in real time or near real time at the target site. Backup could be scheduled every 4 hours at the production site; with completion, it triggers synchronization of backup to the target site, i.e. sub-production.

3.2 Our Selection

We chose NetApp SnapCenter®, a user-friendly data protection solution that encompasses NetApp Snapshot™, NetApp SnapMirror®, and NetApp SnapVault® technologies.

“There’s no team that loves NetApp products and solutions more than NetApp IT. We embrace our products and integrate them with industry and NetApp partner offerings to yield new and innovative business solutions.”

William Miller, SVP and CIO, NetApp

4 SnapCenter: NetApp Data Protection Tool

We chose NetApp SnapCenter because it integrates with SAP HANA databases through plug-ins to deliver efficient data protection and flexible provisioning of SAP environments. It provided critical functions such as scheduling and managing backups, cloning a sub-production environment with production data, and testing disaster recovery without affecting system availability and performance.

Although there are different deployment options, we chose the standalone version of SnapCenter 4.3. It used one system for managing backups in both production and sub-production hosts. Plug-ins are installed in the respective host systems. To ensure disaster recovery and refresh the SAP system, the production SnapCenter system was set to maintain necessary sub-production systems.

We reviewed the following workflows and implemented them in our landscape.

- **SAP System Refresh** to an *existing* target SAP system, e.g. QA or development, with data from a source SAP system. The host name, instance number, and secure identifier (SID) were different for the source and target systems.
- **SAP System Copy** as the setup for a *new* target SAP system, e.g. a sandbox, with data from a source SAP system. The host name, instance number, and SID were different for the source and target systems

- **SAP System Clone** as the identical copy of the source SAP system. SAP system clones were typically used to address logical corruption or to test disaster recovery scenarios. With a system clone operation, the host name, instance number, and SID remain the same as on the source system.
- **SnapCenter Setup** was the web interface used to set up and configure all CRM hosts, storage systems, application users and credentials, backup policies for storage (non-data), and HANA (data and nondata). We devised a backup schedule, replication requirements, and retention policies for each resource, and we set resources accordingly.

As part of the setup, we engaged our storage team to create SnapVault and SnapMirror storage at the target locations, establish replications for the required volumes between the source and the target, and update the SnapVault retention policy for the tags defined in SnapCenter. This updating applied only to policies in which the replication option was selected.

SnapCenter has a built-in dashboard to monitor backups and can mine logs for troubleshooting. Also, the user interface (UI) allowed us to perform system refreshes, system clones, and system copies. By using the reporting features of SnapCenter, we were able to report on backups, restore, and clones by date range and resource group.

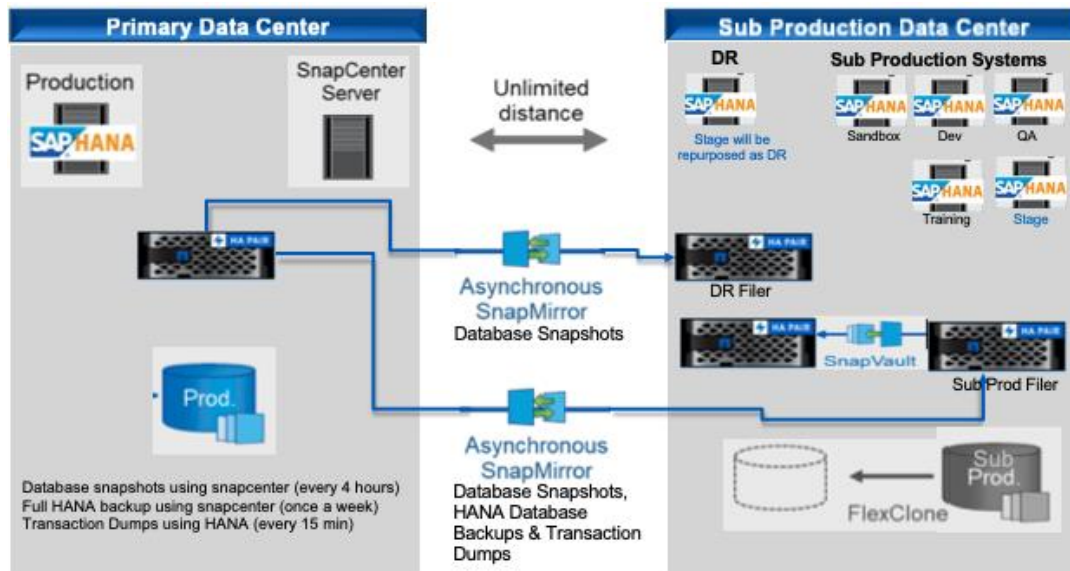
5 Backup Architecture and Success Stories

Our production and sub-production data centers were separated geographically, and we wanted to design a solution that would segregate storage used for disaster recovery and sub-production systems.

5.1 Backup Architecture

As shown in Figure 1, we adopted the architecture and design approach, with some minor modifications, as recommended in NetApp Technical Report ([TR-4646](#)) on SAP HANA Disaster Recovery with Asynchronous Storage Replication.

Figure 1) Architecture highlights.



- Sync Interval to Target Site
- To DR Filer – every 4 hours
 - To Sub Prod Filer
 - DB Snapshot – replicated upon completion of backup
 - Full HANA Backup – every 8 hours
 - Transaction Dumps – every 30 min

We took the following steps with our architecture and design:

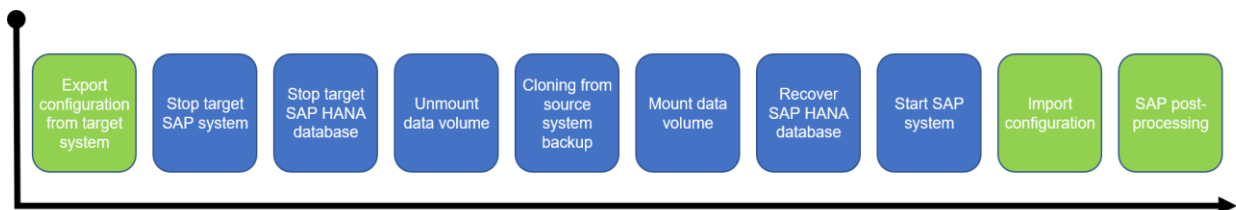
- Established the SnapMirror relationship and synchronization for disaster recovery volumes
- Created SnapVault storage for sub-production volumes. System copy and system refresh flows used this setup
- Scheduled Snapshot copies according to the recommended backup policy
- Scheduled the HANA file system backup weekly. This backup executed a block integrity check of HANA database data blocks.
- Scheduled transaction dumps in HANA and replicated using SnapMirror every 30 minutes to secondary site. During point-in-time recovery, this helps in refreshing sub-production systems to minimize data loss. It would also be helpful during disaster recovery.

NetApp FlexClone® technology allowed us to execute a disaster recovery failover test without influencing or interrupting the ongoing replication to the disaster recovery site. As a result, ran tests without affecting production systems. This [NetApp case study](#) outlines more on our use of FlexClone.

5.2 Results of Disaster Recovery Test

To test recovery steps and demonstrate recoverability of our SAP CRM production data on staging systems, we followed the steps found in NetApp Technical Report ([TR-4667](#)) on Automating SAP System Copies. Specifically, we followed the SAP system refresh workflow shown in Figure 2 and performed point-in-time recovery.

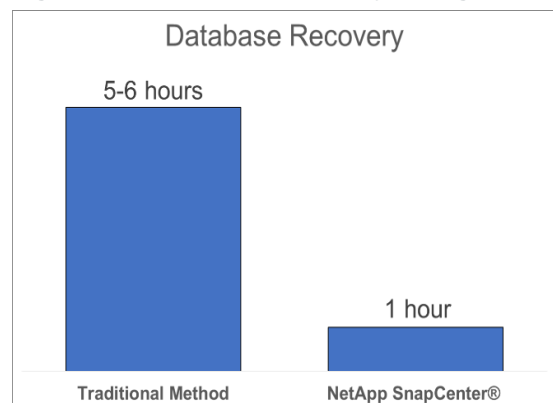
Figure 2) SAP system refresh (existing system)



For the recovery point objective (RPO), the latest production transaction was synchronized with the sub-production storage volumes during the test window. We delivered RPO with 13 minutes of data loss. For the recovery time objective (RTO), we delivered in three (3) hours. A breakdown of the three hours is:

- One hour for database recovery that was comprised of 15 minutes to recover the HANA database and 45 minutes to load data into memory before the application can use the database. As shown in Figure 3, this is a significant time savings compared to the traditional method of a full database recovery that took approximately 5 to 6 hours.
- One hour for post-recovery steps (as required by the application).
- One hour for technical and functional validations.


Figure 3) SAP database recovery savings.



5.3 Provision a New System

In the future, we plan to provision new SAP sub-production systems for evaluating patches and features using NetApp Technical Report ([TR-4667](#)) on Automating SAP System Copies. As shown in Figure 4, we will reference the SAP system refresh for new systems comprising development, QA, staging, and training and support.

Figure 4) SAP system copy (new system).



6 About the Authors and Contributors

Rajan Vaidheeswaran | Author | Systems Architect, Application Platforms Delivery, NetApp

Mr. Vaidheeswaran is focused on identifying, strategizing and delivering critical initiatives within Application Platforms Delivery at NetApp IT

K M S Shreyas | Contributor | Systems Architect, Technology Delivery, NetApp

Mr. Shreyas is focused on delivering solutions in SAP On-Prem Deployments

Nils Bauer | Contributor | SAP Competence Center Manager, Automation and Enterprise Applications, NetApp

Mr. Bauer is responsible for SAP solution prototyping as well as for defining reference architectures, sizing rules and best practices for SAP on NetApp

Ken Lee | Contributor | Tech Lead, Customer-1 Storage Solutions, NetApp

Mr. Lee is focused on storage strategy solutions within NetApp IT

Silar Basha Shaik | Contributor | Information System Engineer, Customer-1 Storage Solutions, NetApp

Mr. Shaik has extensive knowledge of Cluster ONTAP deployment, ONTAP migration, with a focus on SnapCenter

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 2020 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.