Why Your SAP HANA Backup Software Isn't a Disaster Recovery Solution—and What to Do About It

Introduction

Your data-driven organization couldn't survive without you. That's because you help maintain their implementation of SAP HANA. With the in-memory data processing power of SAP HANA, your business can find insights faster than its competition and run processes at a speed that pleases customers.

Of course, when SAP HANA goes down, your business activities will come to a halt. Getting up and running again becomes your top priority. The cost of downtime is no secret to system administrators. The figures will vary widely from one business to the next—not everyone is Facebook, which lost an estimated $90 million in revenue due to a 14-hour blackout on March 13, 2020.[1] But it seems reasonable to accept the 2016 estimate that downtime costs businesses an average of $9,000 per minute.[2]

Nobody disputes the importance of getting back up and running quickly after an outage. What's not clear is how to achieve this goal in a way that helps you meet your business objectives while also respecting your budget.

[1] Brown, Ben. "Facebook's Catastrophic Blackout Could Cost $90 Million in Lost Revenue." CCN.com. October 3, 2020.
[2] Ponemon Institute. "Cost of Data Center Outages." Ponemon Institute Research Report. January 2016.

Why Restoring SAP HANA Takes So Long

Before we go any further, let's establish some assumptions. You're working with SAP HANA on an IBM Power System server that runs SUSE Linux Enterprise. Like most of your peers, your SAP HANA implementation runs on bare hardware with no virtualization. Or, if you want to get technical, you're running on a logical partition (LPAR)—but it's the only form of virtualization on your IBM Power System server.
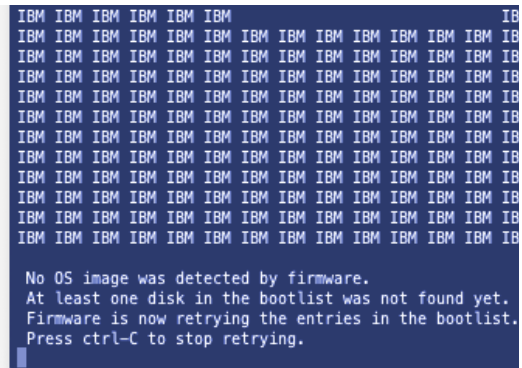
Suppose one of your SAP HANA users comes to you complaining that the system is down. Your first thought is that it's merely a network issue, but you soon realize the network is up and the system is still unresponsive. You attempt to reboot the server but when you try to re-start it, SMS reports "No Operating System Installed" and it doesn't respond.

This server won't be coming back to life. Fortunately, you were using one of the most trusted data back-up solutions on the market, so you're completely confident you can quickly restore your database and bring your business back online.

Your first step, of course, is to restore your SAP HANA server. You remove the dead hard drive from the server, install a new one, and then power up the system. The whole process takes just minutes.

Next, however, you must reinstall SUSE Linux. You and your team begin pulling out DVDs of the operating system and installing them on the server. You then must reapply the 15 security patches you've installed over the past couple of years. Finally, you reinstall a long list of updates for applications and performance tuning.

Eight hours after you began the project, your team brings the server back online, working exactly the way it did before the crash. It's after business hours, and you're all exhausted. But your job still isn't done. You still have to do the data upload.

By the time you finish the upload from the last data backup and are satisfied that everything is working correctly, it's well after midnight. Nobody will even log onto the newly relaunched SAP HANA server for several more hours. Thanks to the crash, your business lost an entire day of productivity. Your employees will spend much of the new workday apologizing to customers for delays and scrambling to catch up on yesterday's work.

There's a reliable, cost-effective way to avoid a scenario like this one, but it's probably not what you think.

What About High-Availability Systems?

In the above scenario, you were doing what you thought you were supposed to do: backing up your data. No serious SAP HANA system administrator would neglect this duty. But there's an even more critical task that few sys admins talk about: backing up the operating system.

If you're trying to restore your SAP HANA operations on a new machine (or on a new hard drive in the old box), you'll of course need to install SUSE Linux on it—and configure it in exactly the way your end users expect. Otherwise, if you install Linux raw, you'll be besieged by user complaints. In a best-case scenario, you'll then have to spend the next several days or weeks gradually installing patches and updates outside of business hours, hoping to disrupt as few users as possible.

At this point, you may be thinking, "Couldn't I avoid this whole problem by using a high-availability (HA) system?" It's a fair question. HA does provide outstanding failover protection. It enables you to mirror your SAP HANA server in another location so that you can resume operations in seconds after a disruption.

But there are still a couple of problems here. First of all, HA systems are outside the IT budget of many organizations. Wall Street stock trading firms can usually justify running HA systems that prevent significant downtime. Most other types of businesses can't.

There's an even larger issue: HA systems replicate everything from the production location to the failover location—even malware.

Malicious cyberattacks remain a significant problem for businesses. In a 2020 Mimecast study, 51 percent of companies had experienced a ransomware attack within the past 12 months. More than 80 percent had experienced downtime from an attack.[3] HA systems offer no protection against these issues. In fact, they can end up replicating corrupt files onto your failover servers.

If your HA system does end up replicating malware onto other servers, you'll need to do more than roll back your data to before the attack. Some forms of malware can attack the operating system. Some can even get into the kernel. Thus, you'll no longer be able to trust any server that was affected by the attack. You'll have to start over on a clean server.

As you can see, the main issue here isn't how effective your backup solution is or whether you can afford an HA system.

The issue is that most SAP HANA backup solutions only back up data.

Yes, data is the lifeblood of your business. Losing it would mean losing countless hours of work and a great deal of your competitive advantage. But in the event of a disaster, the far more time-consuming task is to get your operating system back up and running exactly the way your users expect. That's why it's essential to build a disaster recovery plan that takes into account everything you'll need to do to restore operations fully.

[3] Mimecast. "The State of Email Security 2020." Research report. 2020.

Your Most Important Metric: RTO

Every disaster recovery plan revolves around one important metric: your recovery time objective (RTO). Your RTO should be based on your organization's downtime tolerance, which is simply the amount of time your systems can be down without significant impact to your operations. Your RTO, then, should aim to restore operations slightly before you have to stop serving customers or bring production to a halt.

It used to be that a company's RTO was dictated to some extent by its backup technology. Perhaps you've been in the business long enough to remember the days of tape backups that IT teams had to transport to a secure remote location. Teams would periodically run disaster recovery drills at these remote locations and measure the time it would take to restore computing operations from tape backups.

Nowadays, many companies have shifted to cloud backups. Backing up from the cloud still takes time, of course, and it must be tested periodically. But it can shave hours off a company's recovery time and enable it to set a shorter RTO.
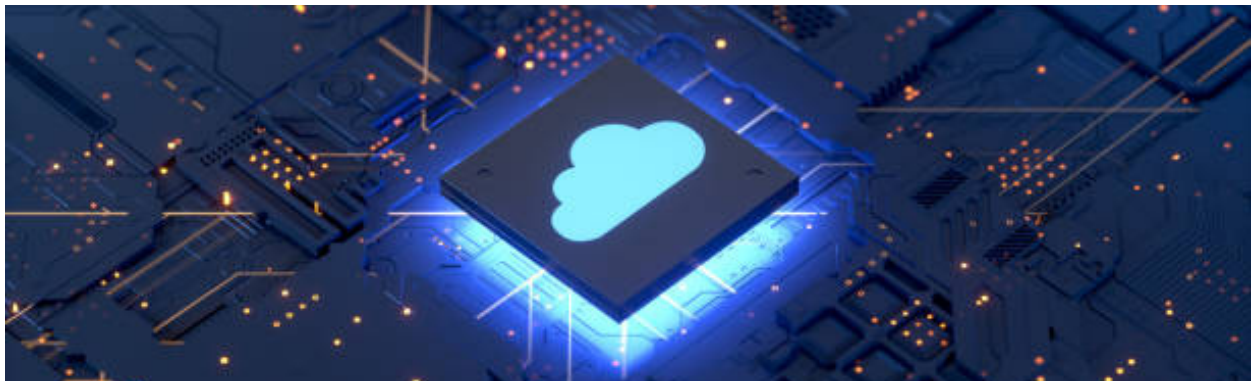
As we mentioned, HA systems enable companies to set and achieve an extremely short RTO. But these systems are expensive—and again, they're still vulnerable to malware attacks.

For most companies that don't require "always-on" systems and want to balance RTO with budget, a bare-metal recovery system is the best option for protecting their SAP HANA implementation.

Introducing SAP HANA Disaster Recovery

When disaster strikes, the last thing you have time to do
is rebuild and configure your operating system to get it to
where it was before the outage. Think of all the hours you've spent fine-tuning your operating
system to yield the best performance. In the aftermath of a dead server, you'll face the
unpleasant choice of either letting that hard work go to waste, or trying to restore it and
causing a significant delay in bringing SAP HANA back online.

A far better approach is to back up your entire system. That way, you'll have no rebuilding work
to do when you install a new server. You'll simply restore one backup and be able to bring SAP
HANA up for your users.

What we're describing is more than SAP HANA backup—it's SAP HANA disaster recovery. That's what Storix SBAdmin provides.

At Storix, our approach is to back up the operating system and data at the file system level.
That way, you'll be able to get SAP HANA up and running within minutes after an outage. If
you're a typical business with an RTO of one hour, you'll most likely find that **Storix SBAdmin
can meet your needs—and eliminate much unnecessary work and stress for your IT team.**

Request Your Proof
of Concept Today

Nearly every IT team is working under serious budgetary constraints—especially nowadays. As you consider the cost of investing in full system backup, compare it to investing in multiple failover servers at a cost of $15,000 to $30,000 per box. By contrast, you could simply purchase one high-end server, protect it with full system backup, and rest assured that your operations will never be interrupted for more than an hour.

*Our team at Storix would appreciate an opportunity to provide you with a **free proof of concept** for full system backup of your SAP HANA implementation. To get started, **call us today at (877) 786-7491 or email sba@storix.com.***

STORIX®
S O F T W A R E

9150 Chesapeake Dr
Suite 290
San Diego, CA 92123
www.storix.com