# Staying operational

Dave Bermingham, Director of Customer Success, SIOS Technology looks at how we can ensure operational continuity for building management systems

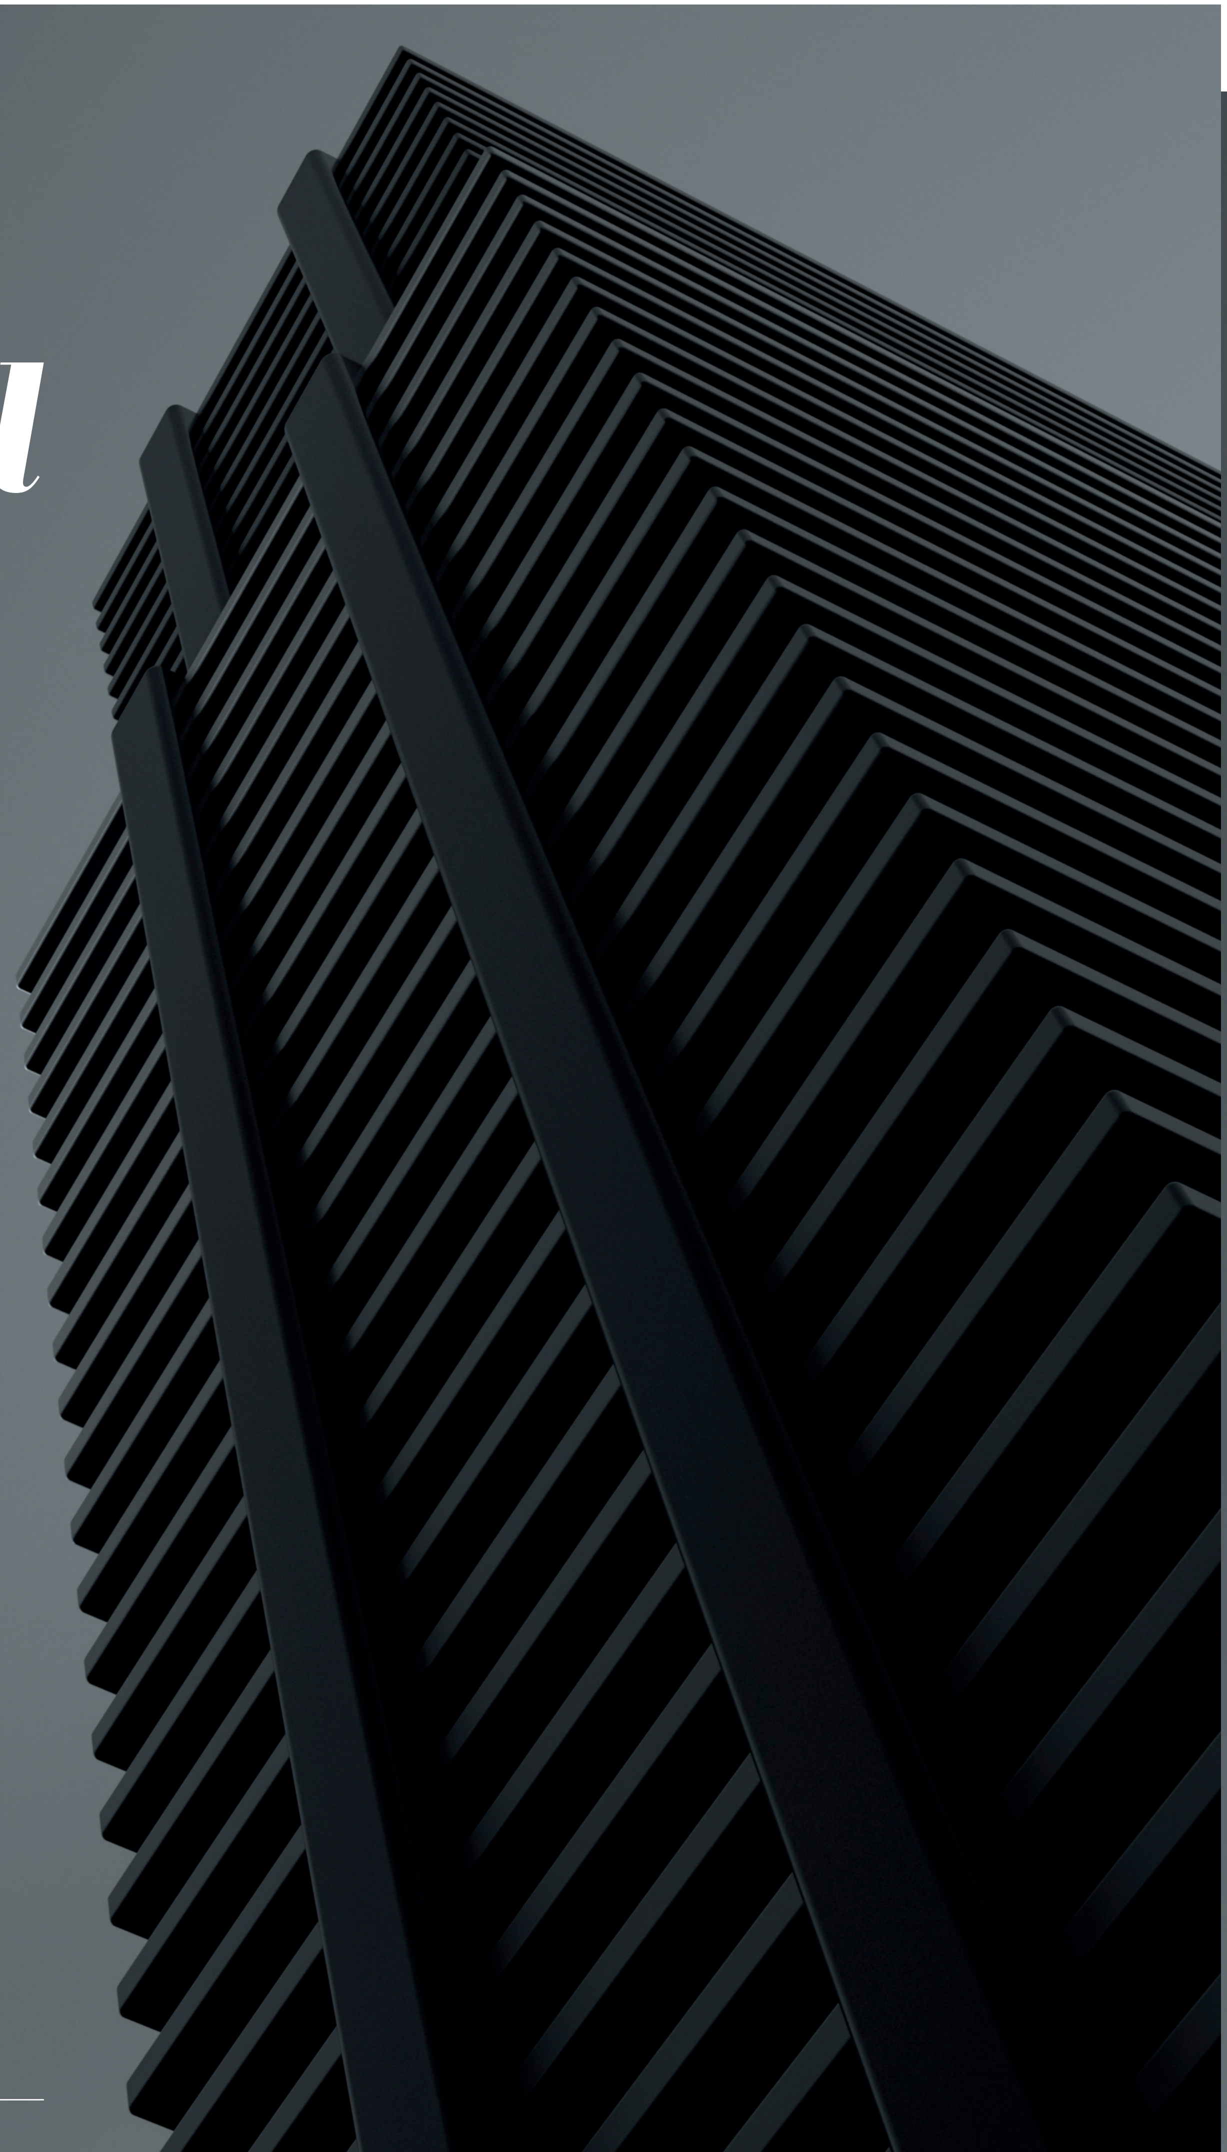

In today's built environment, a problem in the computing infrastructure supporting an integrated building management system (BMS) is more than an inconvenience. If the system supporting your BMS suddenly goes dark – because of a disk crash, a software fault, a power failure or anything else – then the automatic doors it controls may not open. The elevators may stop between floors. The cameras monitoring entrances and hallways may go dark, putting people and property at risk.

It's crucial to ensure that your BMS will operate with minimal interruption – and uninterrupted power supply systems or backup generators are, as they say, necessary, but can often be insufficient. The computing infrastructure running the applications tasked with ensuring the security of the building and the safety of residents should be configured for high availability (HA) – meaning that they are guaranteed to be online and accessible no less than 99.99% of the time.

How does an HA configuration thwart something as catastrophic as a disk crash? Simple: In an HA configuration, there is always another server ready to take over the responsibilities managed by the BMS. Integrated software designed to support HA monitors the active BMS infrastructure and, the instant it detects that the BMS is not performing as expected – and cannot be brought back to full working order immediately – it executes what is known as a failover to the secondary infrastructure. The secondary instance of the BMS then immediately takes over – opening doors, monitoring elevators, watching the hallways and more – as though there were no interruption at all.

## Configuring for high availability

Most BMS have two key components: An application that contains the logic behind the tasks it manages and a database containing a record of operational parameters and settings. A video management system such as Milestone XProtect, which monitors surveillance cameras, might, for example, be running on a Windows-based server and relying on a database running on Microsoft SQL Server or SQL Server Express.

By creating a second instance of this same configuration and integrating the two instances using software designed to ensure HA, a fault in the active system will prompt failover to the secondary system, which, within seconds, will pick up right where the other left off. The secondary system won't be working with a shared instance of the primary system's database; instead, it will be using its own complete copy of that database, ensuring no loss of data as the secondary instance takes over the tasks formerly managed by the other instance.

There are two ways to configure such an infrastructure for HA. Certain editions of Microsoft SQL Server support Availability Groups (AGs), a service that can synchronously replicate one or more SQL Server databases between multiple SQL Server instances. The AG will automatically fail over to a secondary instance of SQL Server if the primary instance becomes unresponsive.

However, the AG approach to HA has its quirks: If you're using the Basic AG functionality that comes with SQL Server Standard Edition, you can only replicate a single SQL Server database to a single secondary. If you have more than one SQL Server database to replicate, you'll either need to configure multiple AGs or you'll need to use the Always On AG functionality that comes with the more expensive SQL Server Enterprise Edition.

If you're using SQL Server Express Edition, you're out of luck, as SQL Server Express does not include the AG feature. Note too that the AG feature doesn't replicate anything other than SQL Server databases, so any other data that you might want to have available on the secondary infrastructure will require the integration of a separate replication mechanism.

> “AN ALTERNATIVE TO THE AG APPROACH TO HA INVOLVES USING A THIRD PARTY SANLESS CLUSTERING SOLUTION.”

An alternative to the AG approach to HA involves using a third party SANless clustering solution. Using a SANless clustering solution, such as SIOS LifeKeeper for Windows, you can create a failover cluster in which all the information in storage is synchronously replicated to storage attached to the secondary infrastructure. The SANless clustering software performs block level replication, so it doesn't distinguish between file types. It simply replicates any updates written to primary storage to the secondary system.

Indeed, your storage system might hold multiple databases supporting multiple BMS components, which makes the SANless clustering approach ideal. Those databases may or ▸

may not be SQL Server databases, but the SANless clustering approach will ensure that all the databases are replicated to the secondary infrastructure and are ready to support all your BMS requirements. In the event of a failure on the primary system, the SANless clustering management software immediately fails over to the integrated secondary infrastructure, ensuring the continuity of your BMS.

> ON-PREMISES LOCATIONS THAT ARE WITHIN A FEW MILES OF ONE ANOTHER OR CLOUD AZS THAT ARE WITHIN THE SAME REGION ARE IDEAL.

One other thing to keep in mind: If your BMS infrastructure is running on Linux rather than Windows, a SANless clustering tool such as SIOS LifeKeeper provides a level of versatility that the AG functionality of SQL Server does not; SIOS LifeKeeper can provide system monitoring and data replication support for Linux-based BMS solutions as well as Windows-based BMS solutions.

## The where matters

Is your BMS running on-premises or on a rack of servers in the cloud? Is it part of a Windows Workgroup or an Windows Active Directory (AD) domain? Some AG-based HA configurations require that the underlying infrastructure – whether on-premises or in the cloud – be configured as a member of an AD domain, which may seem like overkill if your BMS is the only system that requires an AD infrastructure. If your infrastructure is configured as a member of a Workgroup, a tool like SIOS LifeKeeper for Windows will be more attractive because it doesn't require an AD domain.

From the standpoint of HA, whether your BMS runs on-premises or in the cloud, it doesn't really matter. What does matter is that your secondary cluster node should be housed in a location that is geographically distinct from the one in which your primary node is housed. For a BMS running on-premises, that secondary could be running in a remote office or colocation facility. For a cloud-based infrastructure, that secondary could simply be configured to operate in a separate availability zone (AZ) – an option readily available from all cloud providers.

In both scenarios, the physical distance between the nodes should be great enough that anything affecting the location of the primary node is unlikely to affect the secondary location. At the same time, that secondary location should not be so far away that the distance itself would cause a delay in the synchronous replication of data between the infrastructures. On-premises locations that are within a few miles of one another or cloud AZs that are within the same region are ideal.

By configuring your BMS infrastructure for HA, you can ensure that the applications supporting the operations, safety and security of your buildings remains operational, even if something occurs to compromise the infrastructure supporting those applications. An environment configured for HA ensures that your applications and data can quickly fail over to infrastructure that is already integrated and poised to take over in the event of an emergency. ▪