Understanding the Complexity of High Availability for Business-Critical Applications

Lawrence Miller

CONTENTS

IN THIS PAPER

In today's highly competitive, always-on global economy, downtime has become more costly than ever before for modern businesses. In addition to lost productivity and revenue, organizations risk losing customers when their business-critical systems, databases, and applications aren't available to deliver a reliable and superior customer experience.

This tech brief from SIOS explains the complexity of achieving high availability in business-critical applications.

Highlights include:

- High availability basics and design principles
- Clustering and availability group concepts
- Application and database clustering challenges

Minimizing downtime in systems, databases, and applications is the key to maximizing productivity. Modern organizations rely on business-critical systems, databases, and applications—such as enterprise resource planning (ERP), customer relationship management (CRM), e-commerce, financial systems, and supply chain management—to operate efficiently and deliver superior customer experiences. When a system, database, or application fails, high availability protection restores operation to keep the business up and running.

What Is High Availability?

High availability (HA) is an attribute of a system, database, or application that is designed to operate continuously and reliably for extended periods. The goal of HA is to reduce or eliminate unplanned downtime for critical applications by incorporating redundant components and other technologies to address single points of failure in a system, database, or application.

Simply stated, HA ensures that your system, database, or application operates *when* and *as expected*: "when" refers to the percentage of time the system, database, or application must be up and running as expected—meaning that the application operates the way users expect and meets their needs in a timely manner.

IDC MODEL

Service-level agreements (SLAs) for HA help ensure that key components of the IT infrastructure are operational and available during business hours. IDC has created an SLA model for HA that defines five levels with the following uptime requirements:

- **AL4 (Continuous Availability—System Fault Tolerance):** No user interruption and a total maximum of no more than 5 minutes and 15 seconds of planned and unplanned downtime per year (99.999% or "five-nines" availability).
- **AL3 (High Availability—Traditional Clustering):** Minimal user interruption and a total maximum of no more than 52 minutes and 35 seconds of planned and unplanned downtime per year (99.99% or "four-nines" availability).
- **AL2 (Recovery—Data Replication and Backup):** Some user interruption and a total maximum of no more than 8 hours, 45 minutes, and 56 seconds of planned and unplanned downtime per year (99.9% or "three-nines" availability).
- **AL1 (Reliability—Hot Swappable Components):** All service stops and a total of 87 hours, 39 minutes, and 29 seconds of planned and unplanned downtime per year (99% or "two-nines" availability).
- **AL0 (Unprotected Servers):** All service stops, and no uptime SLAs are defined.

Your HA requirements depend on the criticality of the overall system, the application, and numerous other factors, including:

- How critical the applications are to the business
- Whether customers notice an impact
- How often the applications run
- How many users are affected by downtime
- How quickly a database or application must fail over to the redundant system to avoid disruption
- How much data loss is tolerable

Five-nines availability is typically reserved for applications that require continuous "stateful" operation. For business-critical applications four-nines availability is standard. Non-critical systems and applications, you may only require two-nines availability. When determining acceptable downtime, it's important to consider:

- Unplanned downtime (that is, hardware or software failures)
- Planned downtime for routine hardware and software maintenance
- Uptime at the application and database level

Various HA solutions can help businesses achieve their SLA objectives for different systems, databases, and applications. Although continuous availability (AL4) may seem like the most appropriate goal for business-critical deployments, it's important to find the right balance between cost and availability.

HA METRICS

In addition to uptime and availability, *Recovery Time Objectives* (RTOs) and *Recovery Point Objectives* (RPOs) are important metrics used to assess HA (as well as disaster recovery) in a system, database, or application.

RTO is the maximum tolerable duration of any outage. Online transaction processing applications generally have the lowest RTOs, and those that are business-critical often have an RTO of only a few seconds.

RPO is the maximum amount of data loss that can be tolerated when a failure happens. For disaster recovery, a typical RPO for an application and its associated data may be 24 hours. Nightly backups ensure that any changes to data over the past 24 hours can be restored in the event of a disaster. However, for HA applications and data, the RPO is often zero. That is, there should be no data loss under any failure scenarios.

Traditional Clustering

HA clusters are groups of server nodes (and other components) that support business-critical applications that require minimal downtime and continuous availability. Clustering software lets you configure your servers as a cluster so that multiple servers can work together to provide HA and prevent data loss. IT organizations rely on HA clustering to eliminate single points of failure and minimize the risk of downtime and data loss.

A traditional, on-premises HA cluster is a group of two or more server nodes connected to shared storage (typically, a storage area network, or SAN) that are configured with the same operating system, databases, and applications (see **Figure 1**).

One of the nodes is designated as the primary (or active) node and the other(s) are designated as secondary (or standby) nodes. If the primary node fails, clustering allows operation of a system, database, or application to automatically fail over to one or more secondary nodes and continue operating as normal with minimal disruption. Since the secondary node is connected to the same storage, operation continues with zero data loss. The benefits of this cluster architecture are reduced downtime, elimination of data loss, and protected data integrity.

However, there are many scenarios in which shared storage is not wanted. A failure in the shared storage will take all of the clusters off line, making it a single point of failure (SPoF) risk. SAN storage can also be costly and complex to own and manage. Lastly, using shared storage in the cloud can add significant, unnecessary cost and complexity. Some clouds do not offer a shared storage option at all.

As shown in **Figure 2**, SANless or "shared nothing" clusters are the best alternative to shared storage. In these configurations, every cluster node has its own local storage. Efficient host-based, block-level replication is used to synchronize storage on the cluster nodes, keeping

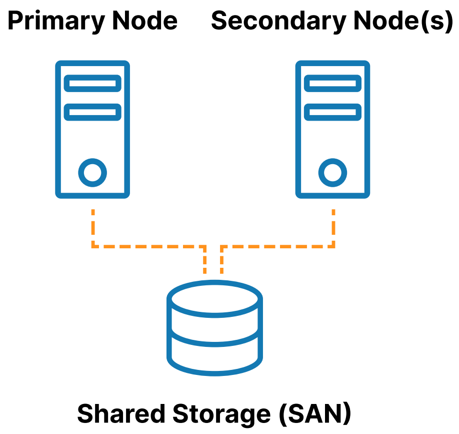

Figure 1: Traditional server clustering with shared storage

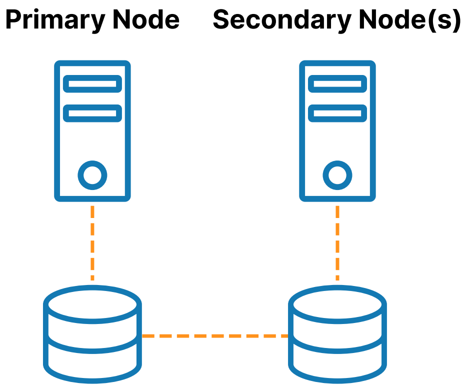

Figure 2: HA clustering with SANless or shared-nothing storage

them identical. In the event of a failover, secondary nodes access an identical copy of the storage used by the primary node. The benefits of this cluster architecture are elimination of a SPoF, elimination of SAN cost and complexity, ease of use and cost savings in the cloud, reduced downtime, and mitigation of data loss.

DESIGN PRINCIPLES

The most advanced HA clusters incorporate the following design principles:

- They automatically and quickly fail over to a redundant system when an active component fails
- They maintain application-specific best practices during and after the failover
- They provide the ability to manually switchover and switch back to enable efficient testing and "rolling" maintenance with minimal planned downtime
- They can automatically detect application-level failures in network, storage, OS, hardware, or application
- They prevent data loss in the event of a system failure
- They failover across geographically separated nodes for disaster recovery

HA Clustering

A variety of clustering software solutions are available for Windows, Linux distributions, and various hypervisors (virtual machine solutions). One group supports only a single operating system, such as the following:

- Windows Server Failover Clustering (WSFC): Provides HA and disaster recovery for hosted applications such as Microsoft SQL Server and Microsoft Exchange
- SUSE Linux Enterprise HA Extension (HAE): Supports clustering of physical and virtual Linux servers with policy-driven clustering and continuous data replication
- Red Hat Pacemaker (Pacemaker): Creates single-site clusters for performance, HA, load balancing, and scalability

None of the solutions listed here can protect SAP running on Oracle Linux operating system for example. Thus, each solution limits your flexibility and deployment options. More advanced HA solutions, such as SIOS Protection Suite for Linux, provide application-aware protection in major Linux distributions, including Oracle Linux, Red Hat, and SUSE.

> **IT organizations rely on HA clustering to eliminate single points of failure and minimize the risk of downtime and data loss.**

In addition, every application, database, and ERP system has its own requirements for configuration, failover orchestration, and ongoing management. To meet these requirements, HAE and Pacemaker typically require a high degree of technical skill, and complicated manual scripting, which introduces the likelihood of human error and unreliable failover.

Some examples of business-critical applications, databases, and ERP systems commonly protected with failover clustering include SAP S/4HANA, SQL Server, and other applications and databases.

SAP S/4HANA

Several Linux vendors offer open source HA extensions for SAP in their "Enterprise for SAP" subscriptions. SAP S/4HANA environments comprise multiple services such as ABAP SAP Central Service (ASCS), Evaluated Receipt Settlement (ERS), and other SAP components, that need to be maintained in the right locations and started up in the right order. In open source clustering products, such as SUSE HAE and Red Hat Pacemaker, manually configuring and managing clusters in these complex environments can be time-consuming and prone to human errors that increase the risk of catastrophic downtime and data loss.

Specific deep expertise in the applications and database is also required to create an application-aware HA solution. In contrast, SIOS Protection Suite for Linux includes application recovery kits for SAP and HANA that ensure failovers maintain application best practices.

SAP also offers HANA System Replication, a feature that comes with the HANA software. It provides continuous synchronization of an SAP HANA database to a secondary location in the same data center, at a remote site, or in the cloud. The data is replicated to the secondary site and pre-loaded into memory. When a failure occurs, the secondary site takes over without a database restart, which helps to reduce the RTO. However, failback to the primary node must be manually triggered. HSR needs to be paired with an application-aware clustering software such as SIOS Protection Suite that can detect failures and orchestrate failovers if necessary.

> HAE and Pacemaker typically require a high degree of technical skill, and complicated manual scripting, which introduces the likelihood of human error and unreliable failover.

SQL SERVER

Many companies rely on SQL Server as the back-end database for key applications supporting important business functions. Microsoft WSFC is commonly used to support Always On availability groups (AG) and SQL Server Failover Cluster Instances for SQL Server applications. However, WSFC with AG requires costly SQL Server Enterprise Edition licensing. In addition, With FCI, the entire instance is failed over to the standby node. With AG only the databases in the group are protected. Using SIOS DataKeeper with WSFC allows you to provide advanced HA protection for SQL Server using cost-efficient Standard Edition licensing.

OTHER APPLICATIONS AND DATABASES

SIOS software can be used to protect a wide range of business-critical applications, databases and ERPs, including Oracle, MaxDB, MySQL, PostgreSQL, and DB2.

SIOS software enables clustering and disaster recovery.

> HA is increasingly crucial for modern operations.

PURPOSE-BUILT FOR HA

HA is increasingly crucial for modern operations. But with the myriad platforms available, complexity ramps up significantly. That’s why a application-aware solution makes so much sense. What you need is a trusted partner who has extensive expertise in high availability—a partner like SIOS, which has the technological know-how to ensure that your business stays up and running.

Don’t wait for a disaster to find out if you have enough resiliency. Schedule a personalized demo today at https://us.sios.com to see what SIOS can do for your business.