

This is a repository copy of *High-Availability Clusters: A Taxonomy, Survey, and Future Directions*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/181882/>

Version: Accepted Version

Article:

Somasekaram, Premathas, Calinescu, Radu orcid.org/0000-0002-2678-9260 and Buyya, Rajkumar (2021) High-Availability Clusters: A Taxonomy, Survey, and Future Directions. *Journal of Systems and Software*. 111208. ISSN 0164-1212

Reuse

This article is distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs (CC BY-NC-ND) licence. This licence only allows you to download this work and share it with others as long as you credit the authors, but you can't change the article in any way or use it commercially. More information and the full terms of the licence here: <https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

High-Availability Clusters: A Taxonomy, Survey, and Future Directions

Premathas Somasekaram^a, Radu Calinescu^a, Rajkumar Buyya^b

^aDepartment of Computer Science, University of York, Deramore Lane, York, YO10 5GH, UK

^bCloud Computing and Distributed Systems (CLOUDS) Lab, School of Computing and Information Systems, The University of Melbourne, Australia

Abstract

The delivery of key services in domains ranging from finance and manufacturing to healthcare and transportation is underpinned by a rapidly growing number of mission-critical enterprise applications. Ensuring the continuity of these complex applications requires the use of software-managed infrastructures called *high-availability clusters* (HACs). HACs employ sophisticated techniques to monitor the health of key enterprise application layers and of the resources they use, and to seamlessly restart or relocate application components after failures. In this paper, we first describe the manifold uses of HACs to protect essential layers of a critical application and present the architecture of high availability clusters. We then propose a taxonomy that covers all key aspects of HACs—deployment patterns, application areas, types of cluster, topology, cluster management, failure detection and recovery, consistency and integrity, and data synchronisation; and we use this taxonomy to provide a comprehensive survey of the end-to-end software solutions available for the HAC deployment of enterprise applications. Finally, we discuss the limitations and challenges of existing HAC solutions, and we identify opportunities for future research in the area.

Keywords: Clustering, dependability, enterprise system, high availability, high availability clusters, reliability

1. Introduction

High-availability clusters (HACs), also called *failover clusters*, are software-managed systems that support the reliable execution of complex enterprise applications (EAs) or of their key layers¹ and components. They play a major role in ensuring the high levels of availability required of today’s mission-critical EAs [2, 1, 3]. They comprise physical servers, storage, communication and other hardware infrastructure, together with sophisticated HAC-management software. This software is responsible for continuously monitoring the protected EA layers (e.g., the application server and database layers), and for seamlessly mitigating EA component failures through failover or through automatically restarting, repairing or relocating failed components. As such, HACs enable the delivery of essential EA services with minimal downtime.

A key requirement for HACs is to ensure the continued operation of the single-point of failure (SPOF) components of the protected EA layers. Such SPOF components may include databases, distributed transaction coordinators, software load balancers, and storage. Due to the diverse high availability (HA) needs of mission-critical EAs, HACs must comply with a wide range of additional requirements. These requirements differ significantly from one EA to another. For example, enterprise resource planning (ERP) EAs facilitate transactions (e.g., online transaction processing), and therefore must be deployed on HACs capable of ensuring the atomicity, consistency, isolation and durability (ACID) requirements associated with transaction processing [3, 4]. In contrast, enterprise analytics EAs

(e.g., online analytical processing) do not have these strict ACID requirements, since they tend to operate with read-only data, while transactional EAs are typically read-write intensive [5]. Moreover, the focus of the analytical EAs is to manage large data sets in multiple steps (e.g., staging, transformation, processing, and reporting), and this is typically reflected in the architectural layers and the components that are part of such layers [6, 7, 8]. Hence, the SPOF components of analytical EAs differ from those of transactional EAs. Therefore, the HACs used to protect such solutions vary significantly.

Further HAC requirements arise from the need to monitor and maintain the “health” of the cluster itself. An essential monitor called a *heartbeat* [9, 10] is required to periodically check the health of individual cluster nodes (i.e., servers) so that the appropriate failover procedure can be initiated when node failures are detected. At cluster level, a *quorum* system [4, 11] is needed for scenarios where the cluster ends up divided into cluster partitions that can no longer communicate with each other. In these scenarios, a voting protocol is enacted to select a single partition that will continue to run the EA. In this way, quorum systems prevent the occurrence of a *split-brain* [1, 3], i.e. a situation in which multiple partitions attempt to use the EA resources at the same time, potentially corrupting important EA data [9].

This diversity in HAC uses and requirements has led to significant research on the techniques underlying the operation of HACs. At the same time, the ability of HACs to run critical EAs with minimal downtime prompted the development of multiple end-to-end HAC solutions. Our article provides a taxonomy and a survey of this large body of work. The taxonomy clarifies and formalises the often overlapping or conflicting terminology

¹We use the term *layer* for a logical component of an application, and the term *tier* to denote a physical structure, as recommended in [1].

and classifications used by HAC researchers and developers. The survey comprises two parts. The former part, covered when we define the taxonomy, represents an extensive coverage of the HAC research landscape. This part offers insights into the capabilities and limitations of the techniques used to achieve HA for critical EAs. The latter part covers end-to-end HAC solutions, supporting the developers and users of these solutions.

Addressing HAC limitations and extending the applicability of HACs to new computing paradigms is an active field of research. Recent advances in areas such as machine learning [12, 13] and self-adaptive systems [14] provide new avenues for addressing current HAC challenges, while recent technologies such as containerisation [15, 16] require the development of new types of HACs. A discussion of these new directions for HAC research and solutions is also provided in the article. This is the first taxonomy and in-depth survey that focus on HACs. While a few previous studies proposed taxonomies for availability in the cloud [17, 15] and dependable and secure computing [18], these taxonomies are complementary to our work, as they do not consider essential HAC characteristics such as heartbeat, quorum, topology and type of cluster.

The contributions of our article are organised as follows. Section 2 explains how HACs are used to protect different layers of critical EAs, and introduces a generic HAC architecture and key HAC terminology. Section 3 presents the HAC taxonomy and the techniques underpinning core HAC operations such as monitoring, heartbeat, quorum, failure detection, and component failover. Section 4 uses the taxonomy to survey end-to-end EA HAC solutions available commercially or from open-source projects. Section 5 discusses HAC limitations, open challenges, and research opportunities. Lastly, Section 6 concludes the article with a brief summary.

2. Uses and Architecture of High-Availability Clusters

2.1. Key Concepts and Terminology

The ISO/IEC 25010 standard defines availability as the ‘*degree to which a system, product or component is operational and accessible when required for use*’ [19]. Availability is calculated as the ratio between the time when a system is operational and the total time over which the system was observed. Equivalently, availability can be computed as the ratio between the *mean time between failures, MTBF*, and the sum of the mean time between failures and the mean time to recover after failures, i.e., the *mean time to repair, MTTR*: $availability = MTBF / (MTBF + MTTR)$ [20, 21].

Component failures lead to *downtime* (i.e., periods when the system is not operational or accessible), and to a decrease in availability. As such, HACs are responsible for reducing both the frequency and the duration of failures, and thus their impact on the availability of the protected EAs. Discharging the first responsibility involves monitoring specific EA components, to identify and resolve *faults* before they lead to *errors*, and errors before they trigger *failures*, i.e., violations of requirements observable to EA users [20].

A fault can occur in any resource (i.e., atomic component) of an EA, and the critical resources are usually combined into one or several SPOFs (or SPOF groups). If such a resource fails irrecoverably, it will lead to the failure of its associated SPOF as well. When an SPOF fails, it may bring down an entire application. Achieving high availability requires that the SPOFs of an application are entirely or partially eliminated, or masked. Consequently, HACs discharge their second responsibility by relocating SPOF-related resources to a secondary server after irrecoverable failures. In this way, they mask the failures of resources, and thus of application SPOFs.

In this context, HACs employ a threefold strategy for failure management:

1. HACs avoid EA downtime, even in the presence of failures of individual resources. To achieve this, HACs reinitialise or restart resources after faults and errors (increasing MTBF) and after failures (reducing MTTR).
2. HACs promote the failure management to a resource group level if the failure at a resource level cannot be resolved locally. This leads to a failover of the concerned resource group to another node. A resource group can also be reinitialised on the same node if there are no available secondary nodes. A resource-group failover is faster than a complete system failover. Therefore, the likely outcome is that the failover does not cause downtime.
3. If there are dependencies between the resource groups and after critical failures, a complete system failure may occur. In this event, the complete system is failed over to another node.

In the first scenario, components are restarted, whereas in the other two scenarios components are first stopped and then started, in a specific order determined by their interdependences.

2.2. Enterprise Application Layers

EAs such as ERPs are transaction-intensive and require stateful communication. Moreover, data consistency and data integrity are vital for such applications. Additionally, modern EAs are highly integrated, which means that data corruption in one application may lead to data corruption in other integrated systems. Therefore, data corruption and data loss must be prevented even when failures occur. To identify and achieve HA holistically for an EA, it needs to be broken down into a set of essential building blocks that are referred to as *layers*. Critchley [4] proposes a layered architecture in describing an IT environment. Somasekaram [22] suggests a similar approach of separating the layers of an IT solution for outsourcing purposes.

When all the layers of an EA are identified, an appropriate solution for ensuring the HA of each layer can be devised. Multiple solutions are typically possible for each layer, including the use of a HAC. As such, different EA critical layers can each be protected by a separate HAC. Alternatively, a single HAC can be employed to protect several critical layers of an EA. In either case, any EA layer not protected by HAC(s) may

require other types of HA solutions (e.g., redundancy or fault tolerance). In the special case of applications with only one critical layer (e.g., firewalls), HA can be ensured through using a *single-layer HAC* [23, 1, 24].

Based on the solutions that can ensure their availability [4, 22, 25, 26, 27, 28, 29, 30, 31, 32, 33], the components of an EA can be organised into the nine layers from Table 1. For each layer, the table shows the typical role(s) that the layer can play within an EA, the solutions available for ensuring its availability, and whether an *application HAC* (i.e., a multi-layer HAC) is among these solutions—providing protection for the whole layer or only for its client resources. As indicated in this table, an application HAC can protect the *application server*, *application core*, and *database* layers of an EA, as well as the client resources associated with the *EA network* and *storage* layers.² In contrast, a HAC is not typically used to protect the *operating system*, and the *virtual machine* (VM) or *server* layers of the EA (i.e., layers 4–6 from Table 1), as a failover always involves relocating the application environment to a different VM or server, respectively. The protection of the *data centre* layer is also beyond the scope of a HAC. However, the HAC still needs to monitor critical elements from layers 4–6 in order to identify critical issues such as unacceptably high levels of CPU utilisation for a server.

A few research initiatives have addressed the challenges of achieving HA solutions for multiple EA layers from Table 1. Bajohr et al. [25] have devised an HA framework for Springer Verlag’s Online Conference Service. Their framework combines different solutions for several layers of this multi-tier applications, including an N+M HAC (the terms are explained in the *topology* section of our taxonomy) for application servers (layer 1), and a master-slave configuration for the database (layer 3). Similarly, Sun et al. [51] present an HA architecture for a multi-tier application in which multiple HA solutions are combined to enable HA for the application. However, most research to date has focused on HACs for single EA layers. For instance, Cheng et al. [52] developed an application cluster service (APCS) scheme comprising separate methods that support state recovery and failure management, respectively. APCS assumes that the state of a shared-storage database layer does not change, and therefore focuses on the protection of the application layer of a three-tier architecture. In many other approaches, the layer protected by different types of HA solutions is the database layer [53], as in the case of Riley et al.’s HA cloud for research computing [54]. Built using the OpenNebula cloud computing platform, this solution employs an active-active HA MariaDB cluster (layer 3) to support the storage of cloud objects.

In summary, modern EAs require a combination of HA solutions to achieve the required levels of end-to-end availability. More often than not, the infrastructure components of EAs have

²The network and storage EA layers are part of the EA infrastructure, and present both a server view and a client view. As an example, a storage system in itself is part of the server view, while its individual disks associated with a server or with a virtual machine become part of the client view, and thus need to be protected by the application HAC.

their own HA setups, and thus HACs typically focus on ensuring the availability of the actual applications.

2.3. HAC Architecture

Figure 1 shows the high-level architecture of a generic HAC operating on $n \geq 2$ nodes distributed across one or multiple locations (i.e., data centres). The HAC is responsible for the management of an EA whose resources are depicted organised into $m \geq 1$ *resource groups*, out of which only the resource groups on the *primary node* 1 are active. The HAC uses three dedicated private networks—a *cluster network* for its communication across cluster nodes, a *quorum network* to connect all nodes to a quorum device (i.e., a facilitator of the quorum service), and a *heartbeat network* whose role is explained later in this section. To HAC modules deployed on each cluster node [9, 10, 55] are described below, using the terminology summarised in Table 2.

I. Cluster management is the core HAC module, responsible for overseeing the operation of the other modules, and including the following sub-modules:

1. **Cluster data**, comprising the data stores managed by a cluster, and shared by all nodes.
 - **Configuration** data comprises static data (e.g., HAC configuration parameters).
 - **Runtime** data consists of dynamic data (e.g., current status of the cluster components).
2. **Communication**, which manages the communication between the HAC modules on the same node and between the cluster nodes, and the heartbeat communications.
 - **Cluster** communication (also known as intra-cluster or inter-node communication) deals with communication between cluster nodes.
 - **Heartbeat** is an essential health monitor that checks the health of member nodes, notifying the HAC when the heartbeat of a particular node fails [3, 10]. During such an event, the HAC consults with a quorum to ensure that there are enough votes to continue to run the cluster. If it is the active node that has failed, this will result in a failover, provided that the cluster can reach a quorum [9].
 - **Node** communication (also known as intra-node communication) manages communication within a node.
3. **Resource management**, responsible for managing two main groups of EA resources:
 - **Base** resources, which include key components such as CPUs and disks.
 - **Application** resources, which are resources specific to the HAC-protected applications.

II. Failure detection and recovery is the HAC module responsible for managing failovers and recoveries, and includes the following sub-modules:

1. **Monitoring**, which monitors the EA resources and notifies other HAC components (e.g., resource management) about any problems.

Table 1: Enterprise application (EA) layers with possible high availability (HA) solutions

No	Layer	Typical Role(s) within the EA	Possible HA solution(s)	HAC [†]
1	application server (e.g., web servers [34, 35, 36])	key tier in multi-tier EAs, e.g., presentation layer	use multiple instances with load balancing [9, 37]	optional
2	application core (e.g., ERP central [38, 39])	coordination of distributed transactions, application servers	use application HAC [39, 1, 40]	yes
3	database (e.g., Oracle, DB2, HANA [1, 41])	databases to support the main application	high-availability features provided by database, such as replication and mirroring [1, 42, 40] which can be used with application HAC	yes
4	operating system (e.g., Linux, UNIX)	operating environments	redundant server environment	no
5	virtual machine (VM)	VM (e.g., virtualisation platform)	VM cluster [43] (a HAC can be combined with a VM cluster [44])	no
6	server	server hardware	redundant servers and fault tolerance	no
7	network (e.g., private, public networks)	local area network (LAN), virtual LAN (VLAN)	redundant network devices, fault tolerance, hardware HACs (e.g., for routers and load balancers) [29, 45, 46]	only client resources
8	storage (e.g., the different type of storage systems)	storage area network (SAN), NAS, direct attached storage (DAS) [3]	redundant devices, fault tolerance, storage HAC [3, 26, 47]	only client resources
9	data centre (e.g., essential data centre components)	supporting utilities such as UPS, power distribution unit (PDU) [48], cloud operating systems [49], and backup infrastructure	redundancy by multiple sites, and redundant data centre equipment, such as UPS and fault tolerance for components, HA for the individual data centre components [29, 50]	no

[†]High availability of layer can be ensured by an application HAC

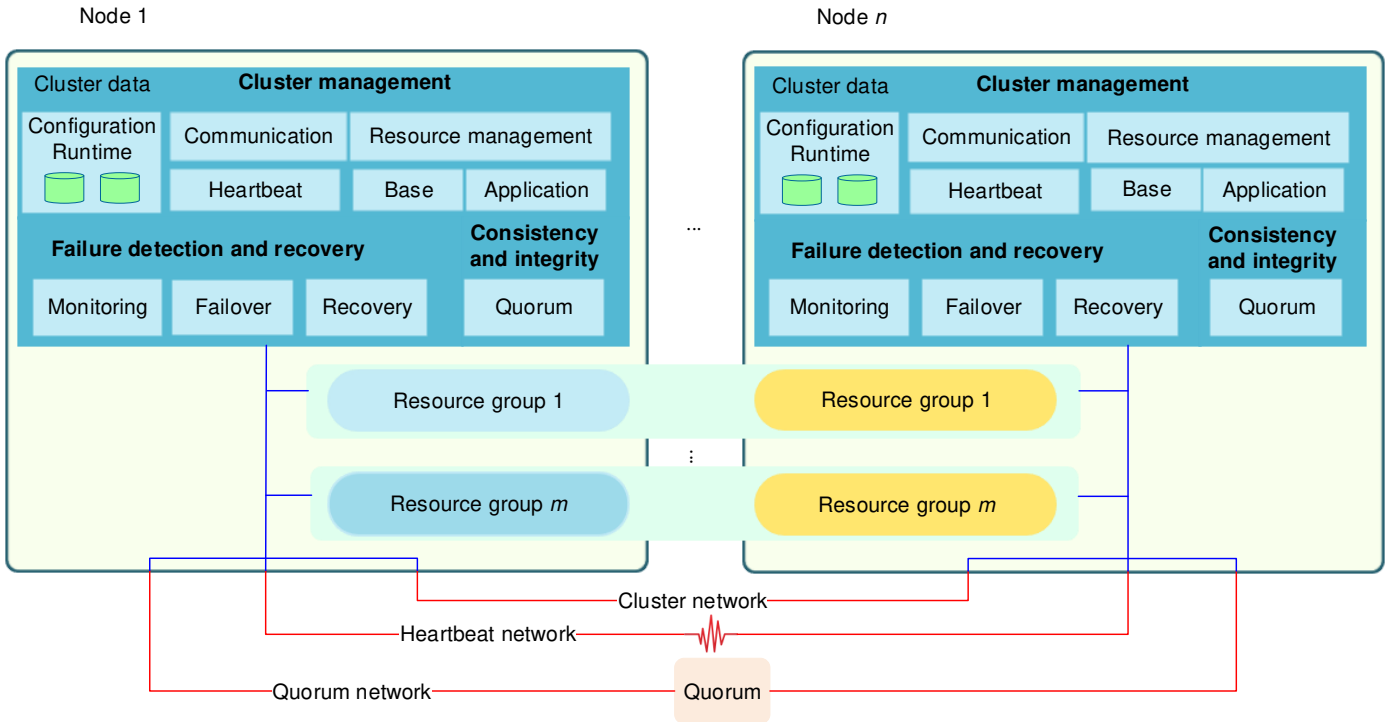


Figure 1: Architecture of a high availability cluster (HAC) with $n \geq 2$ nodes.

2. **Failover**, which is responsible for moving resource groups to a secondary node. Depending on the failure type, a failover can be at resource group or system (i.e., complete application) level. The latter involves moving all

resource groups that belong to an application [3, 4].
 3. **Recovery**, which decides whether failures need to be resolved at resource, resource group or node level by considering their criticality and resource dependencies. When

Table 2: HAC terminology

Term	Description
Resource	A logical or physical component of an EA layer (e.g., an IP address used by a database, or an application component) that is managed as an atomic entity by a HAC, and is either fully operational or unavailable. Resources have interdependencies that can be described by a hierarchical map [9, 10, 3].
Resource group	A set of logically related resources that can be relocated to a secondary node as one entity. As such, each resource can only belong to one resource group. An EA may comprise numerous resource groups, each of which may represent a significant part of the EA, such as a database [1, 9, 10].
Split-brain	A condition that occurs when a cluster ends up divided into partitions that perform conflicting operations on the same resources, typically causing data corruption [3, 56].
Amnesia	A condition that occurs when cluster nodes operate with different configurations, e.g., because nodes that are rebooted resume operation with an older configuration. If such nodes are to become primary, a problem is created because they will run with an out-of-date configuration [57].
Switchover	The manual migration of resources from one node to another [4].
Shared storage HAC	A HAC whose cluster members have access to the same storage. However, when it comes to EAs, typically only one node at a time can allocate the shared storage resources, so that data integrity is not affected [9].
Dependency	Resources and resource groups have dependencies that must be taken into account during a failover and the subsequent restart of services. These dependencies can be modelled using a acyclic directed graph termed a <i>dependency configuration</i> [9].

a failure cannot be resolved, the failover sub-module is notified, so that failover can be initiated.

III. Consistency and integrity is the HAC module that ensures consistency and integrity across all cluster nodes through the following sub-modules:

1. **Fencing**, which is a protection mechanism that isolates a resource or node that experienced failures, removing its ability to connect to any of the critical EA resources [4, 58].
2. **Quorum**, which is a voting system for determining which partition is allowed to run a cluster when a split of the cluster occurs [4, 10, 59]. The partition that has the quorum is considered quorate, and can be used to run the cluster without causing a split-brain.

These terms are described in greater detail in the taxonomy section below.

3. Taxonomy of High-availability Clusters

Our taxonomy applies to single-layer HACs and multi-layer HACs, which it organises into the eight top-level classes shown in Figure 2. The first four classes capture how HACs are deployed (*deployment patterns*), which EA layers are protected by HACs (*application areas*), how this protection is achieved (*type of cluster*), and how the HAC nodes are structured and interconnected (*topology*). The next two classes reflect how HACs manage the resources of the protected EA (*cluster management*) and perform detection of and recovery from failures of these resources (*failure detection and recovery*). Finally, the last two classes indicate how the HACs preserve the consistency of the EA data and the integrity of the cluster (*consistency and integrity*), and how the EA data are synchronised across cluster nodes (*data synchronisation*).

The taxonomy aims to achieve a balance between: (a) considering virtualized resources explicitly where their use makes a significant difference in how a HAC aspect is implemented or operates (e.g., shared-storage or network); and (b) keeping taxonomy classes and subclasses non-prescriptive about the use of virtualised or physical resources where this choice has limited impact (e.g., for monitoring, fencing, heartbeat or quorum).

Given the significant industrial relevance of HACs, we developed the HAC taxonomy based not only on research papers but also on a wide range of technical guidebooks [1, 3, 4, 59, 60, 61, 62], technical reports [63, 64], white papers [65, 66, 67, 68], HAC product documentation [39, 69, 40, 70, 71, 72, 73] and best practices [36, 44, 74, 75].

3.1. A: Deployment Pattern

The *deployment pattern* of a HAC represents the platform **where** the HAC solution is deployed. As shown in Figure 3, we distinguish between the deployment *environment*—which case be public cloud, or can be on-premise, fog or edge IT infrastructure of the organisation using the HAC, and the type of *host* used for the cluster—which can be physical, virtual or container.

The deployment pattern decides, along with business requirements and technical capabilities, **what cluster type** can be implemented for an *application area*. Table 3 describes the relationship between deployment patterns, application areas, and the rest of the taxonomy. A *cluster type*, on the other hand, decides **how** an *application area* can be protected and the subsequent topology and related configuration. An example of this is as follows: if the deployment pattern is a set of virtual servers in a single data centre, it will not be possible to deploy topologies such as *metro* or *continental* (described under *Type of cluster*). Thus each HAC solution comes with deployment restrictions (e.g., whether it can be deployed in a public cloud or not).




Figure 2: Top-level classes of the HAC taxonomy.




Figure 3: Deployment patterns.

Cloud environments impose restrictions that can cause problems for a HAC because many of the infrastructure elements that a HAC needs to monitor and manage may not be available for a cloud deployment. However, a distinction needs to be made between private (i.e., on-premises) clouds and public clouds because private clouds offer much more flexibility, and functionalities may be identical to an on-premises physical environment. Furthermore, the roles and responsibilities of different stakeholders play an essential role when deploying a HAC in a cloud. Table 4 describes the roles and responsibilities of customers and cloud providers for private clouds and for the service models available in public clouds [76, 77, 78], showing that multiple stakeholders may need to collaborate to support the different layers of a HAC in a public cloud.

Both the fog and edge computing paradigms have the potential to improve response times for EAs. However, ensuring HA could be a challenge considering the limited infrastructure components available; some of them are required for implementing HACs (e.g., shared storage) in these deployment environments [79, 80]. Moreover, the deployment hosts can also change, such as using containers to host an application, in which case a HAC must understand the implementation in order to deliver the required HAC functionalities. A HAC requires continuous monitoring for the protected application and the operating environment to detect failures and resolve them at a granular level. For example, suppose that an application is deployed in a container using the deployment environment edge. The HAC must then ensure that the key resources of the application deployed in the container can be monitored and that the HAC can collaborate with the responsible container orchestration system to ensure that failure mitigation actions can be performed. This requires the HAC and orchestration system to collaborate and support each other. We discuss this in detail in Section 4.4.

3.2. B: Application Areas

Application areas are the different IT solutions that can be protected by HACs, and a list of typical applications areas is presented in Figure 4. Considering the application area has dual purposes: (1) to identify if HACs can support the multiple layers that an application is composed of; and (2) to address all

areas that are part of an IT solution, so that HA requirements for those areas can be achieved. For instance, the application area *enterprise system* may require other related areas, such as *application server*, *database*, *server*, *network*, and *storage* to be included to ensure that the *enterprise system* is protected across all critical layers. Some layers can be protected by an application HAC while others may require a different set of options which may include application area specific HACs (presented in Table 1) [33]. Moreover, application areas with fewer layers may need to protect fewer components [81]. For instance, a HAC in the context of a distributed system (e.g., high-performance computing—HPC) may need to protect fewer components than an EA HAC. In case of an HPC, a head node (principal node) is identified as a SPOF. Thus, a HAC can be deployed to protect the head node [82]. Therefore, the application areas of a solution are determined dynamically during an implementation phase, and the numbers of protected resources will change with the type of primary application to be protected.

Several recent projects have implemented HACs that support multiple application areas, as also discussed in Section 2.2. Xiong et al. [53] present a HAC for a relational database in a multi-cloud environment which supports the requirements of both HA and Disaster Recovery (DR). Engelmann et al. [83] have experimented with a HAC to protect the head nodes of an HPC environment. Addressing complex systems that consist of multiple layers, hence also several application areas, is a challenge. Wang et al. [32] address the challenge by proposing an HA solution for a comprehensive medical system which consisted of several layers hence also multiple application areas. The proposed solution used a multitude of HACs to enable HA across the different layers.

3.3. C: Type of Cluster

The type of cluster plays a vital role in selecting the right topology and related configuration for a HAC. An important characteristic is the distance between nodes, and therefore the number of sites (e.g., data centres). The type could be chosen to meet business requirements, such as business continuity or DR.

Table 3: Connection between deployment patterns, application areas and the rest of the taxonomy

	Deployment Patterns	Application Areas	Rest of the Taxonomy
Objectives	Where to deploy the solution?	What application or application components need to be protected?	How should the solution be set up to meet the requirements?
Examples	data centre locations, public cloud, virtual server, container	enterprise system, NAS, network appliance (e.g., firewall), storage system	cluster type, topology, replication, mirroring

Table 4: Roles and responsibilities for service models in a public cloud, and for on-premises deployment

No	Layers	On-premises	IaaS	PaaS	SaaS
1	Application server	C	C	C	AP
2	Application core	C	C	C	AP
3	Database	C	C	C	AP
4	Operating system	C	C	CP	CP
5	Virtual machine	C	CP	CP	CP
6	Server	C	CP	CP	CP
7	Network	C	CP	CP	CP
8	Storage	C	CP	CP	CP
9	Data centre	C	CP	CP	CP

Key: IaaS – Infrastructure as a Service; PaaS – Platform as a Service; SaaS – Software as a Service;
C – Customer; AP – Application provider; CP – Cloud provider (who may also be application provider for the SaaS service model SaaS)

A DR solution requires at least two data centres with a sufficient distance between them and a related configuration. When a HAC solution is explicitly deployed to support DR, it has to comply with restrictions (e.g., low network latency between data centres). Moreover, supplementary mechanisms must be used to guarantee data integrity during failovers. Therefore, the type of cluster should be treated as the starting point for HAC selection, along with the two top-level taxonomy classes presented previously. There are four types of clusters, as shown in Figure 5 and Table 5. Based on a rule of thumb derived from [3, 1, 65, 84, 85], we assumed a communication speed of 3 ms per 160 km to calculate network latency. The distances described in the table may differ due to the use of different technologies. Moreover, the different HAC solutions can also come with specific recommendations.

The limitations presented in Table 5 are due to the different classes of DR that can be supported for the different types of HAC [1, 62], which in turn are related to the distance between the data centres. For example, the DR class “data centre network failure” can be supported by the campus, metro and continental cluster types. However, a DR class of flooding in the area can only be supported by metro or continental. On the other hand, the continental cluster type can support all DR classes, including the most severe ones, such as an earthquake in the region. Therefore, the most severe DR classes are only supported by the continental cluster type.

C.1: Local. A local HAC is hosted in one data centre and uses one storage system, usually shared. When there are two data centres, the distance between data centres is often less than one km [58]. In such case, there exist two options. Option 1 is to distribute the HAC nodes across two data centres, with all nodes utilising shared storage from one of the data cen-

tres. Option 2, on the other hand, uses two storage systems in the two data centres, with the HAC becoming a shared-nothing cluster. However, because data integrity is crucial for EAs, either replication or mirroring must be enabled to synchronise data between the two data centres. The two-data centre setup with replication or mirroring is also a feasible solution for other types of cluster. Since there is usually one data centre associated with a local cluster, the setup is not compliant with DR requirements.

C.2: Campus. A campus cluster is usually deployed across two or more data centres, and the distance between the data centres is less than 30 km [86, 58]. Since a campus HAC has a redundant setup for data centres and related components, it can comply with DR requirements (e.g., it can handle DR scenarios such as a data centre failure). However, the distance requirement between data centres means that businesses may opt for other types of HACs which are optimised for longer distances. Nevertheless, campus clusters can support longer distances when combined with other HAC types, becoming *hybrid* clusters — for instance, multiple interconnected campus clusters with one campus cluster functioning as the primary. This setup enables failover locally for most incidents but will trigger a failover to a different site only when a DR scenario takes place at the primary site.

C.3: Metro. In a metro cluster, the nodes are distributed across a distance of up to 300 km. Although there is no definite cut-off for this distance, the restrictions come from the techniques that are employed to synchronise data [85, 87]. For example, in some cases, the distance can be extended to 400 km by employing Wave Division Multiplexors (WDM) [88].

C.4: Continental. When cluster nodes are geographically dis-




Figure 4: Application areas of HACs.

Table 5: Type of clusters and potential configurations

Type of HAC	Distance (in km)	Network Latency (ms)	Data Centres	Storage Systems	Disaster Recovery Support
Local	≤ 1	≤ 1	≥ 1	≥ 1	No
Campus	≤ 30	< 1	≥ 1	≥ 1	Limited due to short distance
Metro	≥ 30	< 5	≥ 2	≥ 2	Limited due to distance
Continental	≥ 300	> 5	≥ 2	≥ 2	Yes




Figure 5: Type of cluster.

persed, usually at a distance of more than 300 km, the cluster is characterised as a continental cluster. [89]. A continental cluster can also be referred to as a global cluster or geo-cluster.

3.4. D: Topology

The topology (or *redundancy model* [90]) of a HAC represents the way in which the HAC nodes are structured and linked. The topology of a HAC (Figure 6) depends on multiple characteristics of its nodes, on the roles of these nodes (primary or secondary), and on its communication devices, networks, storage systems, and supporting tools (e.g., quorum devices).

D.1: Symmetric. In a *symmetric* topology, all cluster nodes can be utilised concurrently: there is no standby node.

D.1.1: Active-active. While symmetric active-active describes that all nodes are utilised, there have been research efforts to implement variations of the topology to address the specific needs of distributed systems. Engelmann et al. [91, 92] implemented a prototype with a symmetric active-active topology that operated on more than two nodes to provide HA for an HPC. The prototype employed two replication mechanisms, internal and external, using reliable and totally ordered message delivery. The internal replication provided synchronisation for the HPC file system metadata service, while the external replication supported the same for the HPC job and resource manager [92]. The evaluation of the prototype showed that the availability could be improved significantly as more nodes were added to the cluster. Therefore, depending on how applications are hosted on such HACs, we distinguish between *symmetric application-based* and *symmetric server-based* topologies.

D.1.1.1: Application-based. In a symmetric application-based topology, an application is active on all available cluster nodes. This topology requires application support because managing transactions across multiple nodes is only possible by using additional mechanisms, such as distributed lock management. A component of an application, for instance, a database, may provide these mechanisms, which can then be combined with a HAC solution [69]. For example, IBM Purescale supports parallel access to IBM DB2 databases [41], and Oracle provides active-active concurrent access support for Oracle databases using the Oracle Real Application Clusters [51, 70, 93].

D.1.1.2: Server-based. A symmetric server-based topology is frequently referred to as an active-active topology, and this implies that multiple applications are hosted on all server nodes of a cluster; hence, the servers are fully utilised [40]. Since all servers are utilised, the topology is considered active-active. When a failover takes place for one or more applications, they failover to one or more of the available servers, implying that a standby node is not required.

D.1.2: N-to-N. In the *symmetric N-to-N* topology, multiple applications share the same set of N servers, like for the symmetric server-based topology. Upon failure of a primary node for an application, the application is failed over to one of the predefined member nodes of the cluster [94]. The new server will then host both the application that has failed over, and the previously running application [40]. The topology supports failing over multiple applications to multiple nodes.

D.2: Asymmetric. An asymmetric topology is an active-passive configuration in which one node is active while one or more nodes are in a passive or a standby mode [95].

D.2.1: Active-passive. An active-passive topology is the typical asymmetric topology consisting of a two-node cluster setup in which one node is active while the other node is passive or standby. This topology is sometimes referred to as 2N redundancy [64]. Today's HACs make a distinction between the different layers of an application. In protecting a layer 3 component (i.e., database), a HAC can either manage it by employing a database-specific extension (agent) or utilising replica or mir-




Figure 6: HAC topology.

roring features that are offered natively by the database [96]. Most database vendors provide a replica or mirroring option to set up standby databases of primary databases [97], and this configuration can effectively be integrated with a HAC. The prerequisite in such a case is that the HAC has support for the specific feature so that the HAC can recognise and support it as part of its operations.

Several variants of the active-passive topology exist, depending on the set up for a standby database and for the secondary node [41, 4, 1], as shown in Table 6. While the standby modes from this table are often used with databases, other application layers may also employ a similar configuration. For example, a layer two component (i.e., application core), employs the active-warm (or warm) standby mode due to the limited need for data synchronisation. However, a common setup by a HAC is to employ either active-hot (or hot) standby or active-warm because otherwise failover time and MTTR will increase and, as a consequence, availability will go down. The primary reason for using the active-cold (or cold) standby or active-warm standby is cost, as using an active host node is associated with higher costs. The standby modes are usually not explicitly supported by modern HAC solutions; instead, the different standby modes of databases and related features that are supported are specified [41, 1]. Moreover, the standby modes are frequently used to refer to the modes of the data centres, particularly in the context of establishing DR for a system [99].

D.2.2: N-to-1. In an N-to-1 topology, multiple applications are supported by one dedicated standby node. Hence the name N-to-1 [94, 40]. If a node fails, the application is failed over to the standby node and made available there temporarily. However, while the application is active on the standby node, there will be no HA for that application until the primary node is back online. Another aspect of an N-to-1 topology is that such a standby node must be able to host all N applications simultaneously. Hence, sufficient capacity must be available on the standby node.

D.2.3: N+1. In an N+1 topology, one passive (spare) node supports multiple active applications, similarly to the N-to-1 topology. However, unlike the N-to-1 topology, the N+1 topology employs a *rotation scheme* for failovers [40]. This means that, during a failover, an application is failed over to the standby node, but the failed node, once the problems are resolved, effectively becomes the standby node. Hence, any node in the cluster can become a standby node. A variant of the N+1 topology that

uses 2+1 nodes (with two active nodes and one node operating as a standby or backup) has been referred to as *asymmetric active-active* in the context of HACs for HPC [100].

D.2.4: N+M. The N+M topology refers to HACs that comprise N active nodes and M passive nodes in the cluster, and is called an N+N topology when the number of passive nodes equals the number of active nodes. The topology is employed when one passive node is not sufficient, and $M > 1$ passive nodes are required for failovers [90, 101].

3.5. E: Cluster Management

The cluster management module of a HAC is responsible for managing the resources, resource groups, nodes, heartbeats, cluster data, and failovers of a cluster, directly or through other modules. The characteristics used to distinguish between different types of HAC cluster management are shown in Figure 7 and described below.

E.1: Cluster data. Two types of cluster data are relevant to HACs: configuration and runtime. *Configuration* data contains configuration details of a HAC while *runtime* data stores status of the cluster components. Cluster data can be stored in three types of repositories: disk, file and memory.

A repository can be either local or shared. However, a prerequisite for a repository is that it is accessible by all cluster nodes. Hence, if a repository is local, a replication mechanism is used to replicate it between the nodes at regular intervals. However, in some cases, when persistent files are employed, the replication is a manual activity. Both in-memory and file repositories are local. However, there are differences in what cluster data type they support. Configuration data is static and is commonly stored in files, while an in-memory repository is generally used to store runtime data to capture changes in real-time. This means there is a rigorous requirement for in-memory repositories to replicate data to other nodes. Therefore, a designated process governs the synchronisation of the runtime data, for instance, Designated Coordinator (DC) in the case of a Pacemaker-based HAC [10, 58]. The coordinator ensures that one master repository exists in the primary node while a copy of it, a replica, is distributed across all the member nodes. A shared repository (e.g., disk or file share), on the other hand, stores both configuration and runtime data. In many cases, a quorum repository, which is shared, can support the requirements. Cluster data in a repository is organised using an information model. For instance, a Cluster Information Base (CIB)

Table 6: Active-passive topology variants

Standby Mode	Recovery Time	Data Synchronisation Method	Description
Active-Cold	Hours	Backup/restore	A secondary node is installed and configured but brought up only when the primary node is down. Subsequently, the related services are started, as well [98].
Active-Warm	Minutes	Mirroring, shared storage	The secondary node is installed and configured and is running. Related application services are started upon failure of the primary node.
Active-Hot	Seconds	Mirroring, replication, shared storage	The secondary node is fully installed and configured, and services are also started. The secondary node takes over responsibilities immediately upon failure of the primary node.

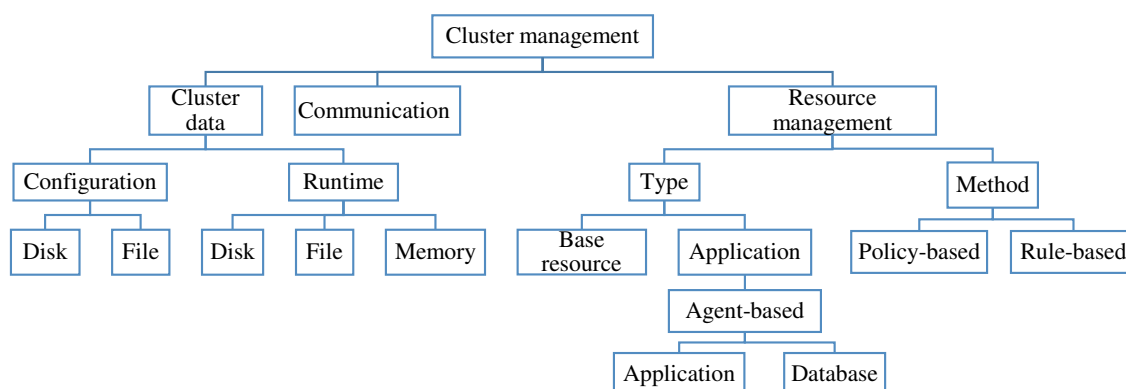


Figure 7: Cluster management.

uses an XML-based object model to represent both configuration and runtime data. However, there are no standardised information models for dealing with cluster data, and, as such, HACs use different information models. Open Service Availability Framework (OpenSAF), for example, employs Information Model Management (IMM), and objects represent the two types of data: configuration and runtime [61].

E.2: Communication. HACs can use different communication types and methods (Figure 8).

E.2.1: Type. A communication type describes the different kinds of communications that a HAC employs and is split further into three subclasses: heartbeat, node and cluster.

E.2.1.1: Heartbeat. A heartbeat is a form of intra-cluster communication. However, it is separated in the taxonomy to highlight its importance and use of additional resources, such as a dedicated network. The type and content of heartbeat messages differ from solution to solution. In some cases, a heartbeat message could be a simple ping or a keepalive to provide the status of a cluster node [102, 95]. Heartbeat communication use a LAN-based or a disk-based method [3, 40, 58, 103].

E.2.1.1.1: LAN-based heartbeat communication uses a Transmission Control Protocol/Internet Protocol (TCP/IP) network [3, 1]. Since heartbeat is a key component of a HAC, the recommendation for business-critical solutions is to set up a dedicated network, such as a virtual LAN, to facilitate heartbeat

communication [3, 40, 103]. With this approach, the heartbeat traffic is not disturbed or delayed by other kinds of traffic in a network, which could be the case if the network is shared. Furthermore, adding redundancy to a heartbeat network by using multiple networks is also a good option so that a single network does not become a SPOF.

E.2.1.1.2: Disk-based heartbeat uses a shared disk and also the SAN fabric as a means to facilitate communication [103]. In some cases, LAN-based and disk-based heartbeat types can be combined to create a full heartbeat service. If a heartbeat mechanism is not employed, an alternative and robust mechanism is required to detect node failures. Cheng et al. [52] propose a HA solution that employs a module that can detect whether a node is sick or not and subsequently forecast the time of failure. This renders the heartbeat setup to be nonessential in such cases. However, there is no information regarding how such a solution works when there are many nodes in a cluster.

E.2.1.2: Node communication is referred to as *intra-node* and deals with communication within a cluster node. The node communication uses internal communication schemes, for instance, inter-process communication (IPC) within a server. Two types of such communication exist in HACs: user interface and resource management.

E.2.1.2.1: User interface communication refers to the different means to connect to the cluster on a particular node, including Graphical User Interfaces (GUIs) for cluster adminis-




Figure 8: Cluster communication.

tation.

E.2.1.2.2: Resource management communication can belong to two subclasses: base resource and agent. *Base resource* describes the communication between the cluster resource management and those resources that are available as a standard (e.g., IP, CPU of a server). An *agent* describes the communication between the cluster resource management and the agents that are responsible for managing application-specific resources (e.g., database, EA components) [104, 40].

E.2.1.3: Cluster communication, also termed *intra-cluster*, *inter-node* (or resource group) communication, describes communication between cluster nodes. For a HAC, internal cluster communication is crucial. It is required for continuous communication between nodes regarding changes in configuration, the health status of nodes, quorum status, and failure notification. Furthermore, since cluster communication is often a basis for making necessary decisions by a cluster, the requirement for cluster communication is that it is enabled using an atomic (ordered) and reliable messaging scheme. Even though several HAC solutions use different types of cluster communication, a strict definition can be used to distinguish the two main types: runtime and configuration. Thus, cluster communication deals primarily with the synchronisation of cluster *configuration* and cluster *runtime* data (e.g., the status of the nodes).

E.2.2: Method. The types of communications utilise different transmission methods, and these methods can employ different protocols, such as UDP and TCP. Some HAC solutions employ custom protocols to meet the HAC-specific requirements, and an example is the Transparent Inter-Process Communication (TIPC) protocol, used by OpenSAF [61, 105]. The methods are further divided into four subclasses: multicast, broadcast, unicast, and IP socket.

E.2.2.1: Multicast. Multicast enables transmission from one node to multiple nodes. Thus, it can be characterised as a one-to-many (1:m) method. The receivers are usually a group of nodes, which means that a subset of cluster nodes can also be addressed [106, 107, 3].

E.2.2.1.1: Atomic. Atomic multicast (or total order multicast) implies that all nodes receive the same message in their sent

order [108].

E.2.2.1.2: Virtual synchrony. Virtual synchrony is an atomic multicast technology that supports reliable inter-process messaging. Corosync, the open-source communication protocol, employs the Totem Single-Ring Ordering and Membership (TOTEM) protocol, which is an example of implementation of virtual synchrony [104]. Engelmann et al. [109] present a multi-node HAC solution for HPC that employs virtual synchrony to support state machine replication between the nodes in a symmetric active-active topology.

E.2.2.2: Broadcast. This method supports one-to-all (1:n) transmissions. While multicast supports transmission to a group of nodes, broadcast transmits to all nodes [106, 1].

E.2.2.3: Unicast. This method facilitates transmission between two nodes, and it is characterised as a one-to-one (1:1) transmission [3, 106].

E.2.2.4: IP socket. An IP socket can also be used in some cases to facilitate communication between cluster nodes [106, 110]. However, the majority of the HACs do not support this method but rely on other methods.

Typically, cluster communication employs either multicast or broadcast, although unicast is also used in some cases. On the other hand, heartbeat communication employs either unicast or multicast [3]. The different communication types and methods may have limitations in cloud, fog and edge deployment environments. For example, routing a private IP address across subnets may not be possible in these environments, although such routing is required to failover from one subnet to another in a two-node cluster hosted in two data centres [111, 112]. A solution could be to use an overlay network technique, which provides a network, e.g., an auxiliary network, on top of the main network, allowing routing across subnets [113], thus enabling seamless failovers.

E.3: Resource management. Resources are structured hierarchically to form a resource group, and links between the resources define the relationships between the resources [3].

E.3.1: Type. HACs can manage two types of resources: base resources and applications (Figure 7).

E.3.1.1: Base resource. A base resource is a standard building block (e.g., IP address, file system) [4, 40, 58, 69]. A HAC can manage base resources without requiring additional tools. Hence, a distinction is made between base and application resources. While managing base resources is supported by all HAC solutions to different degrees, application support must be provided explicitly.

E.3.1.2: Application. Application management is the capability to manage application-specific functionalities and features. Since each application must be handled individually, an extension to a HAC is usually required [4, 114, 115]. Such addition is provided in the form of either an extension or an agent.

E.3.1.2.1: Agent-based. Agents manage two main types of applications: application and database. *Application* agents deal with managing several application-specific layers (e.g., application core of an ERP as in layer 2). *Database* agents manage database-specific components (layer 3). The application agent functionality connects application-specific (this includes both types: database and application) configuration and procedures with the resource management module of a HAC and supports functionalities including [114, 115]:

- Monitoring application-specific components
- An application-specific configuration, which can recognise the architecture of the application components
- Complying with application-specific dependencies
- Logging
- Procedures – to stop and start related application components in a specific order
- Supporting Application Programming Interfaces (APIs) or specification by the application vendor

However, not all HAC solutions can support all applications, as each will require separate lifecycle management. When an application changes, for instance, when it is upgraded, the HAC application agent may also need to be updated to reflect the changes. Likewise, when the HAC solution is upgraded, the application agent may also need to be updated. Thus, supporting a large number of application agents could be connected to much effort. Furthermore, such support may be subject to licensing conditions, and HAC vendors could treat individual application support as an extension to license terms.

E.3.2: Method. Two main methods are used when managing resources: policy- and rule-based. *Policy-based* resource management uses policies to configure conditions, and, when a particular condition is satisfied, appropriate action is triggered [55, 40, 115]. On the other hand, *rule-based* resource management uses one or more rules to make decisions and act upon them [3].

3.6. F: Failure Detection and Recovery

Failure detection implies detecting failures by monitoring and analysing monitoring output [116]. If the monitoring identifies

a status change in a resource or a resource group, it invokes recovery management to initiate a recovery. If the recovery is not successful, the recovery manager may initiate a failover of a resource group or even a system; therefore, failover is part of recovery management. Figure 9 depicts the top-level class with its subclasses.

F.1: Monitoring. HAC failure detection and recovery monitoring can be further organised into subclasses depending on its area, type and method. The area describes the monitored domains, while the type of monitoring addresses monitoring from a configuration point of view. In most cases, HACs can provide support for specific monitoring metrics; however, if there is no support, a custom approach where HAC users define their own monitoring metrics is adopted. The monitoring scope may also vary and can range from the simple state monitoring of a resource to the monitoring of a resource in a detailed manner [3, 1, 40, 103, 72]. Several research initiatives refer to the monitoring aspect of HACs as means to detect failures. Cheng et al. [52] present a state-based internal monitoring approach for the experimental cluster APCS+PEV, while Leangsuksun et al. [55] employ threshold-based monitoring for the cluster HA-OSCAR.

F.1.1: Area. The monitoring areas that a HAC can support play an important role in the overall solution. This is because monitoring is the process that collects details regarding monitored elements from different areas and delivers that data to the cluster management to make appropriate decisions. The areas that a solution can support can roughly be split into three subclasses: server, cluster, and application.

F.1.1.1: Server. Server-specific metrics focus on critical and non-critical monitoring elements of an operating system and a server level. Examples of metrics are CPU utilisation and memory utilisation [40, 69, 103, 117].

F.1.1.2: Cluster. Cluster monitoring implies that monitoring is enabled, even for the internal components of a HAC, including cluster-related processes and objects [117, 118, 40]. This approach enables a HAC to distinguish between failures of cluster and application elements, thus preventing making incorrect decisions.

F.1.1.3: Application. Application monitoring is usually administered by an application-specific agent or an extension that is specifically designed to support a particular application and its architecture. This implies that an application agent is aware of the internals of the application [93, 115].

F.1.2: Type. A monitoring type describes how the state of the resources is measured. There are two types of monitoring: state- and threshold-based.

F.1.2.1: State-based monitoring uses the state of a resource as a monitoring metric, and the states can be as simple as “up” and “down,” or the monitoring can be more elaborate and contain more states [40, 58, 103, 115].

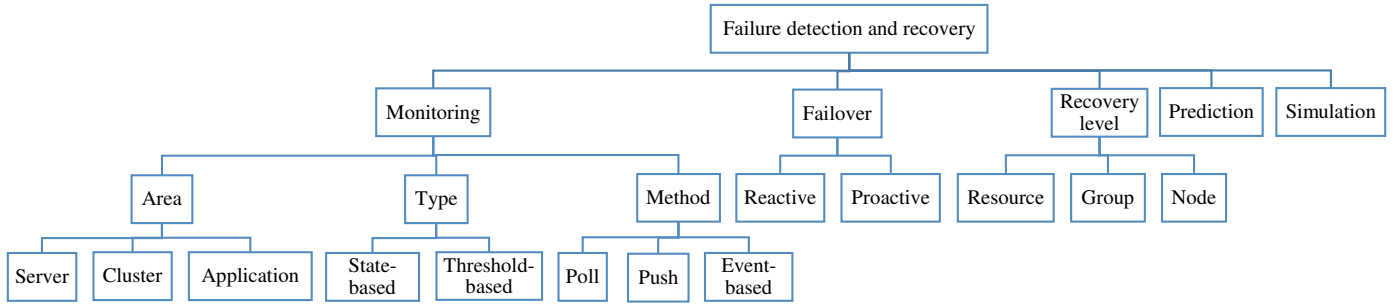


Figure 9: Failure detection and recovery.

F.1.2.2: Threshold-based monitoring uses a set of threshold values related to metrics [119]. As such, alerts with different severity levels can be generated, depending on which threshold is exceeded. While the threshold-based type gives the flexibility to configure monitoring at a granular level, it also adds complexity as the HAC must interpret all the different values and severity levels and act accordingly. One advantage is that a HAC will have more data that can be analysed, and decisions can be made at a granular level.

Even though state-based monitoring is the common type of monitoring, both monitoring types (and others) are sometimes combined. For example, OpenSAF HACs combine threshold-based monitoring with a type called *watermark* monitoring. The threshold-based monitoring is used to monitor system resources, while the watermark monitoring is employed to register the highest and lowest utilisation per configured resource [64].

F.1.3: Method. There are mainly three methods for monitoring, and they are push, poll, and event-based. *Polling* implies that the monitoring module of HAC and agents poll for state changes of resources periodically [119, 15]. On the other hand, *push* implies monitoring data is pushed to the monitoring module or agents [15, 119]. Such a setup will require additional enablers to interact with resources and push monitoring data to HAC agents. Polling is the most common method of monitoring [15], and it usually employs synchronous communication. However, this procedure is associated with a specific overhead. Therefore, other methods are studied by both industry and academia, and a technique that applies an *event-based* design is viewed as less resource-demanding. One type of event-based monitoring employs an intermediate module that interfaces with an operating system to capture instantaneous notifications relevant, for instance, the state change of a process. It passes that to an appropriate module of a HAC. An example of such a setup is the Intelligent Monitoring Framework (IMF) by Veritas [40]. The IMF has a monitoring feature integrated into an operating system for a particular resource so that state changes are captured instantaneously, and a relevant HAC agent is alerted. However, this approach requires specific development towards an operating system for a particular set of resources. Thus, IMF is not available for all types of resources but is being released gradually for different applications. Another recent development is to enable a HAC to interact with

the monitoring feature of an operating system directly [103], which means that the HAC needs only a slim variant of the monitoring module. The downside of this approach is that the HAC becomes highly dependent on the operating system and its developments.

F.2: Failover. Failover management includes procedures for failover and fallback, and all such actions are usually policy-driven. *Policy-based* indicates that policies can be associated with events so that the appropriate policies are triggered whenever a related event occurs [120]. Policies can be used, for example, to determine the target node for failover. Furthermore, policies can encode application-specific requirements, such as the order for starting up or shutting down resources. Failover management can further be split into two subclasses: reactive and proactive [121].

F.2.1: Reactive. A reactive measure uses policies to ensure the correct failover actions. There are two types of policies: static and dynamic. A static policy is created during the implementation or when applying manual changes, while a dynamic policy is created automatically by HACs to enforce policies based on runtime failure cases [122].

F.2.2: Proactive. A proactive measure assumes that a predictive model is employed to ensure that a failover can be initiated based on predictions [121]. The predictions can, in turn, use policies to trigger the required actions [55]. However, the proactive approach could be a challenge in HAC environments that deal with complex EAs because all relevant layers must be addressed in such cases while evaluating the HAC behaviour. Therefore, all active HAC solutions employ only the reactive mechanism.

F.3: Recovery level. The threefold strategy to manage failures is implemented using three recovery procedures: resource, resource group, and node (system) level [1, 40].

F.3.1: Resource. A resource-level recovery deals with recovery attempts on a resource level, implying re-initialisation of a failed resource while adhering to the dependency rules between resources. However, if this step fails, the failure is propagated to a resource-group level [1].

F.3.2: Group. A group-level recovery attempts to failover the entire resource group to a secondary node. However, if there

are no available secondary nodes, an attempt to reinitialise the resource group within the same node can also be initiated. If a resource group have dependencies on other resource groups, it may lead to a node (system) recovery [1].

F.3.3: Node. A node-level (system) recovery deals with failing over the resource groups to a secondary node. Moreover, a resource or resource group failure can also have a cascading effect due to dependencies, and, in such cases, it might lead to recovery on a node level. Since the previous node is labelled as "failed," policies may prevent any resources from being started there until that node is repaired [1].

F.4: Prediction. Prediction in the context of HACs implies that prediction approaches are used to provide prediction at the different levels to improve the operations of HACs. One such example is to predict resource failures; however, current HACs do not commonly employ prediction. Some research initiatives explore the area of predicting failures, but often within a limited scope. An example of such an initiative is the HA-OSCAR project which assesses prediction by evaluating hardware component failures [55]. The research team used a Hardware Platform Interface (HPI), as specified by the Service Availability Forum, to identify hardware events and, subsequently, to analyse such data in order to provide predictions [55]. Similarly, Lee et al. [123] propose a stochastic prediction model for node failure or a node-switch interconnected system failure of HA-OSCAR head nodes. Leangsuksun et al. [2] have also explored a failure-repair model for predicting the availability of HA-OSCAR cluster by using Stochastic Reward Nets (SRNs). While many of the prediction models use HA-OSCAR as the platform, some initiatives explore other platforms. For example, Cheng et al. [52] have used a module for a custom cluster solution that detects sick nodes and subsequently uses a prediction method to forecast the time-to-failure of nodes.

Both Veritas InfoScale Availability and Oracle Clusterware provide functionalities to simulate failures and observe potential paths to failovers [40, 124]. However, the objective of the subclass *prediction* is to ensure that the wealth of information that HACs produce can be incorporated to predict failures or optimise failovers. An example of such an approach could be using prediction to optimise the quorum voting process by dynamically evaluating scenarios.

F.5: Simulation. A simulated cluster or cluster simulation is a feature to run cluster simulations to study the potential failover and recovery paths when failure is simulated at a resource, multiple resources, resource groups, or a system level. Simulating failures and studying the cluster behaviour is becoming increasingly important for the configuration and optimisation of complex HAC solutions. The feature can be provided: (1) as part of HAC software; (2) as a separate tool; and (3) as part of the system management software.

Most HAC solutions provide a simulation feature with varying capabilities, and many simulators can be executed without interfering with the running cluster solutions. OpenSAF

provides a simulated cluster (method 1) with the source code distribution. The resulting five-node simulated cluster can be brought up quickly to evaluate the cluster and perform functional and API tests. Further examples of simulated clusters include the simulator functionality provided by the ClusterLab stack (method 1), which can be used to simulate failures to study the potential failover paths. The simulator feature continues to add more features to enable HACs to deliver optimal services. Red Hat Enterprise Linux HA Add-On provides a command-based tool (method 1) to simulate recovery scenarios. Veritas provides a standalone simulator (method 2) for the Veritas cluster server (VCS) to simulate and test different failover situations and use the results to optimise the cluster configuration [40]. IBM PowerHA SystemMirror for AIX delivers an advanced graphical cluster simulator as part of the IBM Systems Director, which is used to manage systems (method 3) [125]. The simulator supports creating and saving simulated cluster topologies. It enables conducting experiments with different configuration options, and one of the saved simulations can then be used to deploy the cluster when all the required components are in place.

3.7. G: Consistency and Integrity

Figure 10 presents the top-level class *consistency and integrity* and its subclasses. A HAC employs measures, such as a cluster lock or quorum, to preserve the data integrity of cluster resources and, most importantly, the clustered application by preventing harmful situations, for instance, *split-brain* and *amnesia*.

G.1: Cluster lock. Cluster lock is a technique used to lock cluster resources to a particular node, thus preventing other nodes from claiming the same resources. While a quorum-based approach could also be viewed as a cluster lock, a distinction is made to separate a quorum from a cluster lock. A cluster lock is a technique that does not employ a quorum-based approach but uses other means, such as a software-based lock mechanism. In the case of HACs, a distributed cluster lock is one such option, and an example is OpenSAF, which uses a global lock service to manage shared resources and ensure that only one node can access the resources at any given time [61]. Such configurations are deemed *quorum-less*.

G.2: Quorum. A HAC quorum serves two purposes: 1) maintaining cluster consistency by storing configuration and runtime data (e.g., cluster data) [126], and 2) managing a voting system required in the event of a cluster partition. For the latter purpose, the quorum hosts a voting mechanism in which every healthy and active node has a vote [127, 4]. Furthermore, the quorum also has a vote, a potential decider, hence the alternative name *tiebreaker*. Other names that are used to refer to the quorum mechanism are *arbiter*, *witness* and *voting system* [1]. When a partition of a cluster occurs after the failure of one or more nodes, a quorum is gathered to decide which partition should have the quorum. To reach a quorum, a partition must have a majority of votes [128]. The quorum service casting its vote can ensure that one of the partitions achieves this majority.




Figure 10: Consistency and integrity.

Ultimately, the majority cluster is allowed to run the cluster. If a quorum cannot be reached, the surviving nodes will shut down to ensure cluster consistency. The quorum collaborates closely with the heartbeat mechanism, as the heartbeat is the method used to identify unhealthy nodes. Additionally, the quorum or a similar service is required for fencing, as the two often collaborate to determine a quorum and subsequent fencing.

A quorum consists of a device and a process [4, 10, 84]. A device describes where quorum elements are stored, and a device facilitates the process, which uses an algorithm to calculate votes to achieve a quorum. The process employs a mode to determine what policy to use when performing the quorum voting.

G.2.1: Device realisation. Three types of devices can be used by a HAC: server, disk, and file share. A quorum *server* is a service that runs on a server that is usually hosted outside a cluster configuration [4]. The cluster is subsequently configured to connect to the quorum server. A *disk-based quorum* is based on a disk, which can be either local or shared [1, 4]. A *file share* uses a shared file location, and it can be ideal for geographically distributed HACs since member nodes do not have access to a shared disk [73]. The prerequisite for all quorum devices is that they support concurrent access by all cluster members.

G.2.2: Mode. Four modes are possible: server, node, disk and file share. The modes and the devices are an integral part of the quorum solution. However, the supported combination of devices and modes are specific to the different HAC solutions. The mode server uses the device *server*, and the device *Disk* is used by the mode *Disk*. Similarly, the device *File share* is used by the mode *File share*. While the devices *disk* and *File share* imply that they are storage points that are managed by the quorum process, a *quorum server* indicates an advanced device type. The mode *node* is implemented implicitly. Hence, it does not require any additional devices but uses the number of available nodes to decide, and the arrangement is referred to as the ‘majority node’ mode.

A quorum can be set up in different ways, and, in some cases, the several modes of a quorum can be combined. For example, Windows Server Failover Clustering (WSFC) supports a combination of devices and modes, as detailed in Table 7 [73].

However, the same combination is not always supported by other HAC solutions, and an example of this is that the Service-guard HAC does not recommend combining a quorum server with a quorum disk [84]. There are new quorum device types introduced to meet the advances in IT. For example, Microsoft has introduced recently a new quorum device called *cloud witness*, and the purpose is to support a server-based quorum in the Azure cloud, which could be ideal for cloud-based solutions [129].

The standard for all explicit quorum devices is that they are placed outside a HAC to avoid creating a quorum device as a SPOF. Moreover, redundancy of quorum is also preferred because the quorum is a critical HAC functionality. For this reason, most current HAC solutions support a dynamic reconfiguration procedure for quorum devices, which enables adding or removing quorum devices without impacting the running clusters. While quorum is crucial for a two-node cluster, it can also be opted out using a different mechanism. Furthermore, when a cluster has more than two nodes, an explicit quorum device could become optional because that cluster can survive the failure of a single node. However, a configuration using the mode *node* is still required to achieve a quorum. The research in this area focuses on enabling probabilistic approaches. For instance, Malkhi et al. [130] have explored a probabilistic approach to address both benign server failures and arbitrary (Byzantine) ones.

G.3: Dynamic quorum. While a quorum deals with static votes, a dynamic quorum calculates the number of votes and adjusts the quorum dynamically upon the failure of one or more nodes [129]. Thus, if a node is unavailable, it will effectively be out of the quorum voting process. This gives more flexibility to continue running a cluster even when other nodes fail. For example, dynamic quorum enables WSFC to run a cluster when only one node and a quorum device are available [131].

G.4: Isolation. HACs may “isolate” a particular node from the rest of the cluster, i.e., prevent it from allocating any resources. The objective of node isolation is to preserve data integrity by employing several mechanisms, such as putting a fence around a node (fencing) or shutting down a node.

G.4.1: Fencing. There are two types of fencing, node-level and

Table 7: Quorum implementation with Windows server failover cluster (WSFC)

Combination of Modes	Devices Realisation	Formula for Number of Nodes	Purpose
Majority Node	Node only (implicit device)	$n = 2k + 1$ (odd numbers)	Survive failures of $(n - 1)/2$ nodes.
Node and Disk Majority	Node and disk	$n = 2k$ (even numbers)	Survive failures of $n/2$ nodes when disk is available.
Node and File Share Majority	Node and file share	$n = 2k$ (even numbers)	Survive failures of $n/2$ nodes when file share is available.
No Majority: Disk Only	Disk only	-	Survive failures of $n - 1$ nodes when disk is available.

resource-level [4, 40, 58, 103]. The common implementation is to employ the node-level fencing [40, 58, 103].

G.4.1.1: Resource. Resource-level fencing isolates one or more critical resources and, by doing so, renders a node unusable because the node cannot allocate resources. Resource-level fencing can be based on a SAN switch, allowing only one node to connect to the SAN-based storage or SCSI. SCSI-based fencing often uses a SCSI-3 option called persistent reservation, which means there can be only one SCSI-3 persistent reservation per disk at any given time, making it an efficient method for isolating disks [40, 58, 103, 132]. Since resource-level fencing is based on storage input/output (I/O), it is sometimes called *I/O fencing* [133].

G.4.1.2: Node. On the other hand, node-level fencing acts at a node-level and isolates or quarantines the node completely [134]. In some cases, the node can be shut down instead, but the fencing functionality still manages the operation. Furthermore, the state of the fenced node is effectively changed so that it is no longer recognised as an active node by the cluster. Thus, the isolated node is not participating in any cluster operations.

G.4.2: Shutdown. A node shutdown is different from the shutdown procedure managed by the fencing functionality because it operates outside the fencing mechanism. This can be achieved by a HAC module that interacts with operating systems or servers using industry-standard specifications. Examples of APIs based on specifications are: Intelligent Platform Management Interface (IPMI) and vendor-specific embedded technology, such as Integrated Lights-Out (iLO) by HPE [58, 64, 117, 135].

3.8. H: Data Synchronisation.

Data synchronisation refers to the means, technologies and methods used to synchronise data between cluster nodes. The different layers of EAs require that data are synchronised to ensure consistency across all cluster nodes. Although a diverse range of synchronisation methods can be employed at the different layers, the overall responsibility for all layers managed by HACs lies with the HACs because they are responsible for failover management and ensuring data integrity. HACs may

employ additional tools or features that come with the application components to facilitate data synchronisation. Hence, we identify three principal areas of data synchronisation:

1. *Client-state* (i.e., session state replication) deals mainly with client connectivity (e.g., sessions), which means the client state of an application running on a primary node is synchronised with other cluster nodes [136, 137, 138]. Subsequently, a failover can occur seamlessly and without losing any connection data or affecting any active connections. Hence, other nodes can continue to support the connections instead. Client-state synchronisation is widely employed in HAC for the application area *network appliances* (e.g., firewalls) [139, 24, 45, 140].
2. *Cluster-state* employs different methods to synchronise the state of a cluster, and it can be considered as an intra-cluster activity. For instance, OpenSAF uses a checkpoint service to record the state of an application or a service. Subsequently, states are replicated to a standby application or service that is hosted on the secondary node [61]. A more advanced approach is a *State Machine Replication* (SMR) which creates replicas of client and process states to one or more nodes deterministically [141], which can even support more comprehensive solutions such as databases [142, 96]. An example of SMR concerning a HAC is an implementation of a HAC for HPC, which employed SMR to synchronise states between nodes in a symmetric active-active topology [83]. A variation of SMR is *Replicated State Machine* (RSM) employed by VCS to synchronise the resource status across all nodes [40].
3. *Application-state*, on the other hand, implies that the data of an application that a HAC protects are synchronised to one or more nodes to support a possible failover. Hence, data synchronisation in this taxonomy refers implicitly to application-state synchronisation. Such synchronisation can occur at different levels, such as on an application or a file system level.

The top-level class *data synchronisation* is shown in Figure 11. It is further divided into two storage technologies, shared storage and shared-nothing. They both can be connected to a sub-

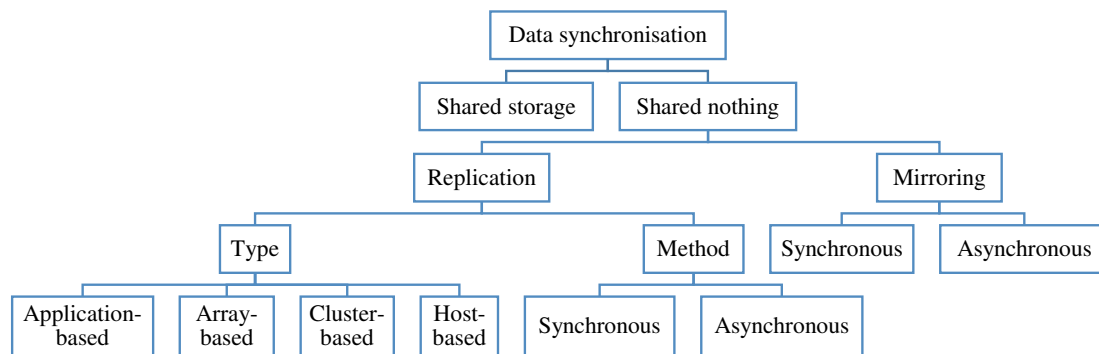


Figure 11: Data synchronisation.

class of file systems which, in turn, can influence the configuration of a HAC. An example of file systems related to shared storage is presented in Figure 12. Both cluster and distributed file systems support concurrent access and are ideal for sharing data between multiple nodes [143]. A distributed file system can be deployed on the top of either shared storage or shared-nothing, and some file systems can be deployed on both. For example, IBM Spectrum Scale (formerly the General Parallel File System (GPFS)) file system can be deployed using both storage technologies [144].

H.1: Shared storage. HAC solutions use several forms of shared storage. However, requirements for such implementation usually come from business requirements, such as supporting DR or geographically dispersed user groups. While shared storage might be ideal for the cluster types local and campus, metro and continental clusters require a different solution due to extended distances between nodes. Shared-nothing is an option in such cases. A hybrid approach is also possible, which means that shared storage and shared-nothing can support a combination of a local cluster (or campus) and continental cluster, as discussed in the topology section.

H.2: Shared-nothing. This setup assumes that there is no shared storage. Instead, each cluster node is connected to separate storage, which could either be SAN-based or based on local storage (e.g., Direct-attached storage (DAS)) [3, 1, 4]. However, EAs must explicitly support these kinds of setups. Moreover, there are also challenges with accessing shared storage in new and emerging technologies. For example, shared storage is limited in the public cloud, fog and edge computing environments [145]; hence, it becomes difficult to set up a HAC using shared storage. In such cases, replication between the individual storage units is required, and this has led to the new term *SANless* (SAN-Less) [146]. There are two techniques associated with the synchronisation of data in a shared-nothing setup: replication and mirroring.

H.2.1: Replication. Replication describes the process of replicating from a primary node to other nodes so that data is synchronised and consistent across all participating nodes [3, 1, 4, 40]. Since there are different kinds of replications in HACs, we group them by type and method. The type describes the replication approaches, while the method specifies the execu-

tion technique.

H.2.1.1: Type. Four types of HAC replication are possible: application-based, array-based, cluster-based, and host-based.

Application-based replication is set up at an application level, and it uses replication features that are provided natively by an application [47]. One of the nodes will be active in such a setup, while other nodes will be either warm standby or hot standby. To include an application in a HAC, explicit support for the application-specific replication feature by the HAC is required. Databases employ application-based replication to synchronise with standby databases [53, 147, 69, 148]. *Array-based replication* is set up on a storage system level to enable synchronisation between two storage systems (e.g., SAN- or NAS-based) [4, 47]. Additional software may be required to facilitate array-based replication. In *cluster-based replication*, the replication functionality is within a HAC and entirely administrated by the HAC [117, 149]. It means that the solution is independent of the operating environment or any other tool; instead, it relies on a high-speed network connection. *Host-based replication* uses software tools on a host (server or nodes) to perform replication. An example is using a Linux logical volume manager (LVM) to set up replication between two logical volumes across two nodes [4, 85]. Tools that operate on an operating system level and are similar to host-based replication can also be included in this category [132]. For instance, Gómez et al. [150] use a software-based Distributed Replicated Block Device (DRBD) solution to enable replication between two volumes at a block-level in a virtual cluster setup.

H.2.1.2: Method. HACs can perform the replication synchronously or asynchronously.

Synchronous replication waits until a write is completed and an acknowledgement is received from the other replication end, guaranteeing consistency between the two replication points [4]. From a transaction viewpoint, synchronous replication can support all the ACID properties. Therefore, no data loss is usually associated with it [151, 97]. However, synchronous replication is challenging with extended distances. Nevertheless, modern techniques may offer solutions. For instance, Schmidt [1] means that connections up to a distance of 100 km can be achieved by using dark fibre. This results in latencies of 0.5 μ s, which is adequate for synchronous replication. On the other




Figure 12: File systems related to shared storage. Key: ext4 – Fourth extended file system, ZFS – Z File System, OCFS2 – Oracle Cluster File System, vxVFS – Veritas Cluster File System, NFS – Network File System, IBM Spectrum Scale – A distributed file system, formerly called the General Parallel File System (GPFS)

hand, *asynchronous* replication does not wait until the writing is completed but gets an acknowledgement as soon as data is received at the second point [4, 148]. As such, it may not comply with the ACID properties entirely, which, in turn, may result in data loss.

Different factors influence the selection of a method, and some of the critical factors are [3, 4]: the distance between two nodes; the volume of data transported between nodes; type of data; frequency (continuous or burst); business requirements, such as DR. There is a network latency recommendation for synchronous replication, as it implies real-time mirroring, while asynchronous replication does not have the same kind of rigorous requirement [4, 97].

H.2.2: Mirroring. In some cases, the terms replication and mirroring are used interchangeably. For example, a host-based mirroring of a file system can also be referred to as file system replication. However, in other cases, a few differences can be observed; for instance, mirroring may differ by not having a running instance on the standby node [4, 1]. Mirroring can be performed synchronously or asynchronously [3, 4, 1].

Synchronous mirroring ensures that the mirroring process waits until a write is completed and committed on the standby node and an acknowledgement is sent back. This method secures consistency of data between two nodes. An *asynchronous* mirroring process, on the other hand, does not wait until a write is ended on the secondary node. This approach may result in data loss when the primary node fails abruptly.

4. Survey of High-availability Clusters

4.1. Selection of HACs for the Survey

We selected the relevant HACs for our survey using a systematic approach that comprised the following six steps.

Step 1. Identification of HACs that support enterprise applications (EAs). Our survey focused on HACs that can protect EAs. However, only a limited number of HACs support EAs due to the complex composition of EAs, which are multi-tiered and multi-layered. We identified likely candidate HACs using comprehensive research reports [63, 156] and articles [55, 91, 157, 158], resulting in 23 candidate HAC solutions.

Step 2. Identification of relevant EAs and databases. In this step, we gathered information for assessing the applicability

of each HAC solution to distinct layers of enterprise applications. To this end, we used relevant research and analysis reports, e.g. [159, 160, 161], to identify the databases and EAs listed next to questions Q1 and Q2 from Table 8.

Step 3. Elimination of HACs not supported by EA vendors. In this step, we used the lists of supported HACs released regularly by enterprise application vendors, e.g. [162, 163], to check which HAC solutions are supported (and sometimes certified) by these EA vendors. “Supported” HAC solutions are solutions that fulfil the requirements of the application vendor for a specific application, with the added implication that support channels have been established between the vendors.

We assessed the candidate HAC solutions using the following criteria to narrow down the list:

1. Does the HAC solution being assessed focus on only specific IT solutions (such as HPC or Hadoop)?
2. Is the HAC solution no longer active, implying that the product lifecycle has ended or the research project that developed it has ended?
3. Is the information available to analyse the HAC solution properly insufficient?
4. Are EAs supported by the HAC solution, or is it the case that the information available cannot be used to conclude whether EAs are supported or not?

We eliminated all the candidate solutions for which one or several of these questions were answered affirmatively. As a result, six HAC candidates were removed in this step, and we proceeded with the remaining 17 candidates. We made an exception for two of the candidate HAC solutions for the reasons described below:

- OpenSAF does not provide enterprise support directly. Nevertheless, we retained OpenSAF because of its stability as a HAC [164, 165]. Besides, application support can be developed individually with OpenSAF, meaning that an OpenSAF HAC can be used to support enterprise-class databases and applications.
- Similarly, the ClusterLabs stack does not support enterprise applications on its own. However, we retained the ClusterLabs stack because it provides the core components for two other selected solutions, SUSE Linux Enterprise High Availability Extension and Red Hat High Availability Add-On. This implies that customisation and further developments are possible using it.

Table 8: Selection questions

No	Question	Evaluation Parameters
Q1	Support for enterprise class databases?	SAP ASE, DB2, HANA, Informix, MySQL, Oracle, PostgreSQL, SQL Server, Teradata
Q2	Support for EAs?	Oracle Siebel Customer relationship management (CRM), Oracle†, SAP†, Others†, WebSphere
Q3	Multi-tier support for EAs?	X-Yes, N-no, ?-no information
Q4	Enterprise support provided?	24x7x365
Q5	Application features can be supported by further developments?	X-Yes, N-no, ?-no information
Q6	Support for disaster recovery?	X-Yes, N-no, ?-no information
Q7	Support for virtualization?	X-Yes, N-no, ?-no information
Q8	Cloud support?	X-Yes, N-no, ?-no information
Q9	Support for enterprise operating environments?	AIX, HP-UX, IBM i, Linux, Solaris, Windows
Q10	Support for multiple platforms?	Power, SPARC, x86
Q11	Support for large-scale clusters?	Number of nodes
Q12	Support for multiple topologies?	Active-active, application-based, server-based, N-to-N, active-passive, N+1, N+M, N-to-1
Q13	Support for availability level?	Minimum 99.9%
Q14	Active lifecycle management?	X-Yes, N-no, ?-no information

† - any of the business suite EAs (e.g., ERP)

Table 9: Eliminated HAC solutions in the six-step approach for selecting HACs for survey

Product	Reason(s) for Elimination
Apache Mesos [152]	focus on specific IT solutions (HPC)
DxEnterprise [153]	lack of EA support; insufficient information available to evaluate the solution properly
everRun [154]	lack of EA support; insufficient information available to evaluate the solution properly
HA-OSCAR [155]	no longer active
Kimberlite [55]	no longer active
Linux FailSafe [55]	no longer active

Step 4. Retention of only HACs that support automatic failover. We used this filter to retain only the HAC solutions that support automatic failover, which is crucial for an EA to minimise downtime. All 17 candidates support automatic failover; hence, all were retained.

Step 5. Design of additional questions for the selection and evaluation of HACs. In this step, we created the questions to evaluate the HACs. The queries reflected the typical requirements of EAs [166], and the objective was to select those HACs that could respond to most of the questions positively. The set of questions is listed in Table 8.

Step 6. Selection of the set of HAC solutions for the survey. We selected all the HAC solutions that can support EAs and fulfil the additional criteria from the questions Q1–Q14 in Table 8, where a positive response for any of the “evaluation parameters” from questions Q1, Q2, Q9, Q10 and Q12 was deemed sufficient to consider that a HAC solution met the criterion associated with that question. The result of the HAC selection is presented in Table 10, which comprises 17 HAC solutions for which we obtained positive responses to all queries and products, while noting the following exception:

- DR support (question Q6) for the following solution was unclear or not available: ApplicationHA, Clusterware 12c, Primecluster, RSF-1, SafeKit and WSFC.

For completeness, we also provide, in Table 9, a list of the six

HAC solutions considered initially but eliminated in Step 3 of our selection approach. For each of these solutions, Table 9 also provides a summary of the reasons for its elimination from the survey.

4.2. HAC Analysis Methodology

We used a hybrid methodology for the analysis of the 17 HAC solutions from Table 10. As a first step, we created a comprehensive spreadsheet and an online questionnaire covering our entire HAC taxonomy, which we used as a basis for the survey. In the second step, we populated the spreadsheet entries for 17 HACs by analysing product documentation, technical white papers, case studies, books, and articles as a primary analysis method. We noticed several inconsistencies between the different materials for the same edition and version of a HAC solution. To resolve these inconsistencies, we crosschecked the results by using a diverse set of materials (e.g., reference guides, technical manuals and documentation) whenever inconsistencies were observed. Finally, as a secondary method, we prepared an email that described what we were trying to achieve. We sent it to all the vendors of selected HACs, particularly to the experts responsible for the HAC products. After two weeks, we sent a reminder to those who did not reply to our original invitation; a second reminder was sent after an additional two weeks. After six weeks, we collected the data provided by the vendors and transferred it to a spreadsheet.

Table 10: Evaluation of selected HAC solutions (X – Yes; N – No; ? – No information)

Question No	ApplicationHA	Clusterware	EXPRESSCLUSTER X	InfoScale Availability	OpenSAF	ClusterLabs stack	PowerHA SystemMirror	PRIMECLUSTER	Red Hat High Availability Add-On	RSF-1	SafeKit	Serviceguard	SIOS Protection Suite	Solaris Cluster	SUSE Linux Enterprise High Availability Extension	Tivoli System Automation for Multiplatforms (SA MP)	Windows Server Failover Clustering (WSFC)
Q1	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Q2	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Q3	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Q4	X	X	X	X	N	N	X	X	X	X	X	X	X	X	X	X	X
Q5	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Q6	?	?	X	X	X	X	X	?	X	?	?	X	X	X	X	X	?
Q7	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Q8	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Q9	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Q10	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Q11	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Q12	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Q13	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Q14	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X

Despite assurance from multiple vendors, we managed to obtain a response from only one vendor, High Availability for the HAC RSF-1. Subsequently, we transferred all collected data to the spreadsheet for conducting further analysis.

4.3. HAC Survey Results

We used the taxonomy to establish the characteristics of the 17 end-to-end HAC solutions selected for the survey. The outcome of the survey is presented in Table 11, starting with general information about each HAC solution (i.e., version and vendor) in the second and third row. The remaining rows from the table present the main results of the survey, organised in the same way as our HAC taxonomy. The results are analysed in Section 4.4.

The surveyed HAC solutions usually consist of multiple editions with varying features, some of which are subject to additional licensing. Our survey covers only advanced editions that include most of the features. As even advanced editions do not support all the features when different operating systems and platforms are considered, we provide details about the limitations relating to the individual HACs where applicable (as footnotes at the end of Table 11).

As discussed, a HAC vendor may enforce further constraints by stating explicitly what version and edition of an EA are supported. Likewise, an EA vendor may list what HACs are sup-

ported by a particular EA version and edition. Many combinations of EA version, database version, HAC version and edition, operating system version, and platform make it challenging to crosscheck every single combination. Therefore, only the relevant EAs and databases are included in Table 11.

4.4. Analysis of the Survey Results

The distribution of the operating system and platform support for the surveyed HACs is shown in Figure 13 grouped by the operating system. Linux is the dominating operating system, and 15 solutions support Linux, out of which 12 support SUSE Linux on an x86-based platform, seven support SUSE Linux on Power-based platforms. Similarly, Red Hat Linux supports 13 HACs on the x86 platform and seven on the Power platforms. Oracle Linux is supported by 8 HAC solutions on x86 platforms, while only two support it on the SPARC platform. Solaris operating system is supported by seven HACs on the SPARC platform, while only four support Solaris on the x86 platform. Seven solutions support windows, and the platform is always x86. Five HACs support AIX on power, and only two HACs support HP-UX on the IA64 platform. Lastly, the rare environment is the IBM i operating system on the Power platform, which is only supported by one HAC solution.

The surveyed HAC solutions can be divided into two groups. The first group, comprising 14 of the 17 surveyed solutions, comprises the HACs marked with a star ‘*’ in Table 12.

Table 11: Outcome of the survey

Taxonomy	ApplicationHA ¹	Clusterware	EXPRESSCLUSTER X	InfoScale Availability	OpenSAF	ClusterLabs stack	PowerHA SystemMirror ¹	PRIMECLUSTER	Red Hat High Availability Add-On	RSF-1 ²³	SafeKit	Serviceguard	SIOS Protection Suite	Solaris Cluster	SUSE Linux Enterprise High Availability Extension	Twoli System Automation for Multiplatforms (SA MP)	Windows Server Failover Clustering (WSFC)
Version	6.2	12c	3.3	7.3.1	5.17.07	2.3.2	7.2.1	4.5	7.0	3.9.10	7.2	A.12.20 ²⁴	9.2	4	12	4.1	2016
Vendor	Veritas	Oracle	NEC	Veritas	SA Forum	ClusterLabs	IBM	Fujitsu	Red hat	High-Availability	Evidian	HPE	SIOS	Oracle	SuSE	IBM	Microsoft
A: Deployment Patterns²⁷																	
OS and platform																	
AIX on Power	X ₁	X ⁴	NS	X	NS	NS	X	NS	NS	NS	NS ²²	NS	NS	NS	NS	X	NS
HP-UX on IA64	NS	X	NS	NS	NS	NS	NS	NS	NS	NS	X	NS	NS	NS	NS	NS	NS
IBM i on Power	NS	NS	NS	NS	NS	NS	X	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS
Oracle Linux on SPARC	NS	NS	NS	NS	X ⁷⁸	X	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS
Oracle Linux on x86_64	X ₁	X ⁴	X	X	X	X	NS	NS	NS	X	NS	NS	X	NS	NS	NS	NS
Red Hat Enterprise Linux on Power	NS	NS	X	NS	X ⁷⁸	X	X	NS	X	X	NS	NS	NS	NS	NS	X	NS
Red Hat Enterprise Linux on x86_64	X ₁	X	X	X	X	X	NS	X	X	X	X	X	X	NS	NS	X ¹⁰	NS
Solaris on SPARC	X ₁	X	NS	X	NS	NS	NS	X	NS	X	NS	NS	NS	X	NS	X	NS
Solaris on x86_64	NS	X	NS	X	NS	NS	NS	NS	NS	X	NS	NS	NS	X ¹¹	NS	NS	NS
SUSE Linux Enterprise Server on Power	NS	NS	X	NS	X ⁷⁸	X	X	NS	NS	X	NS	NS	NS	NS	X	X	NS
SUSE Linux Enterprise Server on x86_64	X ₁	X	X	X	X	X	NS	X	NS	X	NS ²²	X	X	NS	X	X ¹⁰	NS
Windows	X ₁	X	X	X	NS	NS	NS	NS	NS	NS	X	NS	X	NS	NS	NS	X
Support for virtualized environments	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Supported virtual solutions (E-Xen, H-Hyper-V, K-KVM, O-Others, V-VMware)	E, H, K, O, V	O	E, H, V	H, K, O, V	K, O, V, X	E, H, K, V	O ⁹	K, O, V	K	?	H	H, K, V	E, H, K, V	O	E, K	K, O, V	H, V
Maximum number of nodes per cluster	?	64 ₅	32	128	100	32	16	16	16	64	?	16/32 ₂₅	32	8/16 ₂₆	32	32/130 ¹⁴	64
B: Application Areas (EA category, B-Business-critical, T-telecom (carrier-grade))	B	B	B	B	T, B	B	B	B	B	B	B	B	B	B	B	B	B
C: Type of cluster																	
C.1: Local	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
C.2: Campus	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
C.3: Metro	?	X	X	X	x	X	X	NS ⁹	X	X ⁹	NS ⁹	X	X	X	X	x	X
C.4: Continental	?	NS ⁹	X	X	x	X	X	NS ⁹	X	NS	NS ⁹	X	X ⁹	X	X	x	NS ⁹
D: Topology																	
D.1: Symmetric																	
D.1.1: Active-active																	
D.1.1.1: Application-based	NS	X	X ⁷	X ⁷	X ⁷	X ⁷	X ⁷	X ⁷	X ⁷	X ⁷	X ⁷	X ⁷	X ⁷	X ⁷	X ⁷	X ⁷	X ⁷
D.1.1.2: Server-based	X	X	X	X	X ⁷	X	X	X	X	X	X	X	X	X	X	X	X
D.1.2: N-to-N	X ⁷	?	X	X	X	X	X ⁷	X ⁷	?	NS	?	X ⁷	X	X	X	?	X
D.2: Asymmetric																	
D.2.1: Active-passive	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
D.2.2: N-to-1	X ⁷	?	X	X	X	X	X	?	X	X	X	X	X	X	X	?	NS
D.2.3: N+1	?	?	X	X	X	X	X ⁷	?	?	X	?	X ⁷	X ⁷	X	X	?	NS ⁹

D.2.4: N+M	?	?	X	X	X	X	X?	?	?	X	X	X?	X?	X	X	?	NS?
E: Cluster management																	
E.1: Cluster data																	
E.1.1: Configuration (D-Disk or file share, F-File, M-memory)	D,F	F	F	F	F	F	D,F	F	F	D,F	F	F	D,F	F	F	F	F
E.1.2: Runtime (D-Disk or file share, F-File, M-memory)	D, M?	F, M?	F, M?	F, M?	F, M?	F, M	D,F, M?	F, M?	F, M?	D, M?	F, M?	F, M?	D, M?	F, M?	F, M?	F, M?	F, M?
E.2: Communication																	
E.2.1: Type																	
E.2.1.1: Heartbeat	X ¹⁶	X	X	X	NS	X	X	X	X	X	X	X	X	X	X	X	X
E.2.1.1.1: LAN-based	X	X	X	X	NS	X	X	X	X	X	X	X	X	X	X	X	X
E.2.1.1.2: Disk-based	?	NS	X	NS	NS	NS	X	NS	NS	X	NS	NS?	NS?	X	NS	X	NS?
E.2.1.2: Node																	
E.2.1.2.1: User interface	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
E.2.1.2.2: Resource management																	
E.2.1.2.2.1: Agent	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	NS	NS
E.2.1.2.2.2: Base resource	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
E.2.1.3: Cluster																	
E.2.1.3.1: Configuration	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
E.2.1.3.2: Runtime	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
E.2.2: Method																	
E.2.2.1: Multicast	X?	X	?	NS?	X	X	X	?	X	X	NS?	X	NS?	X	X	?	X?
E.2.2.1.1: Atomic	?	?	?	NS?	X	X	?	?	X	?	?	?	?	?	X	?	?
E.2.2.1.2: Virtual synchrony	NS	NS	NS	NS	?	X	NS	NS	X	NS	NS	NS	NS	NS	X	NS	NS
E.2.2.2: Broadcast	?	?	?	X	X	X	X	X	X	?	X	X	X	X	X	X	X
E.2.2.3: Unicast	?	?	?	X	X	X	X	?	X	X	X	X	?	?	X	?	?
E.2.2.4: IP socket	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS
E.3: Resource management																	
E.3.1: Type																	
E.3.1.1: Base resource	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
E.3.1.2: Application	X	X	X	X	x	X	X	X	X	X	X	X	X	X	X	X	x
E.3.1.2.1: Agent-based	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	NS	NS
E.3.1.2.1.1: Application (C-Siebel CRM, O-Oracle, S-SAP, T-Others W-WebSphere)	C, S, T	O, C, S, T	S, O, T, W	O, S	T ¹²	O, S	O, S, T, W	S, T	O, S, W	O, S, T	T	O, S	S, W	C, O, S, W	O, S, W	S	C, O, S, T, W ¹⁵
E.3.1.2.1.2: Database (A-SAP ASE, D-DB2, H-HANA, I-Informix, M-MySQL, O-Oracle, P-PostgreSQL, S-SQL Server, T-Teradata)	D, M, O, S	M, O	A, D, H, M, O, P, S	A, D, H, O, S	M ¹²	D, M, O, P	D, H, O	O	A, D, M, O, P	A, D, I, M, O, P	O, M, P, S	A, D, H, O, P	A, D, I, M, O, P, S	A, M, P, O	D, H, I, M, O, P	D, H, O	A, D, M, O, P, S ¹⁵
E.3.2: Method																	
E.3.2.1: Policy-based	X	X	X	X	X	X	X	X	X	NS	X	X	X	X	X	X	X
E.3.2.2: Rule-based	?	?	?	?	?	X	?	?	?	X	?	?	?	?	?	?	X?
F: Failure detection and recovery																	
F.1: Monitoring	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
F.1.1: Area																	
F.1.1.1: Server	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
F.1.1.2: Cluster	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
F.1.1.3: Application	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X ⁸
F.1.2: Type																	
F.1.2.1: State-based	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
F.1.2.2: Threshold-based	?	X	X	X	X	NS?	NS?	X	NS	X	NS?	?	X	?	NS?	NS?	X
F.1.3: Method																	
F.1.3.1: Poll	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
F.1.3.2: Push	NS	NS	?	NS	?	NS	?	?	NS?	NS	NS	?	NS?	NS	NS	NS?	?
F.1.3.3: Event-based	X ²	NS?	NS?	X	?	NS	NS	X	NS	NS	NS	NS	NS	NS	NS	NS	NS
F.2: Failover																	
F.2.1: Reactive	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
F.2.2: Proactive	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS
F.3: Recovery level																	
F.3.1: Resource	X	X	X	X	X	X	X	X	X	?	?	X	X	X	X	X	X
F.3.2: Group	X	X	X	X	X	X	X	X	X	?	X	X	X	X	X	X	X
F.3.3: Node	X	X?	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
F.4: Prediction	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS	NS
F.5: Simulation	NS?	X	X	X	X	X	X	NS?	X	NS	NS?	X	NS?	NS?	X	X	NS?
G: Consistency and integrity																	
G.1: Cluster lock	NS	?	X	NS	X	NS	NS	X	X	?	?	X	X	NS	NS	X	NS
G.2: Quorum	?	X	NS	X	NS	X	X	X ²⁰	X	NS	NS	X	X	X	X	X	X
G.2.1: Device realisation																	
G.2.1.1: Server	?	NS	NS	X	NS	X	NS	X	X	NS	NS	X	X	X	X	X	X

G.2.1.2: Disk	?	X	NS	NS	NS	NS?	X	NS	X	NS	NS	NS	NS	X	X	X	X
G.2.1.3: File share	NS	NS	NS	NS	NS	NS	X ²¹	NS	NS	NS	NS	NS	NS	X	NS	X	X
G.2.2: Mode																	
G.2.2.1: Server	?	NS	NS	X	NS	X	NS	X	X	NS	NS	X	X	X	X	X	X
G.2.2.2: Node	NS	NS	NS	?	NS	?	?	?	?	NS	NS	?	?	?	?	?	X
G.2.2.3: Disk	?	X	NS	NS	NS	NS?	X	NS	X	NS	NS	NS	NS	X	X	X	X
G.2.2.4: File share	NS	NS	NS	NS	NS	NS	X	NS	NS	NS	NS	NS	NS	X	NS	X	X
G.3: Dynamic quorum	?	?	NS	X	NS	X	X	X	X	NS	NS	X	?	X	X	X	X
G.4: Isolation																	
G.4.1: Fencing	NS	X	NS	X	NS	X	X	X	X	X	NS	X	X	X	X	?	NS
G.4.1.1: Resource	NS	NS?	NS	X	NS	X	X	X	X	X	NS	X	X	X	X	?	NS?
G.4.1.2: Node	NS	X	NS	X	NS	X	?	?	X	X	X ¹⁹	?	X	?	X	?	NS
G.4.2: Shutdown	NS	X	X	NS	X	x ¹³	X	X	X	X	NS	X	X	X?	NS?	X	NS
H: Data synchronization																	
H.1: Shared storage	X	X	X	X	X	X	X	X	X	X	?	X	X	X	X	X	X
H.2: Shared-nothing	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
H.2.1: Replication	X	X	x	X	X	X	X	X	X	X	X	X	X	X	X	X	x
H.2.1.1: Type																	
H.2.1.1.1: Application-based	X ⁶	X ⁶	X ⁶	X ⁶	?	X ⁶	X ⁶	X ⁶	X ⁶	X ⁶	X ⁶	X ⁶	X ⁶	X ⁶	X ⁶	X ⁶	X ⁶
H.2.1.1.2: Array-based	x	x	x	x	x	x	x	x	x	?	?	x	x	x	x	?	x
H.2.1.1.3: Cluster-based	NS	NS	X	NS	NS	NS	NS	NS	NS	NS	X	NS	X ¹⁸	NS	NS	NS	NS
H.2.1.1.4: Host-based	x ¹⁷	x	x	X	x	X	X	X	X	?	NS	X	X	X	X	?	x
H.2.1.2: Method																	
H.2.1.2.1: Synchronous	X	X	X	X	X	X	X	X	X	?	X	X	X	X	X	X	X
H.2.1.2.2: Asynchronous	X	X	X	X	X	X	X	X	X	?	X	X	X	X	X	X	X
H.2.2: Mirroring	x	x	x	x	?	x	x	x	x	x	x	x	x	x	x	?	x
H.2.2.1: Synchronous	x	x	x	x	?	x	x	x	x	x	x	x	x	x	x	?	x
H.2.2.2: Asynchronous	x	x	x	x	?	x	x	x	x	x	x	x	x	x	x	?	x

Key: ? – No information; X – Supported; NS – Not supported; X? – Supported (not explicitly stated in the documentation, but this interpretation has been made by analysing the documentation); x – Supported together with additional components, and an example is replication support by the operating system volume manager; NS? – Not supported (not explicitly stated in the documentation, but this interpretation has been made by analysing the documentation).

- ¹ Supported only on virtualized environments.
- ² Intelligent monitoring framework.
- ³ Replication or mirroring support by additional tools is included.
- ⁴ Supported on both virtual and physical environments.
- ⁵ 64 nodes are supported for the hub, while leaf nodes can support many more.
- ⁶ Replication is provided natively by an application, but a HAC must support the feature.
- ⁷ If an application supports parallel deployments.
- ⁸ OpenSAF provides a generic development package; it can be ported to other UNIX and Linux flavours.
- ⁹ LPARs: 2 logical partitions (LPARs) on IBM PowerVM
- ¹⁰ Supported on System x hardware that is based on the x86 platform.
- ¹¹ Supported only on Oracle's x86 platforms.
- ¹² The implementer can develop application support.
- ¹³ Fencing by STONITH (Shoot the Other Node in the Head).
- ¹⁴ The maximum number of nodes on Linux is 32, and, for AIX, it is 130.
- ¹⁵ Application vendors provide application support for WSFC.
- ¹⁶ Usually, guest heartbeat is passed to a host.
- ¹⁷ Replication features of a virtual machine can also be used.
- ¹⁸ Replication feature is provided by the product DataKeeper, which is part of the SIOS Protection Suite.
- ¹⁹ Fencing as a concept is not employed, but, instead, the node with the problem is put into a waiting state..
- ²⁰ The solution uses a quorum technique called cluster integrity.
- ²¹ Implies repository disk.
- ²² Supported by SafeKit 7.1.3.
- ²³ The vendor provided most details.
- ²⁴ Version for Linux. Current version for HP-UX is A.11.20.
- ²⁵ The maximum number of supported nodes for Linux is 32, while for HP-UX, it is 16.
- ²⁶ The maximum number of supported nodes on Solaris on x86 is 8, and Solaris on SPARC supports 16.
- ²⁷ None of the surveyed HACs support container-based technologies.

Note: Although the operating system version is not stated, it is the most recent version at the time this survey was carried out (The survey was conducted mainly between October 2018 and June 2019. The survey was updated during July 2021, September 2021 and then October 2021).

Each of these HACs functions as middleware, which means that it creates an additional layer on the top of an operating environment. The HACs from the second group, which comprises the remaining three solutions, are tightly integrated with an operating system and make use of the features that are offered by an operating environment. The latter type of HAC functions as part of an operating environment, operating in the kernel mode and directly interacting with operating system functionalities. While such features can make a HAC more efficient, they may also create problems with modularity and portability, and that is why, for example, such HAC solutions only support specific

operating systems. Furthermore, the lifecycle management of such a HAC solution also becomes the operating system's lifecycle management. WSFC has already embraced this approach, and it is entirely integrated with the Windows server operating system.

Major software and hardware vendors have their HAC solutions. However, some of them are supported only by the operating environment and platform from the vendor. An example of this is WFCS, which is only available on the Windows server enterprise edition. On the other hand, some independent ven-




Figure 13: Platform and operating system support of the surveyed high availability clusters (HACs) grouped by operating system.

Table 12: The surveyed HACs, versions and vendors

Surveyed HAC	Vendor	Surveyed HAC	Vendor
ApplicationHA 6.2*	Veritas	RSF-1 3.9.10*	High-Availability
ClusterLabs stack 2.3.2*	ClusterLabs	SafeKit 7.2*	Evidian
Clusterware 12c*	Oracle	Serviceguard A.12.20*	HPE
EXPRESSCLUSTER X 3.3*	NEC	SIOS Protection Suite 9.2*	SIOS
InfoScale Availability 7.3.1*	Veritas	Solaris Cluster 4	Oracle
OpenSAF 5.17.07*	SA Forum	SUSE Linux Enterprise High Availability Extension 12*	SUSE
PowerHA SystemMirror 7.2.1	IBM	Tivoli System Automation for Multiplatforms (SA MP) 4.1*	IBM
PRIMECLUSTER 4.5*	Fujitsu	Windows Server Failover Clustering (WSFC) 2016	Microsoft
Red Hat High Availability Add-On 7.0*	Red hat		

* = solution that functions as middleware

dors specialise in HAC products, and these vendors can support multiple operating systems and platform combinations. Typically, such HACs belong to the middleware group.

Cloud deployment has also come to play an important role. In the early days of cloud computing, a separate development of HAC was considered. This led to the development of specific HAC solutions, such as ApplicationHA by Veritas and vSphere App HA by VMware. However, a better approach is to port existing solutions to the cloud environment, which made developing cloud-specific HAC solutions unnecessary. An such, solutions like App HA by VMware were discontinued. However, HACs in the public cloud comes with limitations. For example, using shared storage is a challenge. On the other hand, this has contributed to developing enhancements to enable deploying HACs in the cloud. One such enhancement is the so-called

storage-less or *SANless* HAC, which allows HACs to operate without shared storage. Moreover, the transition to cloud services models, such as SaaS, PaaS, and IaaS, changes the way HACs are deployed and managed. Likewise, roles and responsibilities for managing a HAC with the different service models also change.

Furthermore, the introduction of *multi-clouds* can also complicate a HAC deployment, not least from a roles and responsibilities perspective. Somasekaram highlights the issues with roles and responsibilities of HA and DR solutions in the context of outsourcing [22]. He argues that the issues are valid even for the cloud environment because the cloud is regarded as outsourcing, and cloud providers are usually responsible for multiple layers (e.g., network and storage). At the same time, other suppliers manage the rest of the layers.

Similar to the challenges described for the deployment environment public cloud and host virtual, the emerging deployment environments fog and edge also face challenges [167, 79]. When used with the host virtual, the challenges are the same as the public cloud. On the other hand, when the host container is used in all deployment environments, ensuring high availability for stateful applications hosted in containers becomes a challenge. Containers run as a process in user space, and this may restrict the implementation of HAC features that require running in kernel space [168]. Moreover, containers typically support a single application or a service in a container which means a HAC cannot deploy agents in the same container to manage the application resources [169]. To overcome this limitation, the commonly implemented container orchestration system Kubernetes provides a sidecar option (i.e., a separate container), which enables deployments of HAC-related components there. The sidecar container runs along with the container that hosts the application. Using this approach, commercial vendors started providing HACs for containers. InfoScale availability (formerly Cluster Server—VCS) for containers from Veritas is one such HAC for Kubernetes. This HAC provides monitoring, integrated I/O fencing, arbitration and shared storage using a Container Storage Interface (CSI) plugin [169] to ensure that the HAC can deliver HA for the application. The solution requires at least two private networks to enable cluster communication and one public network to facilitate heartbeat communication. There are also research projects that explore the use of existing HACs such as Pacemaker/Corosync [170] and OpenSAF [171] to support container-based applications.

Latency over long distances has traditionally been a major problem for HACs. However, the technology has evolved and techniques are currently available to reduce latency considerably, enabling the setting up of HACs across substantial distances. *Atomic broadcast* and *multicast* (total order messaging) are often associated with fault-tolerance in distributed systems; hence, there are persuasive arguments to employ it even for HAC communication [64]. However, it is only employed by some of the HACs today.

On the whole, *prediction* is absent from the surveyed HAC solutions. Most solutions employ a poll-based monitoring mechanism that is often state-based, meaning that only the states of the resources are monitored. Moving towards *industry-standards* has also been observed in some areas, such as when using the IPMI to shut down nodes as part of isolating a problematical node. The *SCSI-3* interface is widely employed to isolate on a resource level, and often as part of *fencing*. The *quorum* concept is commonly employed so that a cluster can take action upon a situation that leads to the partitioning of a cluster.

In conclusion, the current HAC solutions for EAs are dominated by commercial vendors (15 out of the 17 surveyed solutions). This is unsurprising because customers look for HAC solutions for their business-critical applications, and, as such, proper support is paramount. However, this also means that the vendors conduct most of the research. There are, however, some open-source initiatives, and two active initiatives are

OpenSAF and ClusterLabs stack (Pacemaker/Corosync). The open-source initiatives often focus on Linux, and there have been different projects to develop a consistent HAC solution for Linux. While such efforts have been split into other projects or discontinued, some of the components are still active, and the current open-source cluster solutions are a combination of various initiatives. The main components of the current setup of the ClusterLabs stack are Corosync, Pacemaker, DRBD, STONITH, and a diverse range of application agents, which are packaged under a ClusterLabs stack. OpenAIS was an initiative to support implementing Application Interface Specification (AIS) developed by the Service Availability Forum (SA Forum), and Corosync originated from that initiative. Pacemaker is a cluster resource manager (CRM) tool that originates from the Linux-HA project.

Application agents follow the standard API established by the Open Cluster Framework (OCF), which helps standardise the application resource management. Both SuSe Linux Enterprise HA Extension and Red Hat Enterprise Linux HA add-on use Pacemaker, Corosync, the OCF concept and many other open-source components. OpenSAF, on the other hand, focuses on the telecommunications sector, where there is a need to support very high availability for carrier-grade servers that operate in the telecommunication infrastructure. However, there have been initiatives to deploy OpenSAF in a range of environments, such as the cloud. For example, Kanso et al. [164] proposed an OpenSAF based deployment in the cloud, but it too focuses on telecom applications. The challenge with open-source initiatives is to secure proper support, which is crucial for EAs. On the other hand, Red Hat and SuSe provide such support even though they have developed their HACs using mainly open-source components. It must be noted that there have been several projects related to the development of HACs, both commercial and open-source, over the years. However, many of them are no longer active, and examples include FailSafe by Silicon Graphics (SG) in the commercial area, while HA-OSCAR represents an open-source equivalent.

5. Future Directions

We have identified several limitations, challenges and opportunities as part of constructing the HAC taxonomy from Section 3 and conducting the survey from Section 4. The limitations and challenges are from an implementation perspective and an operations viewpoint, while opportunities can improve the overall HAC solutions. Using the identified limitations, challenges and opportunities, we discuss future research directions.

5.1. Limitations

The HAC limitations presented in this section apply to a majority of the HAC solutions that we have studied, with limitations L1, L5, and L7 common for all solutions.

L1. Standardisation. Standardisation of the HA domain, its components, and related processes is missing. For example, the

terminology used by HAC solutions differs considerably. Standardisation could improve research approaches and could enable better discussions and research quality. Furthermore, the lack of standardisation makes it challenging to develop standard APIs that can function with multiple solutions to support specific functionalities, for instance, application-specific agent development. We have addressed this lack of standard terminology using consistent terminology while constructing the taxonomy and performing the subsequent survey.

L2. Virtual environments. The separation between host and guest in virtualised environments complicates some of the functionalities of HACs, such as coordinated monitoring of two operating environments, guest and host, which must be correlated when hosting a critical application. If such a setup is not in place, a guest HAC may not be aware of the host at all. If there are problems in the host which impact all the guests hosted there, the guest HAC may not be able to recognise the problems [172], which could potentially impact the application. Likewise, if the guest application experiences problems, the host may not react since it is unaware of any issues except when hardware resource utilisation significantly increases. Kanso et al. [164] highlighted the problem with a guest HAC that is not aware of the host environment. Some HAC solutions promote a solution by running additional components on the host that also interact with the guest HAC. However, there is no uniformity for deploying such components because they may differ based on the virtual environment, such as VMWare or kernel-based VM (KVM). In KVM, additional tools are typically required on the host, while VMWare comes with a set of accessories that can be used instead so that no additional means are required. Some HAC solutions, such as ApplicationHA on KVM, employ a separate HAC installation on a host machine. For example, ApplicationHA on the guest can interact with the host HAC. This setup can support monitoring of the host and enable the use of features that are not otherwise available in the guest environment due to restrictions. However, a heterogeneous virtual environment with different operating environments for hosts and guests may also complicate the cross-deployment of a HAC as each operating system, platform, and virtual environment comes with restrictions.

L3. Cloud environment limitations. Both private and public clouds come with limitations. In such a cloud environment, particularly in an IaaS model, customers have access to a guest environment (e.g., VMs). To support an EA, a HAC will require access to some host elements well. In addition, the host environment must be monitored as well as part of a holistic HAC approach, which may mean deploying additional tools, as described in L2, on the host. The limitations of a cloud environment may require changes in the architecture of the HAC, hence also the protected application [17].

L4. Public cloud limitations. In addition to what is described in L2 and L3, the public cloud has some additional infrastructure-related restrictions, which are usually different from those of a private cloud. For example, shared storage is not typically supported [31]. Hence any shared-storage-based HAC must find

an alternative solution that implies that shared-disk-based quorums cannot be employed. Further, there could be additional restrictions impacting the core functionality of a HAC, such as multicast or broadcast communication not being allowed [31], which would impact the HAC's ability to communicate. Again, this means alternative solutions must be identified and implemented by adding new tools and procedures, which may, in turn, add more complexity to a solution. Moreover, if an application is deployed in a virtual environment, additional restrictions, described in L2, apply. For instance, the deployment of additional HAC tools on a host, as explained in L2, is usually not possible as hosts are managed entirely by cloud providers in such settings (L3).

L5. Rating of errors. Often a severity rating is not used for errors on a resource level, which means that all errors are treated equally. Adding severity levels would help distinguish between the different types of errors and by the different modules of HACs (e.g., monitoring, failure management) so that actions can be taken accordingly. In addition, multilevel severity would help to improve the recovery process so that, in some cases, errors can be disregarded, indicating that such errors do not result in a complete failover.

L6. Standardisation of error, failure, and event message representation. The current approach is very much individualised to different HAC solutions, implying no standard structure for log messages. This makes it hard to develop a general solution to analyse log messages (e.g., for analytical purposes). Furthermore, several modules of a HAC (e.g., monitoring, failure management) may write the same error messages with the same timestamp when a resource fails, making it challenging to mine the log for distinct error messages. Moreover, log sources can also vary as some HAC solutions may employ more than one log source. The difference presents a challenge in mining data from log sources, as it will require one or more data extraction interfaces for each HAC solution.

L7. Rating of resource and resource group dependencies. Resource and resource group dependencies are not always rated, which means that the same failover and recovery policies are applied to all dependencies, regardless of the strength of the dependencies. The dependency rating describes how the failure of a resource can impact another resource through a dependency connection and on what level, which can ultimately influence the mitigation action.

L8. Application monitoring. Even though many of the current HAC solutions employ application monitoring, application-specific errors (e.g., hang situations), are not usually captured. Furthermore, application-related errors are often difficult to monitor with a HAC. This may require additional modules and steps, such as logging in to an application, to detect such failure. Hence, the current situation is that an application may be completely unresponsive, yet it is still regarded as running by a HAC. Therefore, such errors do not trigger any action until the problem is reported by the users of the application. Consequently, this will also result in incorrect values for MTTR and MTBF since no accurate time of failure is available, thus

providing unreliable figures for availability.

5.2. Challenges

Challenges are associated with functionalities or features that can be implemented to improve the effectiveness of HACs, but that are difficult to realise due to limitations and other constraints.

C1. Roles and responsibilities. HACs work closely with operating environments and infrastructure components to provide the required HAC functionalities, such as heartbeat, monitoring, fencing, and quorum. While the roles and responsibilities of the experts in charge of setting up a HAC change with the different cloud service models [173], it is unlikely that one team can manage a complete HAC implementation and operations. Instead, multiple teams and even organisations must work together to support HAC implementation and operations. Moreover, a heterogeneous virtual environment further complicates the setup because at least two operating systems will be associated with host and guest, which means different teams are usually designated to support the host and guest environments. This means that there must be a support process that links all the different teams together according to a well-defined roles and responsibilities matrix. Moreover, the related support processes, for instance, change and incident management, must be designed accordingly.

C2. Lifecycle management. The combination of many application agents, HAC components, VMs, operating systems, and platforms complicates the lifecycle management of HACs. While having a standard across architecture components (including agents) can reduce the number of combinations, this is extremely difficult to achieve. In particular, in virtual environments, lifecycle management must take into account other elements, such as host and guest operating environments on various VMs, which adds further complexity, as described in **L2**. The number of combinations may prompt more threads of lifecycle management. For example, when an EA vendor releases an update to the application, a HAC vendor must also make sure to release an update of the HAC or agent to support the changes in the application.

C3. Client-state synchronisation. Client-state synchronisation for EAs is a difficulty. However, if achieved by a HAC, it can improve availability significantly because it can transfer user sessions in the event of a failover, which means that no user data is lost. When a failover takes place, all user input that is not saved is lost. When the failover is complete, users can log in again to establish new sessions and start their work from scratch. If an EA supports thousands of users, this means losing countless hours of work. On the other hand, if a client-state synchronisation can be achieved, it will preserve all connections and sessions, saving considerable time. It is also likely that, with faster failovers using client-state synchronisation, users will not even notice that a failover has taken place. Instead, they will be able to continue working as if nothing has happened. However, state synchronisation for an EA is a significant challenge because it requires replicating user connec-

tions, user sessions, user context, session context, user work, and global and local variables. While solutions with a limited scope, such as a firewall, widely employ client-state synchronisation, these are difficult to adopt for the much more complex settings of EAs. Since the problem is about preserving user sessions and related data, in many cases, an applications server layer may also need to be synchronised, as they are the front-ends for user communication in a multi-tier system. Furthermore, applications with a sizable workload require substantial time to stop and start application components in a specific order. Some portion of that time is consumed on ending user sessions gracefully during the stop and establishing non-user (e.g., batch) sessions at the start. However, client-state synchronisation may reduce that time significantly since user data would be already synchronised across the HAC members.

5.3. Opportunities

We have identified a set of opportunities that can improve HAC solutions considerably, typically by overcoming HAC limitations from Section 5.1. For instance, introducing probabilistic and statistical methods as detailed in opportunity **O5** below requires that ratings of errors and dependencies are in place. Hence, the exploitation of opportunity **O5** requires solutions to limitations **L5** and **L7**.

O1. Architecture components. HAC solutions employ different architecture components, and therefore having a standard and modular architecture will help standardise these components. This assumes that such an architecture will consist of standard modules, and that a HAC solution can choose to implement only a subset of modules, but it can always refer such modules to the standard modules. A set of specifications can support defining the roles of the modules and even provide means to develop interfaces (e.g., APIs). Solutions that are developed using the APIs can potentially be used with multiple HAC solutions. Moreover, the approach would aid in simplifying and interpreting architecture components while enabling the development of approaches for new and emerging technologies (e.g., containerisation), standard testing, and benchmarking.

O2. Evaluation of historical data. HACs produce a large volume of data, and such data can be invaluable when analysing past events and mitigations. These data are generated mainly through logging of events, failures, recoveries, and failovers. Historical data, together with current data, can be analysed to identify patterns, enabling proactive approaches to ensuring high availability. Therefore, evaluation of past data and current data can be used to predict failures of a repeating nature and other related failures.

O3. Reliable cluster communication protocols. The reliability of cluster communication can be increased significantly by employing protocols with atomic features such as TOTEM [104]. These features are only supported by a few HAC solutions today. Employing a standard protocol will also enhance development in the areas, as more people can be involved in

the development, which means that issues can be addressed quickly.

O4. Monitoring. Most HAC solutions use a poll-based monitoring method, which is linked to performance problems [40]. If the polling frequency increases, it will improve the monitoring data quality because more up-to-date data will then be available. However, there is an additional overhead associated with frequent polling of many resources, which could be resource-demanding. Furthermore, detecting application-specific errors might also present some challenges as described in **L8**. The monitoring functionality of a HAC may not detect such cases since HACs often focus on monitoring state changes of a resource or a resource group. Therefore, relevant monitoring models should be evaluated to improve the data quality while reducing the performance overhead. The current monitoring type is mostly state-based. However, a different option might be to use a standard API to interact with the operating environments so that the enhanced monitoring features of the operating systems can be utilised. Though this approach may still require an application-specific development, it can be simplified by using standard APIs, as discussed in **L1** and **O1**.

O5. Incorporation of probabilistic and statistical methods. Such methods are not employed currently, but they can improve effectiveness significantly and reduce downtime by analysing data, checking behaviours and providing predictions. In addition, such methods will also improve the quality of the service for HACs and their components, in general, and promote a more robust proactive approach than currently employed by mostly reactive mechanisms. One example of such an improvement is introducing statistical analysis to enable the management of quorum services more intelligently.

O6. Analytical services. Analytical services will help identify patterns in the behaviour of HACs and their components while also providing a consolidated view of total downtime and causes. Analytical services can also incorporate data from multiple sources so that data can be combined to provide reliable analysis and even produce predictions on potential failures. An example is that if some HAC components manifest intermittent failures before complete failure, patterns can be analysed to estimate the subsequent failure or an eventual complete failure.

O7. Benchmark. A standard benchmarking approach that can measure availability at a granular level will improve the performance measurements of HACs, while also enabling more a natural comparison between different solutions.

O8. Security. *HAC security* is a rarely concern. However, unauthorised access to the services of a HAC means effectively that the protected application is also jeopardised because a HAC has typically complete control of the operations of the protected application. Security is of particular concern in cloud environments with shared responsibilities (**C1**), since multiple teams assume responsibility for the different layers, which may present new vulnerabilities without a proper security model in place. Moreover, operating a HAC solution in a public cloud

may also introduce new vulnerabilities [174], mainly when new and alternative solutions must be introduced due to restrictions, as described in **L4**.

6. Conclusions

In this article, we presented a comprehensive taxonomy and a two-part survey of high-availability clusters. The first part of the survey, delivered while describing the elements of the taxonomy, provides an overview of the HAC research landscape. The second part employs the taxonomy to survey end-to-end HAC solutions developed to support enterprise applications. Finally, we detailed HAC limitations, challenges and opportunities identified while constructing the taxonomy and conducting the two-part survey. Using these, we discuss future research directions for high-availability clusters. In particular, an adaption of fully functional HACs for cloud-deployed enterprise applications can significantly improve the availability of these applications. Similarly, exploiting historical data through the use of probabilistic approaches to predicting future failures and other relevant events can improve the effectiveness of HACs. Last but not least, HAC support for client-state synchronisation has the potential to deliver zero downtime for an important range of failures affecting enterprise applications.

References

- [1] K. Schmidt, High availability and disaster recovery: concepts, design, implementation, Vol. 22, Springer Science & Business Media, 2006.
- [2] T. Liu, H. Song, et al., Availability prediction and modeling of high mobility OSCAR cluster, in: 2003 Proceedings IEEE International Conference on Cluster Computing, IEEE, 2003, pp. 380–386.
- [3] E. Marcus, H. Stern, Blueprints for high availability, John Wiley & Sons, Indianapolis, Indiana, 2003.
- [4] T. Critchley, High availability IT services, Auerbach Publications, 2014.
- [5] Y. Mansouri, A. N. Toosi, R. Buyya, Data storage management in cloud environments: Taxonomy, survey, and future directions, ACM Computing Surveys (CSUR) 50 (6) (2018) 91.
- [6] C. Vercellis, Business intelligence: data mining and optimization for decision making, John Wiley & Sons, 2011.
- [7] Y. Demchenko, C. De Laat, P. Membrey, Defining architecture components of the big data ecosystem, in: Collaboration Technologies and Systems (CTS), 2014 International Conference on, IEEE, 2014, pp. 104–112.
- [8] H. Hu, Y. Wen, T.-S. Chua, X. Li, Toward scalable systems for big data analytics: A technology tutorial, IEEE access 2 (2014) 652–687.
- [9] D. M. Ranade, Shared Data Clusters: Scaleable, Manageable, and Highly Available Systems (Veritas Series), Vol. 9, John Wiley & Sons, 2003.
- [10] W. Vogels, D. Dumitriu, K. Birman, R. Gamache, M. Massa, R. Short, J. Vert, J. Barrera, J. Gray, The design and architecture of the microsoft cluster service—a practical approach to high-availability and scalability, in: Fault-Tolerant Computing, 1998. Digest of Papers. Twenty-Eighth Annual International Symposium on, IEEE, 1998, pp. 422–431.
- [11] K. P. Birman, Guide to Reliable Distributed Systems: Building High-Assurance Applications and Cloud-Hosted Services, Springer Science & Business Media, 2012.
- [12] K. P. Murphy, Machine Learning: A Probabilistic Perspective, The MIT Press, 2012.
- [13] X. Gu, H. Wang, Online anomaly prediction for robust cluster systems, in: IEEE International Conference on Data Engineering, IEEE, 2009, pp. 1000–1011.
- [14] R. Calinescu, D. Weyns, S. Gerasimou, M. U. Iftikhar, I. Habli, T. Kelly, Engineering trustworthy self-adaptive software with dynamic assurance cases, IEEE Trans. Software Eng. 44 (11) (2018) 1039–1069.

- [15] P. T. Endo, M. Rodrigues, G. E. Gonçalves, J. Kelner, D. H. Sadok, C. Curescu, High availability in clouds: systematic review and research challenges, *Journal of Cloud Computing* 5 (1) (2016) 16.
- [16] W. Li, A. Kanso, A. Gherbi, Leveraging linux containers to achieve high availability for cloud services, in: *Cloud Engineering (IC2E)*, 2015 IEEE International Conference on, IEEE, 2015, pp. 76–83.
- [17] M. Nabi, M. Toeroe, F. Khendek, Availability in the cloud: State of the art, *Journal of Network and Computer Applications* 60 (2016) 54–67.
- [18] A. Avizienis, J.-C. Laprie, B. Randell, C. Landwehr, Basic concepts and taxonomy of dependable and secure computing, *IEEE transactions on dependable and secure computing* 1 (1) (2004) 11–33.
- [19] ISO/IEC 25010:2011, *Systems and software engineering — Systems and software Quality Requirements and Evaluation (SQuaRE) — System and software quality models*, 2011.
- [20] I. Koren, C. M. Krishna, *Fault-Tolerant Systems*, 1st Edition, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2007.
- [21] P. P. O'Connor, A. Kleyner, *Practical reliability engineering.*, 5th Edition, John Wiley & Sons, Ltd, New York, 2012.
- [22] P. Somasekaram, *A Component-based Business Continuity and Disaster Recovery Framework*, Master's thesis, Uppsala university (jun 2017).
URL <http://uu.diva-portal.org/smash/get/diva2:1108197/FULLTEXT01.pdf>
- [23] P. N. Ayuso, R. M. Gasca, L. Lefèvre, Demystifying cluster-based fault-tolerant firewalls, *IEEE Internet Computing* 13 (6) (2009) 31–38.
- [24] Check Point Software Technologies Ltd, *ClusterXL Administration Guide R80.10* (2018).
URL <https://downloads.checkpoint.com/dc/download.htm?ID=54804>
- [25] M. Bajohr, T. Margaria, High service availability in MaTRICS for the OCS, in: *International Symposium On Leveraging Applications of Formal Methods, Verification and Validation*, Springer, 2008, pp. 572–586.
- [26] L. G. Zhu, D. Z. Han, S. Z. Zhou, C. S. Xie, High availability cluster with combining nas and ISCSI, *Proceedings of the 2006 International Conference on Machine Learning and Cybernetics 2006 (August) (2006) 4455–4460*.
- [27] G. L. Santos, P. T. Endo, G. Goncalves, D. Rosendo, D. Gomes, J. Kelner, D. Sadok, M. Mahloo, Analyzing the IT subsystem failure impact on availability of cloud services, in: *Proceedings - IEEE Symposium on Computers and Communications*, 2017, pp. 717–723.
- [28] D. A. B. Fernandes, L. F. B. Soares, J. V. Gomes, M. M. Freire, P. R. M. Inácio, Security issues in cloud environments: a survey, *International Journal of Information Security* 13 (2) (2014) 113–170.
URL <http://link.springer.com/10.1007/s10207-013-0208-7>
- [29] L. A. Barroso, U. Hölzle, The datacenter as a computer: An introduction to the design of warehouse-scale machines, *Synthesis lectures on computer architecture* 4 (1) (2009) 1–108.
- [30] R. Dukaric, M. B. Juric, Towards a unified taxonomy and architecture of cloud frameworks, *Future Generation Computer Systems* 29 (5) (2013) 1196–1210.
- [31] Amazon Web Services Inc, *SAP on Amazon Web Services High Availability Guide* (2016).
URL <https://d0.awsstatic.com/enterprise-marketing/SAP/sap-on-aws-high-availability-guide.pdf>
- [32] H. Wang, H. Wang, J. Shen, Architectural design and implementation of highly available and scalable medical system with ibm websphere middleware, in: *Proceedings. 17th IEEE Symposium on Computer-Based Medical Systems*, IEEE, 2004, pp. 174–179.
- [33] Z. Wen, Y. Liang, G. Li, Design and implementation of high-availability PaaS platform based on virtualization platform, in: *2020 IEEE 5th Information Technology and Mechatronics Engineering Conference (ITOEC)*, IEEE, 2020, pp. 1571–1575.
- [34] R. Buyya, et al., *High performance cluster computing: Architectures and systems (volume 1)*, Prentice Hall, Upper SaddleRiver, NJ, USA 1 (1999) 999.
- [35] SAS Institute Inc., *SAS 9.4 Intelligence Platform: Overview*, Second Edition (2017).
URL <http://documentation.sas.com/api/docsets/biov/9.4/content/biov.pdf>
- [36] Oracle Corporation, *Operating SAP Landscapes on Oracle Engineered Systems Using ITIL Best Practices* (2013).
- [37] Oracle Corporation, *Oracle Fusion Middleware: High Availability Guide* (2017).
URL <https://docs.oracle.com/middleware/12212/lcm/ASHIA/title.htm>
- [38] H. Lyu, P. Li, R. Yan, H. Qian, B. Sheng, High-availability deployment for large enterprises, in: *2016 International Conference on Progress in Informatics and Computing (PIC)*, IEEE, 2016, pp. 503–507.
- [39] Oracle Corporation, *Solaris Cluster Data Service for Oracle E-Business Suite as of Release 12.2 Guide Oracle Solaris Cluster Data Service for Oracle E-Business Suite as of Release 12.2 Guide* (2016).
URL https://docs.oracle.com/cd/E56676_01/pdf/E60641.pdf
- [40] Veritas Technologies LLC, *Cluster Server 7.3 Administrator's Guide - Linux* (2017).
URL https://origin-download.veritas.com/resources/content/live/DOCUMENTATION/SFDC/000126860/en_US/vcs_admin_73_lin.pdf
- [41] S. Bartkowski, C. De Buitlear, A. Kalicki, M. Loster, M. Marczewski, A. Mosaad, J. Nelken, M. Soliman, K. Subtil, M. Vrhovnik, et al., High availability and disaster recovery options for DB2 for Linux, UNIX, and Windows, *IBM Redbooks*, 2012.
- [42] U. F. Minhas, S. Rajagopalan, B. Cully, A. Abounaga, K. Salem, A. Warfield, Remusdb: Transparent high availability for database systems, *The VLDB Journal—The International Journal on Very Large Data Bases* 22 (1) (2013) 29–45.
- [43] B. Cully, G. Lefebvre, D. Meyer, M. Feeley, N. Hutchinson, A. Warfield, Remus: High availability via asynchronous virtual machine replication, in: *Proceedings of the 5th USENIX Symposium on Networked Systems Design and Implementation*, San Francisco, 2008, pp. 161–174.
- [44] VMware Inc., *SAP Solutions on VMware vSphere Guidelines Summary and Best Practices* (2015).
- [45] Palo Alto Networks Inc., *PAN-OS 8.0 Admin Guide* (2018).
URL https://www.paloaltonetworks.com/content/dam/pan/en_US/assets/pdf/technical-documentation/80/pan-os/pan-os-admin/pan-os.pdf
- [46] M. Sheghdara, J. Hassine, Automatic retrieval and analysis of high availability scenarios from system execution traces: A case study on hot standby router protocol, *Journal of Systems and Software* 161 (2020) 110490.
- [47] H. Saxena, J. Pound, A cloud-native architecture for replicated data services, in: *12th {USENIX} Workshop on Hot Topics in Cloud Computing (HotCloud 20)*, 2020.
- [48] C. Zhang, A. G. Kumbhare, I. Manousakis, D. Zhang, P. A. Misra, R. Assis, K. Woolcock, N. Mahalingam, B. Warriar, D. Gauthier, et al., Flex: High-availability datacenters with zero reserved power, in: *2021 ACM/IEEE 48th Annual International Symposium on Computer Architecture (ISCA)*, IEEE, 2021, pp. 319–332.
- [49] G. B. Heimovski, R. C. Turchetti, J. A. Wickboldt, L. Z. Granville, E. P. Duarte Jr, FT-Aurora: A highly available IaaS cloud manager based on replication, *Computer Networks* 168 (2020) 107041.
- [50] D. Rosendo, D. Gomes, G. Leoni Santos, L. Silva, A. Moreira, J. Kelner, D. Sadok, G. Gonçalves, A. Mehta, M. Wildeman, et al., Availability analysis of design configurations to compose virtual performance-optimized data center systems in next-generation cloud data centers, *Software: Practice and Experience* 50 (6) (2020) 805–826.
- [51] Z. Sun, H. Jin, J. Yong, S. Al-Ismaili, C. Li, J. Shen, A High Availability Application Service Platform for nuclear power enterprises, *2016 IEEE 20th International Conference on Computer Supported Cooperative Work in Design (CSCWD)* (2016) 613–618.
- [52] F. T. Cheng, S. L. Wu, P. Y. Tsai, Y. T. Chung, H. C. Yang, Application cluster service scheme for near-zero-downtime services, *Proceedings - IEEE International Conference on Robotics and Automation 2005 (April) (2005) 4062–4067*.
- [53] H. Xiong, F. Fowley, C. Pahl, A database-specific pattern for multi-cloud high availability and disaster recovery, in: *Communications in Computer and Information Science*, Vol. 567, 2016, pp. 374–388.
- [54] J. Riley, J. Noss, W. Dillingham, J. Cuff, I. M. Llorente, A high-availability cloud for research computing, *Computer* 50 (6) (2017) 92–95.
- [55] C. Leangsuksun, T. Liu, T. Rao, S. Scott, R. Libby, A failure predictive and policy-based high availability strategy for linux high performance computing cluster, in: *The 5th LCI International Conference on Linux Clusters: The HPC Revolution*, Citeseer, 2004, pp. 18–20.
- [56] K. Birman, R. Van Renesse, W. Vogels, Adding high availability and autonomic behavior to web services, in: *Proceedings of the 26th International Conference on Software Engineering*, IEEE Computer Society,

- 2004, pp. 17–26.
- [57] Oracle Corporation, Solaris Cluster Concepts Guide (2014).
URL https://docs.oracle.com/cd/E39579_01/pdf/E39575.pdf
- [58] SUSE LLC, Administration Guide SUSE Linux Enterprise High Availability Extension 12 SP3 (2017).
URL <https://www.suse.com/documentation/sle-ha-12/pdfdoc/sle-ha-12-sp3.zip>
- [59] D. Quintero, V. Balappa, S. Bodily, D. De, S. Casali, J.-G. Dengel, M. Dhandapani, M. Hoshino, J. Kiran, K. Langer, P. Sergio, L. Queiroz, M. Radford, A. Socoliuc, N. K. Tu, IBM PowerHA SystemMirror 7.1.2 Enterprise Edition for AIX, IBM Redbooks, 2013.
URL <https://www.redbooks.ibm.com/redbooks/pdfs/sg248106.pdf>
- [60] D. M. Ranade, Shared Data Clusters: Scaleable, Manageable, and Highly Available Systems (VERITAS Series), John Wiley & Sons, Ltd, New York, 2002.
- [61] M. Toeroe, F. Tam, Service availability: principles and practice, John Wiley & Sons, 2012.
- [62] S. Snedaker, Business Continuity and Disaster Recovery Planning for IT Professionals, 2nd Edition, Syngress Publishing, 2013.
- [63] IDC Corporation, Clustering Solutions for Achieving High Availability for Diversifying Platforms: The Future in Advanced Best Practices, Tech. rep., IDC Corporation, Massachusetts (2016).
URL https://www.nec.com/en/global/prod/expresscluster/materials/203586_EN_IDC_NEC_HA_White_Paper.pdf
- [64] Service Availability Forum, Service Availability Forum Service Availability Interface (2011).
URL <http://devel.opensaf.org/SAI-Overview-B.05.03.AL.pdf>
- [65] Veritas Technologies LLC, A guide to understanding volume replicator, Engineering White Paper, Symantec (2006).
- [66] Red Hat, Inc., Deploying Highly Available SAP NetWeaver-based Servers Using Red Hat Enterprise Linux HA add-on with Pacemaker (2014).
URL <https://access.redhat.com/sites/default/files/attachments/rh-pacemaker-sap-whitepaper.pdf>
- [67] Novell, Inc., SAP Applications Made High Available on SUSE Linux Enterprise Server 10 (2014).
URL https://www.b1-systems.de/fileadmin/content/whitepaper/Technical_Guide_SLES_HA_for_SAP.pdf
- [68] SUSE LLC, SAP on SUSE linux enterprise (2012).
URL https://www.suse.com/media/white-paper/sap_on_sle.pdf
- [69] Oracle Corporation, Oracle Database High Availability Overview (2017).
URL <https://docs.oracle.com/en/database/oracle/oracle-database/12.2/haovw/index.html>
- [70] Veritas Technologies LLC, Storage Foundation for Oracle RAC 7.3 Administrator’s Guide - Linux (2017).
URL https://origin-download.veritas.com/resources/content/live/DOCUMENTATION/SFDC/000126863/en_US/sfrac_admin_73_lin.pdf
- [71] IBM Corporation, Administering PowerHA SystemMirror (2018).
URL http://public.dhe.ibm.com/systems/power/docs/powerha/721/hacmpadmngd_pdf.pdf
- [72] Fujitsu Limited, Primecluster: Installation and Administration Guide 4.5 (2017).
URL <http://software.fujitsu.com/jp/manual/manualfiles/m170002/j2s21672/01enz200/j2s2-1672-01enz0.pdf>
- [73] Microsoft Corporation, Failover Cluster Step-by-Step Guide: Configuring the Quorum in a Failover Cluster (2011).
URL [https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2008-R2-and-2008/cc770620\(v=ws.10\)](https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2008-R2-and-2008/cc770620(v=ws.10))
- [74] Microsoft Corporation, High availability for SAP NetWeaver on Azure VMs on SUSE Linux Enterprise Server for SAP applications (2020).
URL <https://docs.microsoft.com/en-us/azure/virtual-machines/workloads/sap/high-availability-guide-suse>
- [75] Microsoft Corporation, Cluster an SAP ASCS/SCS instance on a Windows failover cluster using a cluster shared disk in Azure (2017).
URL <https://docs.microsoft.com/en-us/azure/virtual-machines/workloads/sap/sap-high-availability-guide-wsfc-shared-disk>
- [76] F. Liu, J. Tong, J. Mao, R. Bohn, J. Messina, L. Badger, D. Leaf, NIST cloud computing reference architecture, NIST special publication 500 (2011) (2011) 1–28.
- [77] Microsoft Corporation, Shared Responsibilities for Cloud Computing (2017).
URL <https://gallery.technet.microsoft.com/Shared-Responsibilities-81d0ff91>
- [78] Amazon Web Services Inc., Shared Responsibility Model (2018).
URL <https://aws.amazon.com/compliance/shared-responsibility-model/>
- [79] A. Yousefpour, C. Fung, T. Nguyen, K. Kadiyala, F. Jalali, A. Niakanlahiji, J. Kong, J. P. Jue, All one needs to know about fog computing and related edge computing paradigms: A complete survey, *Journal of Systems Architecture* 98 (2019) 289–330.
- [80] T. Qiu, J. Chi, X. Zhou, Z. Ning, M. Atiquzzaman, D. O. Wu, Edge computing in industrial internet of things: Architecture, advances and challenges, *IEEE Communications Surveys & Tutorials* 22 (4) (2020) 2462–2488.
- [81] F. Magnanini, L. Ferretti, M. Colajanni, Scalable, confidential and survivable software updates, *IEEE Transactions on Parallel and Distributed Systems* (2021).
- [82] K. Uhlemann, C. Engelmann, S. L. Scott, JOSHUA: Symmetric Active/Active Replication for Highly Available HPC Job and Resource Management, in: 2006 IEEE International Conference on Cluster Computing, IEEE, 2006, pp. 1–10.
- [83] C. Engelmann, S. L. Scott, C. Leangsuksun, X. He, Symmetric active/active high availability for high-performance computing system services: Accomplishments and limitations, in: 2008 Eighth IEEE International Symposium on Cluster Computing and the Grid (CCGRID), IEEE, 2008, pp. 813–818.
- [84] Hewlett Packard Enterprise Development L. P., Designing Disaster Recovery HA Clusters using Metrocluster and Continentalclusters (2011).
URL https://support.hpe.com/hpesc/public/docDisplay?docId=emr_na-c02814903
- [85] IBM Corporation, IBM PowerHA SystemMirror for AIX: Geographic Logical Volume Manager (2017).
URL http://public.dhe.ibm.com/systems/power/docs/powerha/722/hacmpeolvm_pdf.pdf
- [86] I. Shankar, A. Mishra, Enhanced cluster failover management, US Patent 8,484,510 (Jul. 9 2013).
- [87] Oracle Corporation, Oracle Solaris Cluster 4.3 Geographic Edition Overview (2016).
URL https://docs.oracle.com/cd/E56676_01/pdf/E56739.pdf
- [88] Oracle Corporation, Oracle Solaris and Oracle Solaris Cluster: Extending Oracle Solaris for Business Continuity (2010).
URL <http://www.oracle.com/technetwork/server-storage/solaris-cluster/documentation/solaris-cluster-businesscontinuity-168285.pdf>
- [89] Hewlett Packard Enterprise Development L. P., Understanding and Designing Serviceguard Disaster Recovery Architectures (2012).
URL https://support.hpe.com/hpsc/doc/public/display?docId=emr_na-c03604629
- [90] A. Kanso, M. Toeroe, F. Khendek, Comparing redundancy models for high availability middleware, *Computing* 96 (10) (2014) 975–993.
- [91] C. Engelmann, Symmetric active/active high availability for high-performance computing system services, Ph.D. thesis, Department of Computer Science, University of Reading, UK (2008).
URL <http://www.christian-engelmann.info/publications/engelmann08symmetric3.pdf>
- [92] X. He, L. Ou, C. Engelmann, X. Chen, S. L. Scott, Symmetric active/active metadata service for high availability parallel file systems, *Journal of Parallel and Distributed Computing* 69 (12) (2009) 961–973.
- [93] Oracle Corporation, Solaris Cluster Data Service for Oracle Real Application Clusters Guide (2016).
URL https://docs.oracle.com/cd/E56676_01/pdf/E57757.pdf
- [94] S. Distefano, F. Longo, M. Scarpa, Availability assessment of ha standby redundant clusters, in: 2010 29th IEEE Symposium on Reliable Distributed Systems, IEEE, 2010, pp. 265–274.
- [95] Y. Bouizem, N. Parlavantzis, D. Dib, C. Morin, Active-standby for high-availability in faas, in: Proceedings of the 2020 Sixth International Workshop on Serverless Computing, 2020, pp. 31–36.
- [96] A. Magalhaes, J. M. Monteiro, A. Brayner, Main memory database recovery: A survey 54 (2) (2021). doi: 10.1145/3442197.
URL <https://doi.org/10.1145/3442197>
- [97] T. Pohanka, V. Pechanec, Evaluation of replication mechanisms on selected database systems, *ISPRS International Journal of Geo-*

- Information 9 (4) (2020) 249.
- [98] G. Levitin, L. Xing, Y. Dai, Cold vs. hot standby mission operation cost minimization for 1-out-of-n systems, *European Journal of Operational Research* 234 (1) (2014) 155–162.
- [99] T. A. Nguyen, D. S. Kim, J. S. Park, Availability modeling and analysis of a data center for disaster tolerance, *Future Generation Computer Systems* 56 (2016) 27–50.
- [100] C. Leangsuksun, V. Munganuru, T. Liu, S. Scott, C. Engelmann, Asymmetric active-active high availability for high-end computing, in: *Proceedings of 2nd International Workshop on Operating Systems, Programming Environments and Management Tools for High-Performance Computing on Clusters (COSET-2)*, 2005.
- [101] G. E. Gonçalves, P. T. Endo, M. Rodrigues, D. H. Sadok, J. Kelner, C. Curescu, Resource allocation based on redundancy models for high availability cloud, *Computing* 102 (1) (2020) 43–63.
- [102] Z. Hou, Y. Huang, S. Zheng, X. Dong, B. Wang, Design and implementation of heartbeat in multi-machine environment, in: *17th International Conference on Advanced Information Networking and Applications*, 2003. AINA 2003., IEEE, 2003, pp. 583–586.
- [103] IBM Corporation, PowerHA SystemMirror concepts (2016).
URL http://public.dhe.ibm.com/systems/power/docs/powerha/721/hacmpconcepts_pdf.pdf
- [104] S. C. Dake, C. Caulfield, A. Beekhof, The corosync cluster engine, in: *Linux Symposium*, Vol. 85, Citeseer, 2008.
- [105] J. Maloy, Tipc: Providing communication for linux clusters, in: *Linux Symposium*, Vol. 2, 2004, pp. 347–356.
- [106] A. B. Forouzan, *Data communications & networking* (sie), Tata McGraw-Hill Education, 2007.
- [107] D. Dolev, D. Malki, The transis approach to high availability cluster communication, *Communications of the ACM* 39 (4) (1996) 64–70.
- [108] X. Défago, A. Schiper, P. Urbán, Total order broadcast and multicast algorithms: Taxonomy and survey, *ACM Computing Surveys (CSUR)* 36 (4) (2004) 372–421.
- [109] C. Engelmann, S. L. Scott, C. Leangsuksun, X. B. He, Symmetric active/active high availability for high-performance computing system services., *JCP* 1 (8) (2006) 43–54.
- [110] SIOS Technology Corp., SIOS Protection Suite for Windows (2018).
URL http://cdn.manula.com/user/1870/10540_11235_en_1504797428.pdf
- [111] Veritas Technologies LLC, Veritas InfoScale 7.4.3 Solutions in Cloud Environments (2020).
URL https://www.veritas.com/support/en_US/doc/145798412-145798468-1
- [112] SUSE LLC, SAP HANA High Availability Cluster for the AWS Cloud - Setup Guide (v15) (2021).
URL https://documentation.suse.com/sbp/all/pdf/SLES4SAP-hana-sr-guide-PerfOpt-15_AWS_color_en.pdf
- [113] M. Waldvogel, R. Rinaldi, Efficient topology-aware overlay network, *ACM SIGCOMM Computer Communication Review* 33 (1) (2003) 101–106.
- [114] Veritas Technologies LLC, Symantec High Availability Agent for SAP NetWeaver Installation and Configuration Guide (2013).
URL https://sort.veritas.com/agents/download_docs/8304/vcs_sapnetweaver_install
- [115] IBM Corporation, Smart Assists for PowerHA SystemMirror (2017).
URL http://public.dhe.ibm.com/systems/power/docs/powerha/721/smartassist_pdf.pdf
- [116] H. Yang, Y. Kim, Design and Implementation of Fast Fault Detection in Cloud Infrastructure for Containerized IoT Services, *Sensors* 20 (16) (2020) 4592.
- [117] NEC Corporation, EXPRESSCLUSTER X 3.3 for Linux Reference Guide (2017).
URL https://www.nec.com/en/global/prod/expresscluster/en/support/Linux/L33_RG_EN_06.pdf
- [118] NEC Corporation, EXPRESSCLUSTER X 3.3 for Windows Reference Guide (2017).
URL https://www.nec.com/en/global/prod/expresscluster/en/support/Windows/W33_RG_EN_06.pdf
- [119] J. S. Ward, A. Barker, Observing the clouds: a survey and taxonomy of cloud monitoring, *Journal of Cloud Computing* 3 (1) (2014) 24.
- [120] M. Q. Hiep, H. Yang, Y. Kim, Dynamic policy management system for high availability in a multi-site cloud, in: *2020 International Conference on Information and Communication Technology Convergence (ICTC)*, IEEE, 2020, pp. 359–362.
- [121] I. Kaitovic, M. Malek, Impact of failure prediction on availability: Modeling and comparative analysis of predictive and reactive methods, *IEEE Transactions on Dependable and Secure Computing* 17 (3) (2018) 493–505.
- [122] A. Beekhof, *Pacemaker 1.1 Configuration Explained An A-Z guide to Pacemaker's Configuration Options* (2017).
URL http://clusterlabs.org/pacemaker/doc/en-US/Pacemaker/1.1/pdf/Pacemaker_Explained/Pacemaker-1.1-Pacemaker_Explained-en-US.pdf
- [123] Y.-J. Lee, H.-Y. Kim, C.-H. Lee, A stochastic availability prediction model for head nodes in the ha cluster, in: *22nd International Conference on Advanced Information Networking and Applications-Workshops (aina workshops 2008)*, IEEE, 2008, pp. 157–161.
- [124] Oracle Corporation, *Oracle Clusterware Administration and Deployment Guide* (2017).
URL <https://docs.oracle.com/en/database/oracle/oracle-database/12.2/cwadd/clusterware-administration-and-deployment-guide.pdf>
- [125] D. Quintero, S. Bodily, D. Martin-Corben, R. Prathap, K. Singh, A. Thajudeen, W. Zanatta, I. Redbooks, *IBM PowerHA SystemMirror for AIX Cookbook*, IBM Redbooks, 2015.
- [126] C. Gomes, E. Tavares, M. N. d. O. Junior, B. Nogueira, Cloud storage availability and performance assessment: a study based on NoSQL DBMS, *The Journal of Supercomputing* (2021) 1–21.
- [127] M. Naor, A. Wool, The load, capacity, and availability of quorum systems, *SIAM Journal on Computing* 27 (2) (1998) 423–447.
- [128] X. Wang, H. Sun, T. Deng, J. Huai, On the tradeoff of availability and consistency for quorum systems in data center networks, *Computer Networks* 76 (2015) 191–206.
- [129] Microsoft Corporation, *Failover Clustering in Windows Server* (2018).
URL <https://docs.microsoft.com/en-us/windows-server/failover-clustering/failover-clustering-overview>
- [130] D. Malkhi, M. K. Reiter, A. Wool, R. N. Wright, Probabilistic quorum systems, *Information and Computation* 170 (2) (2001) 184–206.
- [131] Microsoft Corporation, *Behavior of Dynamic Witness on Windows Server 2012 R2 Failover Clustering* (2016).
URL <https://blogs.technet.microsoft.com/askcore/2016/03/21/behavior-of-dynamic-witness-on-windows-server-2012-r2-failover-clustering-3/>
- [132] L. Zhu, J. Lin, A scsi3 persistent reservation synchronization solution for iscsi targets cluster hosting ceph rbd with active/active connections, in: *2020 IEEE 4th Information Technology, Networking, Electronic and Automation Control Conference (ITNEC)*, Vol. 1, IEEE, 2020, pp. 1785–1793.
- [133] K. W. Preslan, A. P. Barry, J. Brassow, M. Declerck, A. Lewis, A. Manthei, B. Marzinski, E. Nygaard, S. Van Oort, D. Teigland, et al., Scalability and failure recovery in a linux cluster file system., in: *Annual Linux Showcase & Conference*, 2000.
- [134] T. Lumpp, J. Schneider, J. Holtz, M. Mueller, N. Lenz, A. Biazetti, D. Petersen, From high availability and disaster recovery to business continuity solutions, *IBM Systems Journal* 47 (4) (2008) 605–619.
- [135] Y.-L. Lee, S. N. Arizky, Y.-R. Chen, D. Liang, W.-J. Wang, High-availability computing platform with sensor fault resilience, *Sensors* 21 (2) (2021) 542.
- [136] D. Rossi, E. Turrini, Analyzing the impact of components replication in high available j2ee clusters, in: *Joint International Conference on Autonomous and Autonomous Systems and International Conference on Networking and Services-(icas-isns' 05)*, IEEE, 2005, pp. 56–56.
- [137] A. van der Linde, J. Leitão, N. Preguiça, Practical client-side replication: weak consistency semantics for insecure settings, *Proceedings of the VLDB Endowment* 13 (12) (2020) 2590–2605.
- [138] S. H. Mortazavi, M. Salehe, B. Balasubramanian, E. de Lara, S. PuzhavakathNarayanan, Sessionstore: A session-aware datastore for the edge, in: *2020 IEEE 4th International Conference on Fog and Edge Computing (ICFEC)*, IEEE, 2020, pp. 59–68.
- [139] J. Noble, D. Maxwell, K. X. Hourihan, R. Stephens, *Check Point NG VPN-1/FireWall-1: Advanced configuration and troubleshooting*, Syngress (Canada), 2003.
- [140] P. Fondo-Ferreiro, F. Gil-Castañeira, F. J. González-Castaño, D. Candal-

- Ventureira, A software-defined networking solution for transparent session and service continuity in dynamic multi-access edge computing, *IEEE Transactions on Network and Service Management* 18 (2) (2020) 1401–1414.
- [141] L. H. Le, C. E. Bezerra, F. Pedone, Dynamic scalable state machine replication, in: 2016 46th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN), IEEE, 2016, pp. 13–24.
- [142] F. Pedone, R. Guerraoui, A. Schiper, The database state machine approach, *Distributed and Parallel Databases* 14 (1) (2003) 71–98.
- [143] Y. Shi, J. Zuo, Y. Guo, Y. Lu, Distributed file system multilevel fault-tolerant high availability mechanism, in: *Proceedings of the 2020 International Conference on Cyberspace Innovation of Advanced Technologies*, 2020, pp. 431–438.
- [144] IBM Corporation, *IBM Spectrum Scale Concepts, Planning, and Installation Guide* (2017).
URL https://www.ibm.com/support/knowledgecenter/STXKQY_4.2.1/com.ibm.spectrum.scale.v4r21.doc/pdf/a7604417.pdf
- [145] R. Rani, N. Kumar, M. Khurana, A. Kumar, A. Barnawi, Storage as a service in fog computing: A systematic review, *Journal of Systems Architecture* (2021) 102033.
- [146] SIOS Technology Corp., *SANless Cluster Software* (2017).
URL <https://us.sios.com/products/sios-datakeeper/what-is-a-sanless-cluster/>
- [147] J. Lee, S. Moon, K. H. Kim, D. H. Kim, S. K. Cha, W.-S. Han, Parallel replication across formats in SAP HANA for scaling out mixed OLTP/OLAP workloads, *Proceedings of the VLDB Endowment* 10 (12) (2017) 1598–1609.
- [148] Y. Lu, X. Yu, L. Cao, S. Madden, Epoch-based commit and replication in distributed OLTP databases, *Proceedings of the VLDB Endowment* 14 (5) (2021) 743–756.
- [149] SIOS Technology Corp., *SIOS Protection Suite for Linux Documentation* (2017).
URL <http://docs.us.sios.com/Linux/9.2/LK4L/AllTechDocs/index.htm>
- [150] A. Gómez, L. Carril, R. Valin, J. C. Mouriño, C. Cotel, Fault-tolerant virtual cluster experiments on federated sites using bonfire, *Future Generation Computer Systems* 34 (2014) 17–25.
- [151] R. Kanagavelu, B. S. Lee, R. F. Miguel, L. N. Mingjie, et al., Software defined network based adaptive routing for data replication in data centers, in: 2013 19th IEEE International Conference on Networks (ICON), IEEE, 2013, pp. 1–6.
- [152] R. DelValle, P. Kaushik, A. Jain, J. Hartog, M. Govindaraju, Electron: Towards efficient resource management on heterogeneous clusters with apache mesos, in: 2017 IEEE 10th International Conference on Cloud Computing (CLOUD), IEEE, 2017, pp. 262–269.
- [153] DH2i., *DxEnterprise* (2020).
URL <https://dh2i.com/dxenterprise/>
- [154] Stratus Technologies., *everRun* (2020).
URL <https://www.stratus.com/solutions/platforms/everrun/>
- [155] I. Haddad, C. Leangsuksun, S. L. Scott, HA-OSCAR: the birth of highly available OSCAR, *Linux Journal* 2003 (115) (2003) 1.
- [156] D. Prior, A. MacNeela, I. Brown, J. Krischer, D. Scott, J. Green-Armytage, *Enterprise guide to gartner’s high-availability system model for SAP*, Gartner, December (2001).
- [157] R. Rabbat, T. McNeal, T. Burke, A high-availability clustering architecture with data integrity guarantees, in: *Third IEEE International Conference on Cluster Computing (CLUSTER’01)*, IEEE Computer Society, 2001, pp. 178–178.
- [158] S. Corsava, V. Getov, Intelligent architecture for automatic resource allocation in computer clusters, in: *Proceedings International Parallel and Distributed Processing Symposium*, IEEE, 2003, pp. 8–pp.
- [159] Gartner Inc., *Magic Quadrant for Operational Database Management Systems*, Tech. rep., Gartner, Inc., Stamford (2017).
- [160] Gartner Inc., *Market share analysis: ERP software, worldwide, 2017*, Tech. rep., Gartner, Inc., Stamford (2017).
- [161] Gartner Inc., *Market Share Analysis: Customer Relationship Management Software, Worldwide, 2016*, Tech. rep., Gartner, Inc., Stamford (2016).
- [162] SAP SE, *Certified HA-Interface Partners - SAP Application Server High Availability Interface Certification* (2018).
URL <https://wiki.scn.sap.com/wiki/display/SI/Certified+HA-Interface+Partners>
- [163] IBM Corporation, *Supported cluster management software* (2018).
URL https://www.ibm.com/support/knowledgecenter/SSEPGG_11.1.0/com.ibm.db2.luw.admin.ha.doc/doc/r0051388.html
- [164] A. Kanso, Y. Lemieux, Achieving high availability at the application level in the cloud, in: 2013 IEEE 6th International Conference on Cloud Computing (CLOUD), IEEE, 2013, pp. 778–785.
- [165] M. Khan, M. Toeroe, F. Khendek, Comparing pacemaker with OpenSAP for availability management in the cloud, in: *Edge Computing (EDGE), 2017 IEEE International Conference on*, IEEE, 2017, pp. 106–111.
- [166] U. Franke, Optimal it service availability: Shorter outages, or fewer?, *IEEE Transactions on Network and Service Management* 9 (1) (2011) 22–33.
- [167] J. Singh, P. Singh, S. S. Gill, Fog computing: a taxonomy, systematic review, current trends and research challenges, *Journal of Parallel and Distributed Computing* (2021).
- [168] F. Ramos, E. Viegas, A. Santin, P. Horschulhack, R. R. dos Santos, A. Espindola, A machine learning model for detection of docker-based app overbooking on kubernetes, in: *ICC 2021-IEEE International Conference on Communications*, IEEE, 2021, pp. 1–6.
- [169] Veritas Technologies LLC, *Veritas InfoScale 7.4.3 Support for Kubernetes - Linux* (2020).
URL https://www.veritas.com/support/en_US/doc/145798462-145798468-1
- [170] L. A. Vayghan, M. A. Saied, M. Toeroe, F. Khendek, Microservice based architecture: Towards high-availability for stateful applications with kubernetes, in: 2019 IEEE 19th International Conference on Software Quality, Reliability and Security (QRS), IEEE, 2019, pp. 176–185.
- [171] Y. Alahmad, A. Agarwal, T. Daradkeh, High availability management for applications services in the cloud container-based platform, in: 2018 IEEE/ACS 15th International Conference on Computer Systems and Applications (AICCSA), IEEE, 2018, pp. 1–8.
- [172] S. Loveland, E. M. Dow, F. LeFevre, D. Beyer, P. F. Chan, Leveraging virtualization to optimize high-availability system configurations, *IBM Systems Journal* 47 (4) (2008) 591–604.
- [173] J. R. Vacca, *Cloud computing security : foundations and challenges*, 1st Edition, CRC Press, 2016.
- [174] R. Chow, P. Golle, M. Jakobsson, E. Shi, J. Staddon, R. Masuoka, J. Molina, Controlling data in the cloud, in: *Proceedings of the 2009 ACM workshop on Cloud computing security - CCSW ’09*, ACM, 2009, p. 85.