

Magic Quadrant for Distributed Hybrid Infrastructure

7 September 2026 - ID G00841622 - 38 min read

By Julia Palmer, Elaine Zhang, [and 2 more](#)

Heads of I&O seek a hybrid cloud experience across public and private cloud, as well as edge deployments. The DHI market aims to reduce operational complexity and create a unified cloud-to-edge infrastructure management and control plane to simplify workload management and ensure policy enforcement.

Strategic Planning Assumption

By 2029, 55% of enterprises will initiate proofs of concept for alternative distributed hybrid infrastructure products to replace their VMware-based deployments and embrace hybrid cloud infrastructure delivery, up from 25% in 2026.

Market Definition/Description

Distributed hybrid infrastructure (DHI) is a solution incorporating cloud-native infrastructure principles (e.g., programmability, elasticity, modularity and resiliency). DHI can be deployed and managed, delivering cloud services wherever the customer chooses, including on-premises, at the colocation, at the edge or in the public cloud. This flexibility sets DHI apart from traditional public cloud infrastructure as a service (IaaS), which typically relies on a centralized approach. Vendors fall into three categories, with distinct strategies: cloud platform providers, full-stack hyperconverged vendors and software infrastructure platform providers. Their offerings include software and/or integrated hardware with a unified management and control plane over hybrid cloud infrastructure deployment.

DHI serves as the foundation for deploying applications in a distributed manner, while maintaining a cloud-native or cloud-inspired approach. This approach enhances agility and flexibility for workloads that require deployment beyond the public cloud infrastructure and addresses potential limitations of traditional on-premises setups. DHI offers greater consistency, flexibility, mobility, isolation and availability across customer environments.

Enterprises can benefit from the principles of cloud computing while retaining the ability to deploy resources whenever and wherever they are needed, and provide options to address workload-specific characteristics, such as latency, performance and data sovereignty. Infrastructure and operations (I&O) leaders seeking DHI typically have strong hybrid needs and do not necessarily adhere to a cloud-only or cloud-first strategy, often maintaining substantial on-premises deployments. This strategic approach enables enterprises to harness the advantages of both environments, offering the flexibility needed to standardize and optimize operations, uphold performance and availability service-level agreements (SLAs), and support a variety of workloads.

Mandatory Features

- Ability to deploy the infrastructure services in multiple locations, which must include on-premises customer data centers or colocation sites, and at least one hyperscale public cloud platform.
- Integrated platform services comprising hypervisor, storage and virtualized networking services.
- Vendor-engineered, hybrid cloud infrastructure, resource management control plane.
- Distributed infrastructure management supported by a set of APIs that encompasses all of the distributed infrastructure's capabilities.
- Centralized management portal that provides a secure, automated, full-stack solution for operating and supporting the distributed infrastructure.
- Support for virtual machines (VMs).

Optional Features

- Ability to deliver the DHI solution as a service.
- Ability to provide sovereign cloud support with air gapped operations.

- Vendor-supplied facilities or other deployment options, such as colocation, edge and remote locations.
- Integrated hardware — either supported by the same vendor that offers the software or certified for the integrated software, but managed by the enterprise, the vendor that provides the software, or a third-party partner.
- Additional platform services — e.g., as-a-service, add-on offerings, such as database as a service (DBaaS) and disaster recovery as a service (DRaaS).
- Product integration with traditional enterprise infrastructure systems — e.g., directory services, IT service management (ITSM).
- Metered consumption-based services — customers are charged only for what they consume.
- Scalable and elastic — the offering will provision and deprovision resources to fit the customer's continuous requirements (e.g., bursting to the public cloud or dynamic preemptive ordering).
- Support for graphics processing units (GPUs) and AI accelerators.
- Ability to provision and orchestrate bare-metal servers.
- Ability to provision, support and orchestrate containers, or integrate with third-party container orchestration platforms.
- Ability to deploy services across multiple public clouds.

Magic Quadrant

Figure 1: Magic Quadrant for Distributed Hybrid Infrastructure



Magic Quadrant for Distributed Hybrid Infrastructure



Gartner

Vendor Strengths and Cautions

Alibaba Cloud

Alibaba Cloud is a Challenger in this Magic Quadrant. Its DHI offering focuses on its Apsara Stack and Edge Node Service (ENS) family. Alibaba Cloud operates primarily in China and the Asia/Pacific region, with a growing presence in Europe, Middle East, Africa and Latin America. Key developments over the past year include the acquisition of ZStack to bolster small-scale deployment options, as well as the launch of OpenTrek (an AI agent platform), BCI instances for serverless AI and Apsara DataLake for AI to advance its AI infrastructure. Alibaba Cloud's

DHI offering is a good fit for enterprises requiring sovereign cloud and the ability to execute AI workloads across cloud and on-premises.

Strengths

- **AI-centric innovation:** Alibaba Cloud offers a platform for developing and running AI workloads across distributed environments. Its MaaS (model as a service) platform, OpenTrek, supports agent development and orchestration, allowing organizations to maintain data residency and localized operations.
- **Product simplification:** Apsara Stack supports connected or air-gap deployments and hardware integration options. It allows organizations to configure infrastructure based on security, sovereignty and administrative requirements. This differs from some of the models offered by other hyperscale providers.
- **Sales execution:** Leveraging the acquisition of ZStack and a strategic partnership with ZTE, Alibaba Cloud's DHI portfolio offers a small-footprint deployment and expanded local hardware support services in more countries than in 2025.

Cautions

- **Architectural fragmentation:** ZStack does not share the same architecture with Apsara Stack, requiring a separate management console for edge operations. Also, the native Alibaba Cloud on-premises experience still requires a minimum investment of 17 nodes to initiate an Apsara Stack deployment.
- **Deployment complexity:** Unclear documentation makes Apsara Stack difficult to deploy. Customers express they often require experts to handle the complex configuration process.
- **Inadequate support localization:** International language support for Alibaba Cloud's DHI offering remains restricted to Chinese and English. Minimal enhancement is achieved through the assistance of AI translation agents and specific language support plans in high-demand locations.

Amazon Web Services

Amazon Web Services (AWS) is a Leader in this Magic Quadrant. Its DHI portfolio spans customer-sited hardware and AWS-operated facilities, including AWS Outposts, AWS Local Zones, AWS Dedicated Local Zones and AWS AI Factories. For software-based offerings,

AWS provides Amazon ECS Anywhere and Amazon EKS Hybrid Nodes for cloud-managed containers, and Amazon EKS Anywhere for air-gapped Kubernetes. AWS serves enterprises globally, with customer-sited DHI shipping and parent-region connectivity supported in specific countries. Over the last year, AWS launched AWS AI Factories, expanded AWS Local Zones to new metropolitan areas and added AWS Outposts support for RDS for Oracle. It is best-suited for latency-sensitive, sovereign and hybrid AI workloads.

Strengths

- **Public cloud interoperability:** AWS DHI offerings serve as a natural extension to the AWS public cloud, eliminating the Day 2 operational friction typical of hybrid deployments.
- **Infrastructure scalability:** Scaling from 42U racks up to large-scale data center deployments, AWS DHI hardware solutions utilize standardized AWS compute and storage to preserve consistent deployment patterns locally.
- **Fully managed delivery model:** For its hardware-based DHI, AWS manages physical provisioning, component replacement and hypervisor maintenance, limiting customer responsibilities to site facilities and local networking.

Cautions

- **Tethered infrastructure:** With the exception of local EKS clusters, AWS DHI platforms are dependent on a persistent network tether to a parent AWS region. Any disruption to the service link or parent region itself immediately halts control plane operations locally.
- **Restricted local service catalog:** Even at data-center-scale deployments, AWS DHI offerings deliver a significantly narrower catalog of local services (fewer than 30) compared to competitors that provide more than 100 or full regional parity.
- **Premium pricing:** Reflecting local real estate and power costs, most AWS Local Zones carry a 15% to 35% price premium over their parent regions in expensive metros. Additionally, for AWS Dedicated Local Zones and AI Factories, the fully dedicated nature of the deployments drives substantial cost premiums.

Broadcom (VMware)

Broadcom (VMware) is a Leader in this Magic Quadrant. This evaluation focuses on VMware Cloud Foundation (VCF), which integrates compute, networking, storage, management and security to provide a software platform for VMs and containers across on-premises, cloud and

edge environments. Broadcom's VCF serves customers of all sizes and verticals across all geographies. Over the last year, Broadcom released VCF 9.0 (and recently announced 9.1, which isn't part of this evaluation), which natively integrates VCF Private AI Services, NVMe Memory Tiering and vSAN global deduplication. It also introduces hardware-enforced confidential computing and an enhanced vMotion feature to reduce downtime during live migration of GPU-accelerated virtual machines. It is a good fit for hybrid infrastructure management, edge computing and sovereign workload use cases.

Strengths

- **Broad, deep virtualization capabilities:** Built on the proven ESX hypervisor, VCF delivers a comprehensive virtualization stack that natively consolidates compute, storage and networking into a single, time-tested platform.
- **Formal sovereign cloud ecosystem:** For highly regulated industries, Broadcom offers its VMware Cloud Provider program, featuring over 50 independent, local providers globally.
- **AI-native infrastructure:** VCF 9.0 integrates its VCF Private AI Services, GPU virtualization and Agent Builder service directly into the VCF platform, allowing organizations to build and run sovereign AI natively on familiar infrastructure.

Cautions

- **Communication and sales practices:** VMware has experienced an increased level of negative sentiment from Gartner customers, particularly regarding communication, commercial business practices and delays in support, compared to other market leaders.
- **Limited data services:** VCF does not have mature advanced data services for unstructured file and object storage capabilities compared to other market leaders.
- **Edge licensing constraints:** VCF imposes strict commercial limits for edge deployments, including a mandatory 10-site minimum in the first year and a maximum cap of 256 cores per edge site.

Google

Google is a Visionary in this Magic Quadrant. Google Distributed Cloud (GDC) consists of GDC software-only on VMware and bare metal, GDC connected (appliance-only) and GDC air-gapped (appliance and rack form factors). While both software-only and air-gapped options are available globally, GDC connected is orderable in 29 countries spanning North

America, Europe, the Asia/Pacific region and Latin America. Over the last 12 months, Google introduced NVIDIA H200 and B300 GPU support, made Gemini on GDC air-gapped generally available and pivoted GDC connected into a targeted appliance for consuming Gemini Enterprise Agent Platform. GDC best serves sovereignty-conscious enterprises, regulated public-sector entities and defense organizations deploying localized AI.

Google declined requests for supplemental information. Gartner's analysis is therefore based on other credible sources.

Strengths

- **AI focus:** While capability depth varies by offering, Google embeds one of the market's most comprehensive AI stacks into GDC, enabling localized AI inference, optical character recognition (OCR) and natural language processing (NLP) on classified datasets while meeting strict sovereignty mandates.
- **Open-source foundation:** GDC's architecture is anchored on the Cloud Native Computing Foundation's open-source standards (Kubernetes, KubeVirt and Istio) rather than on proprietary technologies, preserving a degree of technological sovereignty and mitigating lock-in.
- **Sovereign pedigree:** GDC air-gapped holds authorizations for U.S. Secret/Top Secret data and DoD IL6. By satisfying the technical controls for strict frameworks, like NATO D-48 and NIST, it has driven broad adoption across global defense and public-sector agencies.

Cautions

- **Strategic uncertainty:** A pattern of disruptive portfolio churn, evidenced by continuous product rebranding (e.g., Anthos to GDC Virtual to GDC software-only), the deprecation of GDC connected rack and the shift of GDC connected into a targeted Gemini AI appliance, creates long-term roadmap uncertainty for customers seeking a general-purpose platform.
- **Architectural divergence:** While the market has gravitated toward architectural parity, GDC and Google Cloud do not share the same technical substrate. API operations, service maturity and capabilities differ between the two platforms, complicating workload portability.
- **Service fragmentation:** Advanced capabilities, including the specialized GDC Marketplace, AlloyDB and latest frontier Gemini models (which may be substituted with Gemma in memory-constrained deployments) are not uniformly available across GDC offerings.

Huawei

Huawei is a Niche Player in this Magic Quadrant. This evaluation covers its DHI portfolio, including Huawei Cloud Stack Online (HCSO), CloudPond and Huawei Cloud Stack (HCS). Huawei maintains a global presence, excluding North America, for enterprises of all sizes and verticals. Over the past year, Huawei has introduced CloudMatrix384 supernodes for AI infrastructure, incorporated a pooling solution based on the Huawei UnifiedBus high-speed interconnect protocol to optimize throughput, and launched CloudWAN and CloudDC to bolster distributed capabilities. Huawei's DHI offering is well-suited for enterprises requiring hybrid infrastructure, strict data sovereignty and the capacity for large-scale AI workloads.

Strengths

- **AI supercomputing:** Huawei provides an AI stack through its HCSO and HCS hardware-integrated offerings, featuring liquid-cooled CloudDC, CloudMatrix384 supernodes and the ModelArts platform for large-scale training and inference within a sovereign, customer-controlled environment.
- **Hardware-software synergy:** Huawei's DHI portfolio utilizes a vertically integrated model where both hardware and software are developed internally. This architectural synergy optimizes the end-to-end life cycle for better operational agility, deployment speed and overall cost-efficiency.
- **Operations management:** Huawei provides ManageOne as a cloud management platform for oversight of its public cloud, DHI portfolio and heterogeneous third-party environments. The platform provides life cycle operations and resource automation, leveraging open APIs to integrate networking, telemetry and monitoring across enterprise ecosystems.

Cautions

- **Supply chain uncertainty:** Geopolitical tensions drive long-term risks to Huawei's supply chain resilience, restricting access to advanced third-party chips and components. Lead times for Huawei's hardware-integrated solutions may be longer in some regions.
- **AI ecosystem:** The AI market is currently dominated by NVIDIA's CUDA ecosystem. Transitioning to Huawei's proprietary AI infrastructure is challenging, as migrating workloads to the Huawei ecosystem takes time and effort to achieve optimal performance.

- **Geographic service availability:** Huawei's DHI deployment model is limited globally. Specific services like CloudDC are currently only available in China, and international availability of its CloudMatrix384 supernode is expanding slowly.

IBM

IBM is a Challenger in this Magic Quadrant. This evaluation covers IBM's DHI offering, consisting of IBM Fusion HCI, IBM Power, IBM Power Virtual Server (PowerVS), IBM Cloud services and IBM Cloud Satellite. IBM Fusion HCI combines elements of Red Hat OpenShift software (Red Hat is a fully owned subsidiary of IBM) with IBM's Cloud Satellite cloud management, for an x86-based DHI platform supporting virtual machines and containers. The cloud-based IBM PowerVS, together with the on-premises IBM PowerVS Private Cloud, supports Power-based DHI. Since last year, IBM has added streamlined IBM watsonx portfolio integration and improved Ceph and IBM Storage Scale capabilities for IBM Fusion HCI. For IBM PowerVS, it introduced the Power11 generation of servers and added two additional IBM Cloud regions offering PowerVS. IBM Fusion HCI is well-suited for enterprise AI adoption and workload modernization, providing a consistent OpenShift-based experience, while its PowerVS-based DHI is a good fit for organizations that are already running Power workloads or evaluating Power as a landing zone for enterprise workloads such as IBM i, Oracle and SAP.

Strengths

- **Strong AI/ML integration:** IBM has integrated AI capabilities from other product lines into its DHI solutions. This includes IBM watsonx on Fusion HCI and the use of IBM Fusion Content-Aware Storage (CAS) for managing enterprise data in AI workloads like retrieval-augmented generation (RAG). IBM also offers its own open-source AI models called IBM Granite.
- **High-availability servers:** IBM Power and PowerVS platforms support high-availability mission-critical workloads such as IBM i, core banking applications, Epic healthcare databases, and very large in-memory SAP RISE, SAP HANA and Oracle workloads.
- **Comprehensive vertical support:** IBM delivers cloud solutions tailored for regulated industries, (e.g., financial services, healthcare) through IBM Consulting, removing the need for additional system integrators and providing a single contract to the end user.

Cautions

- **Portfolio fragmentation:** Some IBM Fusion HCI components, like IBM Storage Scale, are not yet fully integrated into the Fusion portfolio, complicating some documentation and tool navigation.
- **Software dependency on Red Hat OpenShift:** Fusion HCI relies on Red Hat OpenShift, limiting use for any customers that prefer alternate container platforms.
- **Limited on-premises hardware options:** IBM Fusion HCI is only available on Dell and Lenovo models which may limit its appeal to customers that prefer a broader range of compute options.

Microsoft

Microsoft is a Leader in this Magic Quadrant. This evaluation focuses on Azure Local and Azure Arc. Azure Arc extends Azure management and governance capabilities to Azure Local, multicloud, VMware virtual machines, Windows Server, and other on-premises or edge resources through a common control plane. Microsoft operates globally across North America, Latin America, EMEA and the Asia/Pacific region. It services organizations of all sizes and verticals. Over the last year, Microsoft has expanded Azure Local's capabilities by enabling fully disconnected operations with Microsoft 365 Local and increasing multitrack cluster support up to hundreds of nodes by introducing disaggregated SAN storage architectures for independent compute and storage scaling. Microsoft enhanced its on-premises AI offerings by launching Foundry Local in preview for sovereign AI inferencing and making Azure Migrate for VMware generally available to streamline infrastructure replatforming. Microsoft's DHI is best-suited for hybrid infrastructure management, sovereign cloud and edge use cases.

Strengths

- **Unified single-product architecture:** Unlike other public cloud providers that offer fragmented products for edge, data center and sovereign environments, Microsoft streamlines hybrid cloud with Azure Local as a single unified DHI platform.
- **Market traction:** Microsoft's DHI solution continues to build significant momentum, with its installed base growing rapidly and evidence of increasing interest from Gartner clients.
- **Flexible HCI and disaggregated architecture:** Azure Local now supports both hyperconverged and disaggregated architecture. Customers can scale compute and

storage independently and use existing enterprise storage to accommodate a range of needs, from small edge sites to large data center environments.

Cautions

- **Implementation considerations:** Some Gartner customers report that Azure Local has an intricate setup, requiring hardware integration, networking design and strong Azure expertise to get started, with some features and ecosystem support still maturing.
- **Azure Local scalability:** Azure Local's new multirack, data-center-scale deployment capabilities are very recent, require a minimum four-rack deployment, and involve highly prescriptive hardware and complex third-party SAN storage integration.
- **AI infrastructure immaturity on Azure Local:** On-premises enterprise-scale AI inference and agentic orchestration through Foundry Local on Azure Local are currently in preview and not yet generally available.

Nutanix

Nutanix is a Leader in this Magic Quadrant. This evaluation focuses on Nutanix Cloud Platform (NCP), which is a unified, software-defined hybrid cloud platform designed for centralized management of applications and data across various environments. Nutanix operates globally, utilizing a 100% channel-led sales model through resellers, service providers and hardware OEMs across North America, EMEA and the Asia/Pacific region. Nutanix Cloud Clusters (NC2) are available in AWS, Azure and Google Cloud. Over the last year, Nutanix added support for disaggregated external SAN storage from vendors such as Dell and Everpure, and expanded its sovereign cloud coverage to AWS GovCloud and OVHcloud. The company also introduced Nutanix Enterprise AI with an Agent Gateway for LLM governance, integrated model security with Palo Alto Networks and NKP Metal for bare-metal Kubernetes deployments. NCP is suited to sovereign workloads, hybrid infrastructure management and edge use cases.

Strengths

- **Hybrid cloud operations:** Nutanix has expanded its NC2 footprint across major hyperscalers (AWS, Azure and Google Cloud) as well as sovereign providers, and introduced support for regulated environments. Nutanix NC2 customers can now scale workloads into the public cloud for workload elasticity, disaster recovery and specialized GPU hardware access.

- **Virtualization agility:** Nutanix offers a reliable alternative to legacy hypervisors. Its free Nutanix Move tool automates the migration of virtual machines and network security policies, enabling enterprises to replatform their hybrid cloud infrastructure with minimal disruption and without a steep learning curve.
- **Deployment flexibility and hardware reuse:** Nutanix now supports disaggregated architectures, enabling customers to run its compute software on a wide range of servers and integrate with third-party external storage arrays to maximize existing hardware investments.

Cautions

- **HCI perception:** Nutanix is often seen mainly as a hyperconverged product rather than a hybrid cloud platform, causing prospects seeking cloud-native solutions to overlook its hybrid capabilities.
- **Cost and licensing complexity:** Gartner clients report that Nutanix's licensing and pricing models can be complex and less competitive than other market leaders.
- **Billing constraints:** DHI buyers often expect the flexibility of native public cloud consumption models. Nutanix does not support direct pay-as-you-go billing for on-premises customer deployments, which can reduce its attractiveness for organizations prioritizing cloud-like financial flexibility.

Oracle

Oracle is a Leader in this Magic Quadrant. Its DHI offering consists of customer-sited infrastructure, including Oracle Roving Edge Infrastructure, Oracle Exadata Cloud@Customer, Oracle Cloud Infrastructure (OCI) Dedicated Region, Oracle Cloud Isolated Region and Oracle Alloy, alongside Oracle's multicloud services. Oracle is a global provider with DHI customers that tend to be large enterprises with existing investments in the Oracle ecosystem, as well as government entities. Over the past year, Oracle achieved general availability of Oracle AI Database@AWS, introduced Hosted Dedicated Regions and launched Oracle Multicloud Universal Credits. Oracle's DHI is one of the best fits for organizations with strict sovereignty mandates spanning tactical edge to large-scale, air-gapped needs or multicloud database requirements.

Strengths

- **Public cloud parity:** Unlike competing solutions that expose limited subsets, Oracle's connected data-center solutions deliver an identical OCI service catalog and architecture, with no on-premises financial premium. Its air-gapped offering maintains this exact parity for all non-internet-dependent services.
- **Broad use-case suitability:** Oracle's DHI portfolio spans tactical edge, rack-scale database, Oracle-hosted, air-gapped and white-labeled environments. This breadth enables organizations to more closely match DHI deployments with their physical, connectivity and regulatory needs.
- **Strong multicloud support:** Oracle's multicloud strategy spans two distinct offerings: zero-egress-free interconnects with other hyperscalers, and embedded database infrastructure (Database@) located inside AWS, Azure and Google Cloud. This approach enables native hyperscaler integration, while Oracle Multicloud Universal Credits simplify multicloud consumption under a single OCI contract, alongside existing hyperscale agreements.

Cautions

- **Hardware ownership constraints:** Oracle retains legal ownership of its DHI hardware. This may preclude adoption by sovereign and highly regulated organizations with strict legal mandates requiring outright hardware ownership.
- **Fragmented multirealm administration:** OCI Dedicated Region operates as a distinct, self-contained realm rather than sharing a single management console with public OCI. This requires separate administration for identity, policies and service configurations, and introduces cross-realm limitations for native disaster recovery and data replication.
- **Disparate development focus:** Oracle's developments in the past year have focused primarily on OCI Dedicated Region, Oracle Alloy and its multicloud capabilities. Notably, smaller-scale offerings have experienced fewer new feature releases and service catalog expansion.

Red Hat

Red Hat (a wholly owned subsidiary of IBM) is a Visionary in this Magic Quadrant. Red Hat's DHI platform is Red Hat OpenShift. The software platform integrates container orchestration with cloud-native software deployment and life cycle management. Red Hat OpenShift can manage and automate the life cycle of containers and virtual machines through a common Kubernetes-based control plane. In addition to deploying self-managed Red Hat OpenShift

on-premises or in public clouds, the platform is also offered as a service with many public cloud vendors, including AWS and Azure. Red Hat OpenShift is a good fit for organizations seeking GitOps workflows and deploying cloud-native applications.

Red Hat declined requests for supplemental information. Gartner's analysis is therefore based on other credible sources.

Strengths

- **Cloud-native and container/virtual machine unification:** Red Hat integrates DevOps-oriented features for operational consistency in managing VMs and containers through tools like CI/CD pipelines and GitOps workflows.
- **Innovation:** Red Hat has rapidly adopted new features, such as adding ARM server and vGPU support in OpenShift Virtualization, IPv6 single-stack support and a GenAI virtual assistant in the OpenShift dashboard in 2025. Red Hat has also introduced new joint offerings with key cloud providers (e.g., AWS, Azure and Oracle).
- **Broad compatibility:** Red Hat OpenShift is tightly integrated with multiple public clouds, including extensive technical and go-to-market cooperation and unified billing. It can also be run on top of other DHI platforms, including Broadcom (VMware) and Nutanix. Red Hat also offers migration toolkits for both VMs and containers to move between locations.

Cautions

- **Lack of integrated on-premises product:** Red Hat doesn't offer OpenShift-integrated hardware for on-premises deployment. Instead, it relies on other providers like server OEMs for integration and hardware support.
- **Requires Kubernetes skills:** Red Hat OpenShift is optimized to support container operations, requiring training for administrators accustomed to VM-focused management tools.
- **Storage services lack advanced features:** Gartner clients report that the Ceph-based Red Hat OpenShift Data Foundation doesn't have some advanced storage management features, requiring them to implement third-party storage solutions.

SUSE

SUSE is a Visionary in this Magic Quadrant. Its DHI platform consists of SUSE Rancher Prime with SUSE Virtualization, with other optional packages including SUSE Storage for software-

defined storage and SUSE Security for container security and threat protection. Rancher can be deployed in most public clouds as well as on-premises, with SUSE Rancher Prime providing consistent operations between clusters. SUSE operates globally, with most customers located in North America and Europe. SUSE's strategy is built around open-source software including Linux, Kubernetes, KubeVirt and Longhorn. In the last year, SUSE expanded its edge capabilities by acquiring industrial IoT management company Losant; added AI capabilities such as Liz, which is an SRE-assistant agent; and added an automated virtual machine migration tool. SUSE is a good fit for organizations prioritizing open-source software platforms.

Strengths

- **Container/virtual machine unification:** SUSE's unified container and virtual machine management plane, as well as inclusion of infrastructure-as-code management and continuous-delivery capabilities like Fleet and Argo CD, helps organizations deliver DevOps and automation capabilities.
- **Broad public cloud support:** SUSE's Rancher platform, including SUSE Virtualization, is supported in multiple public clouds, including AWS, Microsoft Azure, Google Cloud and Alibaba Cloud.
- **Sovereign supply chain:** SUSE is an appealing option for organizations seeking sovereign alternatives to U.S.- or China-based DHI providers. It is headquartered in Europe, is an open-source software company and focuses on security-related capabilities (e.g., immutable base OS), and is deployable on EU-based cloud providers, such as OVHcloud.

Cautions

- **Limited portfolio awareness:** Gartner clients know SUSE primarily as a Linux vendor and often lack awareness of the DHI capabilities enabled by SUSE Rancher Suite.
- **Lack of integrated on-premises product:** SUSE does not offer SUSE-integrated hardware for on-premises deployment. Instead, it relies on other providers, like server OEMs, for integration and hardware support.
- **Requires cloud-native skills:** SUSE Rancher Prime is optimized to support Kubernetes clusters, requiring training for administrators accustomed to VM-focused management tools.

Tencent Cloud is a Niche Player in this Magic Quadrant. This DHI evaluation focuses on its Cloud Dedicated Cluster (CDC), Cloud Dedicated Zone (CDZ) and Tencent Cloud Enterprise (TCE) offerings. Tencent Cloud also offers Tencent Cloud-native Suite (TCS) for container management and AI platform services. Tencent Cloud operates primarily in China and the Asia/Pacific region, but has expanded into Latin America, Europe and the Middle East. It serves large enterprises in China as well as businesses expanding abroad in the internet, manufacturing and finance sectors. Over the last 12 months, Tencent Cloud enhanced its AI infrastructure capabilities, including embedding its Tencent Cloud TI-ONE Platform and TACO Kit solutions for model operation and management. Also, it introduced an integrated high-performance intelligent computing suite: a full-stack, AI infrastructure solution designed for large-scale, deep-learning and foundational AI model workloads. Tencent Cloud is a good fit for enterprises requiring strict data sovereignty, robust full-stack AI capabilities and modernization of existing infrastructure.

Strengths

- **Consistent architecture:** Tencent Cloud's DHI solutions feature an architecture consistent with its public cloud, reducing the learning curve for organizations that already use Tencent's public cloud.
- **Flexible deployment:** Both TCE and TCS are software-defined solutions, enabling organizations to repurpose existing hardware or choose certified hardware to deploy across any chosen location.
- **Broad use-case suitability:** Tencent Cloud provides multiple DHI services that can be configured to address different sovereign, technical and operational needs and use cases.

Cautions

- **Global presence constraint:** Tencent's smaller public cloud market share limits the global growth of its DHI products, especially for CDC/CDZ, which are required to connect to its public cloud.
- **Limited international partners:** Tencent Cloud's major footprint is in China and the Asia/Pacific region; the global strategic partner ecosystem is underdeveloped. As a result, customers outside those geographies may face support issues.
- **Hardware reuse complexity:** Reusing hardware in CDC/CDZ environments requires a Cloud Hosting Cluster solution that isn't included in the standard service-level agreements.

Vendors Added and Dropped

We review and adjust our inclusion criteria for Magic Quadrants as markets change. As a result of these adjustments, the mix of vendors in any Magic Quadrant may change over time. A vendor's appearance in a Magic Quadrant one year and not the next does not necessarily indicate that we have changed our opinion of that vendor. It may be a reflection of a change in the market and, therefore, changed evaluation criteria, or of a change of focus by that vendor.

Added

- Red Hat
- SUSE

Dropped

No vendors were dropped from this Magic Quadrant.

Inclusion and Exclusion Criteria

Inclusion Criteria

To qualify for inclusion in this Magic Quadrant, providers:

- Must sell a DHI offering aligning to the DHI market definition. All mandatory features must be generally available as of 1 May 2026.
- Must have demonstrated achievement of 60 enterprise customers deploying products in DHI scenarios, or reported over \$50 million in ARR contract value on 1 May 2026, or be on trajectory to meet either of these criteria within the next 12 months.
- Must have evidence that all DHI production customers' deployments are across on-premises and at least one hyperscale strategic public cloud environment (see [Magic Quadrant for Strategic Cloud Platform Services](#)). For hyperscale cloud providers, there should be evidence for on-premises or edge DHI deployment; for software providers or full stack HCI vendors, there should be evidence in the deployment of the DHI offering in the hyperscale strategic public cloud.

- All DHI offerings must have at least 10 production customers per region in on-premises data centers in at least four out of seven global regions (North America, Europe, Asia/Pacific [excluding China], Latin America, China, Middle East and Africa).
- Must have the following business capabilities:
 - Be the primary developer and IP owner of the software components of DHI (control plane and hypervisor), with an exception of open-source-based products. If any part of the technology is based on open source, the vendor should be one of the top 10 open-source project contributors.
 - Have 24/7 DHI customer support (including phone support). There must be an English-language localization of the contract, service portal, documentation and support.
 - Have a managed service minimum DHI availability of not lower than 99.9%.
- Must have the following technical capabilities relevant to Gartner clients:
 - Ability to deploy and operate the DHI where the customer prefers, including on-premises, at the edge or colocated
 - Ability to deploy and operate the DHI in one or more hyperscale public clouds
 - A vendor-developed DHI control plane that provides a secure, automated and integrated capability for onboarding, operating and life cycle management of infrastructure, and supporting the distributed infrastructure across on-premises and in the public cloud.
 - Ability to address three out of four use cases defined in the companion [Critical Capabilities for Distributed Hybrid Infrastructure](#).

Exclusion Criteria

Offerings are excluded from this Magic Quadrant if they are exclusively marketed and sold as container management infrastructure products and cannot accommodate virtual machine infrastructure. In addition, offerings are excluded that do not have the ability to deploy full DHI services in the hyperscale public cloud.

Evaluation Criteria

Ability to Execute

The Ability to Execute criteria for this Magic Quadrant are as follows:

Product/Service: This criterion covers the assessment of vendor capabilities to deliver and differentiate features and functionality supporting most market use cases, diversification of customer use across the vendor's portfolio and the scope of issues impacting customer experience.

Overall Viability: This criterion covers the assessment of a vendor's key financial, staffing and customer base growth metrics.

Sales Execution/Pricing: This criterion covers the assessment of a vendor's success in the market. Considerations include results of new versus repeat business, growth of new DHI customers and changes in customer investments. Adaptations to sales and presales efforts and levels of pricing transparency are also considered.

Market Responsiveness/Record: This criterion evaluates the vendor's ability to deliver DHI products and capabilities that are first-to-market and differentiating compared to the competition, while also continuing to meet market demands and gaps in their portfolio.

Marketing Execution: This criterion evaluates the vendor's ability to create mindshare, expand to new markets and build a sales pipeline in the DHI market.

Customer Experience: This criterion assesses the vendor's ability to demonstrate ongoing client satisfaction, implement improvements and offer distinctive customer support capabilities. This also focuses on how the vendor is perceived in the market, and how well its marketing programs are recognized. The evaluation focuses on how well the vendor is able to influence and shape perception in the market through marketing activities and thought leadership that drives awareness. An additional indicator for this criterion is how often Gartner clients inquire about a specific vendor in terms of capabilities/reputation or in a shortlist evaluation process.

Ability to Execute Evaluation Criteria	
<i>Evaluation Criteria</i>	<i>Weighting</i>
Product or Service	High

<i>Evaluation Criteria</i>	<i>Weighting</i>
Overall Viability	High
Sales Execution/Pricing	High
Market Responsiveness/Record	Medium
Marketing Execution	Low
Customer Experience	High
Operations	NotRated

Source: Gartner (September 2026)

Completeness of Vision

The Completeness of Vision criteria for this Magic Quadrant are as follows:

Market Understanding: This criterion evaluates the ability of the vendor to understand customer requirements, align those requirements to its products and services, and evolve its product vision, based on its own established perspectives of the market's direction.

Marketing Strategy: This criterion evaluates the clarity of the vendor's marketing vision that highlights competitive differentiation and an understanding of personas engaged in solution selection.

Sales Strategy: This criterion evaluates the vendor's ability to establish and update a sales strategy that aligns with company goals and customer interest. Factors also include the vendor's ability to reach customers directly and expand coverage through its network of partners.

Offering (Product) Strategy: This criterion evaluates the vendor's product planning, emphasizing its alignment to shortcomings, commitment to differentiation and improvement of existing capabilities.

Business Model: This criterion evaluates the vendor’s strategies to sustain its business in the market.

Innovation: This criterion evaluates the plans to bring future differentiated capabilities to market that will enhance the vendor’s ability to interact with customers and drive business. It also evaluates the vendor’s strategy for reinvestment and its differentiating innovations in product design and capabilities, introduction of new technologies, third-party partner relationships, integration, and overall differentiation in the DHI market.

Geographic Strategy: This criterion evaluates the vendor’s strategy to direct resources, skills and product offerings to meet needs across most major geographies.

Completeness of Vision Evaluation Criteria	
<i>Evaluation Criteria</i>	<i>Weighting</i>
Market Understanding	High
Marketing Strategy	Medium
Sales Strategy	Medium
Offering (Product) Strategy	High
Business Model	Medium
Vertical/Industry Strategy	NotRated
Innovation	High
Geographic Strategy	Medium

Source: Gartner (September 2026)

Quadrant Descriptions

Leaders

Leaders distinguish themselves by offering a service suitable for strategic adoption and having an ambitious roadmap. They can serve a broad range of use cases. However, they do not excel in all areas, may not necessarily be the best providers for a specific need and may not serve some use cases at all. Leaders in this market have appreciable market share and many referenceable customers supporting multiple geographies, verticals and deployment models.

Challengers

Challengers are well-positioned to serve some current market needs. They deliver a good service that is targeted at a particular set of use cases, and they have a track record of successful delivery. However, they might not adapt to market challenges quickly enough or do not have a broad scope of ambition. These vendors have the potential to establish themselves across the broader global market, but have not yet done so.

Visionaries

Visionaries are typically vendors that focus on strong innovation and product differentiation, with the potential to significantly disrupt the market if execution improves. Their services are still emerging, and they have many capabilities in development that are not yet generally available. Although they may have many customers, they might not yet serve a broad range of use cases well or may have a limited geographic scope.

Niche Players

Niche Players in the DHI market may be excellent providers for particular use cases or in regions where they operate, but they should ultimately be viewed as specialist providers. They often do not serve a broad range of use cases, deployment models or customer segments, or have a broadly ambitious roadmap. Some may have solid leadership positions in markets adjacent to this market, but have developed only limited global DHI capabilities.

Context

Heads of I&O seeking to modernize IT infrastructure and adopt cloud-native and AI-driven operations across on-premises, cloud and edge environments should take the following

actions:

- **Develop a hybrid platform strategy.** Create a plan that aligns with organizational goals and integrates on-premises, cloud and edge components. In 2026, include a modernization roadmap that supports both virtual machines and containers on a unified stack, and prepares for potential hypervisor licensing changes to avoid vendor lock-in.
- **Analyze workloads and use cases.** Review all workloads and applications for DHI suitability. Assess performance, capacity, SLA and data requirements. Profile workloads for AI readiness, identifying GPU needs, high-performance storage, and networking requirements for large language model training and local AI inference.
- **Engage with vendors and partners.** Select DHI vendors that offer unified solutions across all environments. Work with their ecosystem partners, including AI hardware providers and third-party storage vendors, to ensure the platform can scale compute and storage independently for data-intensive workloads.
- **Assess security and compliance.** Evaluate DHI solutions for security and regulatory needs, including digital sovereignty and resilience to geopolitical disruptions. Require hardware-based confidential computing, automated ransomware recovery and FIPS 140-3 compliance.
- **Leverage edge capabilities.** Use DHI edge solutions for latency-sensitive tasks and local processing of high-bandwidth data. Deploy edge AI and local AI factories to run agent workflows and avoid sending sensitive data to public clouds, thus reducing latency, cost and compliance risks.
- **Evaluate resource management tools.** Review DHI control plane and orchestration tools for hybrid management. Focus on AI Ops and native AI assistants that simplify distributed cluster management. Confirm the platform can handle heterogeneous hardware and dynamic GPU allocation.
- **Test and validate deployments.** Conduct thorough testing and validation of DHI setups to identify performance issues and integration challenges. Ensure the solution can operate with a fully autonomous local control plane during extended cloud disconnections, maintaining application performance, identity management and sovereign AI operations.

Following these steps will help heads of I&O effectively manage DHI and support operational and strategic objectives.

Market Overview

The DHI market was born from the head of I&O's need to deploy standardized infrastructure platforms for any deployment scenario. However, it has evolved tremendously from its earlier market phases.

In 2026, the DHI market has emerged at the convergence of two trends: the adoption of distributed cloud solutions, extending public cloud services to noncloud customer environments; and the deployment of full-stack software-defined platforms or hyperconverged solutions in public cloud environments.

A massive catalyst accelerating this convergence is the widespread migration and modernization of legacy virtualization environments, particularly following Broadcom's acquisition of VMware. Organizations are actively seeking DHI solutions to provide modernization pathways that treat virtual machines as first-class citizens, alongside modern Kubernetes and cloud-native workloads.

In this evolving market scenario, organizations are strategically recalibrating their approach as they consider a diverse range of options at their disposal. Cloud-inspired on-premises solutions and distributed public cloud IaaS (infrastructure as a service) offer compelling DHI avenues.

A defining attribute of the 2026 market is the absolute necessity for a unified control plane that manages diverse infrastructure estates without fragmentation.

This hybrid framework facilitates centralized control across an array of environments, including on-premises, public cloud and edge infrastructure, while addressing the intricate operational needs of modern businesses.

Heads of I&O have a rapidly growing strategic interest in sovereign cloud capabilities, and DHI has emerged as a key enabler in delivering them. In fact, DHI has become the foundation for "Sovereign AI," enabling organizations to train, fine tune and run inference on proprietary and frontier models entirely within their legally controlled boundaries.

Sovereignty has evolved from simple data residency to demanding absolute jurisdictional control over data, AI models, infrastructure and daily operations. To meet complex global regulations, DHI providers are now delivering fully disconnected (air-gapped) operations as a

basic requirement, allowing entire cloud stacks and enterprise productivity tools to run autonomously without public cloud connectivity.

Recently, special emphasis has been placed on the ability to rely on DHI products to enable AI-capable infrastructure delivery. AI workloads require flexible, high-performance environments that can handle intensive compute, large datasets and real-time processing.

DHI enables organizations to run AI training on powerful compute platforms often delivered as fully managed, on-premises AI factories equipped with the latest accelerators (e.g., NVIDIA GPUs, custom AI silicon), petabit-scale networking and agentic AI gateways. This allows organizations to operate directly at the data source, improving efficiency, tackling data gravity, eliminating egress costs and preventing sensitive corporate data from ever leaving the facility.

DHI Solutions Are No Longer Limited to IaaS Without PaaS

Vendors now deliver full-stack cloud services, including robust platform as a service (PaaS) and model as a service (MaaS) capabilities, directly to customer data centers and edge locations. This encompasses scalable compute (e.g., virtual machines, bare metal, containers), storage, networking, managed databases and advanced AI toolchains, deployed wherever the customer prefers.

⊕ Evidence

⊕ Evaluation Criteria Definitions

Unlock more AI value

With unmatched, proprietary AI insights from Gartner

[Explore the AI Hub ↗](#)

© 2026 Gartner, Inc. and/or its affiliates. All rights reserved. Gartner is a registered trademark of Gartner, Inc. and its affiliates. This publication may not be reproduced or distributed in any form without Gartner's prior written permission. It consists of the opinions of Gartner's research organization, which should not be construed as statements of fact. While the information contained in this publication has been obtained from sources believed to be reliable, Gartner disclaims all warranties as to the accuracy, completeness or adequacy of such information. Although Gartner research may address legal and financial issues, Gartner does not provide legal or investment advice and its research should not be construed or used as such. Your access and use of this publication are governed by [Gartner's Usage Policy](#). Gartner prides itself on its reputation for independence and objectivity. Its research is produced independently by its research organization without input or influence from any third party. For further information, see "[Guiding Principles on Independence and Objectivity](#)." Gartner research may not be used as input into or for the training or development of generative artificial intelligence, machine learning, algorithms, software, or related technologies.

[About Us](#) [Search Careers](#) [Newsroom](#) [Policies](#) [Site Index](#) [Featured Insights](#)
[Contact Us](#) [Send Feedback](#)

Gartner

© 2026 Gartner, Inc. and/or its Affiliates. All Rights Reserved.