




PLATFORM9

WEBINAR

Beyond VMware: Introducing Platform9

A practical look at migration, risk reduction and full-stack private cloud with Platform9

September 22, 2026 • 12:00 PM AEST • Microsoft Teams Webinar



Agenda

- **Introduction to Platform9:**
The company history and how they got here
- **Market shifts:**
Why it's more important than ever to your business
- **Positioning:**
How Platform9 compares with other alternatives in the VMware and virtualisation space
- **Beyond virtualisation:**
A full-stack private cloud, including Kubernetes and a preview of the private AI inference platform



Meet The Speakers



Peter Fray
Field CTO, Platform9




Phi Dao
Solution Sales Senior Specialist – Cloud
Crayon

From VMware to What's Next

Same workloads. More freedom.

Modernise | Simplify | Future-proof



Keep your workloads moving



Reduce complexity and cost



Gain flexibility and choice



Build for the future

Different platform. A brighter future.

THE CRAYON DIFFERENCE

The trusted ANZ & SEA partner.

- **01 | UNIQUE BROADCOM POSITION: Only CCM for Broadcom in ANZ & SEA**

- Local market understanding backed by regional capabilities
- One accountable partner(Distribution) across the transition.

- **02 | PLATFORM9 — NOW ON BOARD**

- Platform9 adds a proven VMware replacement path to our portfolio
- Helps customers evaluate migration options with confidence
- Crayon can support the journey from assessment → migration → optimisation

- **03 | PARTNER-FIRST ECOSYSTEM**

We enable partners — not compete with them

- Commercial, technical and migration support
- Extend partner capabilities with Crayon expertise

Platform9 Programs

Three distinct Pricing Models.

01



Enterprise Distribution model

Distributor-led route to market, built on partner margin and NDA-based engagement.

02



Portkey SMB + deployment-led model

Entry-level program with a one-time deployment fee, ideal for smaller partners and customers.

03



PCSP Subscription + onboarding model

Single-SKU subscription that drives long-term platform adoption through partner onboarding. The Platform9 Cloud Service Provider Program (PCSP) is purpose-built for VCSP members who lost the ability to sell VMware-based services when Broadcom reduced its authorized partners.

Your VMware Exit, De-Risked

Your timeline. Your infrastructure. Your Cloud.



Platform9: The Trusted Private Cloud

The Track Record Behind the Trust



Deep Expertise

Founders built VMware vCloud
Director



Proven at Enterprise Scale

1,200 Hypervisors single deployment



Platform9 Named a Visionary

2026 Gartner® Magic Quadrant™ for
Server Virtualization Platforms

Trusted by Global Enterprises & Service Providers



AEROMEXICO



sambanova

Magic Quadrant for Server Virtualization Platforms



Platform9 named a Visionary in the 2026 Gartner® Magic Quadrant™ for Server Virtualization Platforms

"Gartner recognized Platform9 Private Cloud Director for its hardware reuse and in-place migration that avoids costly swing gear, plus native multi-tenancy, domains, tenants, regions, and self-service billing, for service provider and enterprise deployments"



We've Got Your Back

PCD Experts with You Every Step of the Way

One Global Team Behind Your Cloud

Reorganized in 2026 into a single Customer Experience organization:

One Organization

- **Support:** 24/7 production support and proactive operations
- **Solutions Architecture:** Design and deployment by senior architects
- **Professional Services:** Fixed-price implementation and migration delivery
- **Business Critical Services:** Subscription tiers for dedicated resources

Follow the Sun Coverage

- **Americas:** United States and Latin America. In-region, shared time zone, Spanish-capable engineers
- **EMEA:** Bucharest, Romania, with India on EMEA-shift coverage. Daytime coverage across Europe, the Middle East, and Africa
- **APAC:** Pune and Bangalore, India teams. Daytime coverage across Asia Pacific with a deep engineering bench

Nine principles. Every case. Every customer. Every time.

Every interaction with our team is a **moment of truth**: an opportunity to give you something you did not have before.

OWNERSHIP

Take ownership. Tell the truth. Keep promises.

EMPATHY

Understand your point of view. Reflect your urgency.

ENGAGEMENT

Over Communicate. Make every interaction add value.

Navigating Uncertainty

The Reality of the Broadcom Disruption



The Riskiest Move, Is No Move

- **Rising Costs:** VMware costs are outpacing budgets, eating into funds meant for other priorities
- **Technical Debt:** Legacy technology creates administrative debt and automation dead-zones
- **Compounding Dependencies:** Storage, networking, and tooling dependencies complicate future moves — worsened by rising memory costs and longer hardware lead times

VMware Changed Course

Risk abruptly followed

- **Financial Uncertainty:** Forced bundles, reduced flexibility, & rising contract costs.
- **Support Uncertainty:** Rapid SKU retirements
- **Technology Uncertainty:** Few true alternatives with full enterprise parity



Broadcom removes public access to VDDK



Public access to VDDK has been removed

The default mechanism vJailbreak uses to migrate live VMs is no longer publicly available.



Only available to select VMware partners

Alliance and backup vendors have access — Platform9 is not currently on that list.



Multiple paths forward

See the supported migration options on the next slide.

Three supported paths to migrate off VDDK



VDDK + NBD

- Needs VDDK: manual setup required
- Platform9 customers: ask in #vjailbreak for access
- Partners: cache and reuse your VDDK copy



Storage Assisted

FASTEST OPTION

- Cold migration: safe, structured host-based transfer
- Fully compatible with iSCSI and FC
- NFS environments not yet supported



Accelerated Copy


- Cold migration: direct connection to ESX and storage, no NBD
- Warm migration coming soon
- Works across all storage arrays

Architecture




"If you think good architecture is expensive, try bad architecture." — Brian Foote and Joseph Yoder

PCD Architecture

SaaS-hosted or self-managed control plane · KVM/Ubuntu hypervisor · OVN networking · Cinder storage



PF9 Agents ↓



Outbound port 443 only · No inbound firewall rules required

PCD-K: Kubernetes, 100% Managed

Zero control plane nodes to manage · Same management plane as PCD-V · CAPI-based, self-healing, auto-scaling.

ZERO CONTROL PLANE NODES


0

*control plane nodes
for you to manage*

- Automated worker node management
- Self Healing
- Auto Scaling — reduces waste
- Improves uptime
- Simplifies upgrades
- 100% managed by Platform9

Platform9 Management Plane · Outbound 443

WORKER NODES · Automated lifecycle — provision, scale, heal, upgrade



CAPI-based · GPU / AI ready · Bare metal K8s · Shared networking & storage with PCD-V · One management plane



Private Cloud Director: Adoption

The Familiar & Flexible Private Cloud

VM Migration Isn't Simple... But It Is!

Every dimension of complexity — handled. Point and click.

Scale	Operating Systems	Virtualization Change	Driver Change
1,000s of Virtual Machines	- Many versions of Linux - Many versions of Windows	ESX to KVM	- Paravirtual devices - Storage devices - Network devices
Types of Migration	Network Mapping	Complex VMs	Complex Tooling
Cold or Warm	Different virtual networks on source and destination	- Multiple NICs - Multiple Disks	- Terraform - Ansible

vJailbreak handles all of this — point and click, schedulable, live copy, cutover configurable.

VMware Migration Made Easy


Introducing vJailbreak

- Migrate in **weeks**, not months
- Live or cold cluster-level migration
- Reduce migration cost and resource demand
- Built by VMware-native engineers




Non-Disruptive Migration from VMware

In place upgrades



Large datasets, moved at array speed

vJailbreak orchestrates the cut-over. Your storage array does the data copy — using the same snapshot and replication APIs it runs every day.



Why it's fast

- Your storage array already copies its own blocks every night — for backup, snapshots, and replication. vJailbreak just points it at the destination.
- Data moves array-to-array. It never traverses our hypervisors, so the host CPU and network aren't the bottleneck.
- Because PCD doesn't bring its own storage stack (unlike Nutanix), nothing gets in the way of the array doing what it's already good at.

PROOF POINT

3 hours → 90 seconds


A VM migration that took 3 hours over the host data path completed in 90 seconds with storage-assisted migration. The only thing that changed was the data path.

— Tata Cloud, 2026

NETWORKING and Storage


Familiar primitives. Familiar mental model. No new playbook for your team.

Interconnectivity




PCD Private Network on Geneve — East-West Packet Walk

Down to the physical wire: VM → br-int → Geneve VTEP → bond0 → physical NICs → ToR fabric



What actually crosses the physical wire (between the two bond0 IPs)



Packet walk


1. VM1 → br-int/OVN resolves VM2 as remote; hands the frame to the local Geneve VTEP.
2. VTEP encapsulates it: outer Eth/IP (bond0 → bond0), UDP/6081, Geneve header with VNI 5001.
3. bond0 load-balances the encapsulated packet onto eth0 or eth1 per the LACP hash.
4. ToR switches forward on the outer IP/UDP only — they never see VM1, VM2, or VNI 5001.
5. HV2's bond0/VTEP decapsulates by VNI; br-int delivers the original frame to VM2.

Networking: ESX vs. PCD

Capability	VMware ESX / NSX	PCD (OVN)
Switching	vSwitch / vDS portgroups	Open vSwitch — multi-tenant, distributed
VLAN Support	Portgroup per VLAN	Provider VLAN trunks — 802.1Q native
Overlay / Tunnel	NSX VXLAN/Geneve (add-on)	OVN Geneve — built in, no extra license
DHCP / IPAM	Manual or NSX add-on	Built-in DHCP + IPAM (Infoblox integration)
Security Groups	NSX Distributed Firewall	Per-port stateful filtering — OVS dataplane
Load Balancing	NSX LB (add-on)	Octavia L4/L7 — HA, health checks, TLS
IP Spoofing Prevention	Port binding (manual)	Fixed IP per port — OS-independent
Routing	NSX Edge / Distributed Router	OVN distributed gateway — E-W and N-S
BGP Peering	NSX-T Edge BGP	H2 2026 (native) · VyOS bridge available now


vSphere Storage - I

VMFS Based



vSphere Storage - I


Snapshots



- Copy on Write disks created
- VM writes to the current delta disks
- Can have a tree of snapshots
- Incur a penalty on read
 - look at the current delta disk, if not get to the parent disk etc.

vSphere Storage - 2


RDM and vVol Based



- RDM
 - Manual configuration of the LUN connected to VM
 - Snapshot supported only on virtual RDM
 - VMFS
- vVol
 - Automatic configuration of LUN creation for individual VM
 - VASA Provider for storage providers
 - Storage Provider responsible for
 - deduplication
 - thin provisioning
 - other capabilities like replication

PCD Storage


NFS



- Similar to NFS or VMFS
- (optional) Individual Storage provider driver like Tintri, NetApp
- Similar to VMware, the disks are files
- Use QCOW2 image - equivalent to VMDK
- Can be thin or thick provisioned

PCD Storage

iSCSI, FC



- iSCSI/FC
- Similar to vVol
- Individual Storage provider driver
- Each VM gets its own LUN
- Storage provider responsible for
 - thin provisioning
 - deduplication
 - advanced svc - replication
- Snapshots handled by Storage Provider

Storage: VMware vs. PCD

Capability	VMware	PCD (Cinder)
Primary disk format	VMDK	qcow2 (thin or thick)
File-based datastore	VMFS / NFS	NFS via Cinder driver
Block storage (LUN per VM)	vVol	iSCSI / FC via Cinder driver
Raw disk mapping	RDM	Cinder volume (direct LUN)
Snapshots	VMFS delta disks / vVol snapshots	Cinder snapshots (per-volume)
Storage vMotion	Yes — live disk migration	Cinder volume migrate
Thin provisioning	Yes	Yes — qcow2 native
Vendor drivers	VAAI, VASA, vVol providers	Pluggable Cinder drivers
Supported arrays	Broad HCL	NetApp · Pure · HPE · Dell · Ceph

During vJailbreak migration: VMDK → qcow2 converted automatically. RDM → Cinder LUN. Networks re-mapped and IPs preserved.

Networking

T0, T1, segments — the model you already know, on open infrastructure

Three ways to do networking on PCD

Pick the trim that matches your environment. Mix and match across tenants — they share the same management plane.

Simplified

L2 networks — vSphere vSwitch + Port Groups, on KVM

Pure L2 plumbing. PCD provides the virtual switch and VLAN segmentation. You keep your existing DHCP, IPAM, routing, and firewall.

- VLAN-tagged provider networks (802.1Q)
- Maps 1:1 to vSwitch + Port Groups
- Distributed L2 across hosts (VXLAN-backed)
- External DHCP / IPAM (Infoblox, NetBox, etc.)

BEST FOR: VMware shops with existing VLAN topology, IPAM, and routing. Keep what works — plug PCD in.

Software-Defined

PCD OVN — multi-tenant SDN

Tenant overlay networks, distributed routing, per-port security groups, and Octavia load balancing. The full PCD SDN.

- Geneve tenant overlays + distributed E-W routing
- Floating IPs, SNAT/DNAT, security groups
- Octavia L4/L7 load balancing
- BGP, FWaaS, VPNaaS on the roadmap

BEST FOR: Multi-tenant private cloud, dev/test isolation, and the modern operating model.

External Edge

VyOS-managed border routing

VyOS deploys as a per-tenant virtual border router on PCD. Adds the NSX-Edge-style services OVN doesn't have yet.

- Full BGP with ECMP and route policies
- IPsec / WireGuard / OpenVPN termination
- Zone-based stateful firewall
- OSPF, multi-WAN, Q-in-Q stacked VLANs

BEST FOR: BGP-centric NSX-T exits, multi-WAN routing, and tenants that need advanced edge services today.

ALL THREE ARE FIRST-CLASS. Same management plane. Same APIs. Same support. Mix and match across tenants — a single PCD region can run all three.

What each model gives you today

Available today ✓

Coming on the roadmap H2 2026 · H1 2027

Capability	Simplified	Software-Defined (OVN)	External Edge (VyOS)
CORE L2 / L3			
Provider VLAN trunks (802.1Q)	✓	✓	✓
Tenant overlay networks (Geneve)	—	✓	—
Distributed East–West forwarding	L2 only	✓	works alongside
Built-in DHCP / IPAM	external	✓	✓
NORTH-BOUND CONNECTIVITY			
Floating IP / SNAT / DNAT	—	✓ basic	✓ advanced
North-bound BGP peering	via fabric	H2 2026	✓
IPsec / WireGuard VPN	—	H1 2027	✓
OSPF / multi-WAN routing	—	—	✓
SECURITY & SERVICES			
Per-port stateful security groups	—	✓	per-tenant
Zone-based stateful firewall	—	H1 2027	✓
Load balancer (L4 / L7 Octavia)	—	✓	—
ADVANCED			
Multi-region overlay (OVN-IC)	—	H2 2026	—
Q-in-Q (802.1ad) stacked VLANs	—	—	✓
Shared VM + Kubernetes networking	✓	✓	✓

WHY PCD TODAY: All three trims ship today (Simplified GA April 2026). Same management plane, same APIs, same support. Mix trims across tenants — and get BGP, IPsec, and stateful firewall today via the VyOS edge while OVN matures.

BGP Support

"Does PCD support BGP?" · Short answer: today via VyOS bridge, natively in H2 2026.

TODAY · OVN

- Distributed L2/L3 routing
- Provider VLAN trunks (802.1Q)
- Floating IPs / SNAT / DNAT
- Security groups (per-port)
- Octavia L4/L7 load balancing
- DHCP + IPAM (Infoblox)

BRIDGE · VYOS (NOW)

- Full BGP with ECMP
- Route policies & filters
- IPsec / WireGuard VPN
- Zone-based stateful firewall
- OSPF, multi-WAN routing
- Q-in-Q stacked VLANs

ROADMAP · H2 2026

- Native BGP via OVN/FRR
- Distributed gateway scale
- OVN overhaul — perf & scale
- OVN-IC multi-region networks
- VPN-as-a-Service (Futures)
- Firewall-as-a-Service (Futures)

VyOS deploys as a per-tenant virtual border router on PCD. It is a supported bridge — not a workaround — while native OVN BGP matures.

PCD networking, powered by OVN

Native, multi-tenant SDN delivering NSX-class capability on commodity hardware — with known gaps blocking some BGP-centric NSX-exit deals.

Tenant overlay networks

Geneve-tunneled L2/L3 networks per project with full multi-tenant isolation.

Distributed routing

OVN distributed gateway routers for east–west and north–south traffic.

Floating IPs & NAT

SNAT, DNAT, and floating IP allocation against external provider networks.

Security groups

Per-port stateful filtering enforced in the OVS dataplane.

Load balancing (Octavia)

Tenant-managed L4/L7 LBs with HA, health checks, and TLS termination.

DHCP & metadata

Built-in DHCP, IPv6 RA, and metadata service for cloud-init workflows.

Provider VLAN trunks

Direct attachment to existing physical VLANs — no re-IP, no fabric changes.

VM & K8s parity


Same SDN serves VM workloads and PCD-managed Kubernetes clusters.

WHAT'S MISSING TODAY

No native BGP peering · no IPsec/WireGuard VPN termination · basic firewall only (no NSX DFW-style microsegmentation) · OVN scale & error reporting limits.

VyOS as the NSX-T T0 / T1 replacement

Two-tier routing on PCD — BGP edge cluster = T0, per-tenant routers = T1, OVN overlays = NSX segments.



PCD OVN vs. VyOS — capability map

VyOS deploys as a per-tenant virtual border router on PCD — delivering the NSX-Edge-style BGP, VPN, NAT, and firewall capabilities OVN doesn't have yet.

Available today ✓ Coming on the roadmap H2 2026 · H1 2027 Not on roadmap ✗

Capability	PCD OVN	VyOS Edge
Tenant L2 networks	✓ Native — Geneve overlays	<i>Works alongside</i>
Tenant L3 networks (logical routers)	✓ Native — distributed routers	<i>Works alongside</i>
Distributed East–West routing	✓ Native	<i>Works alongside</i>
Provider VLAN trunks	✓ Native 802.1Q	✓ 802.1Q
Q-in-Q (802.1ad) stacked VLANs	✗ <i>Not on roadmap</i>	✓ Native S-VLAN / C-VLAN
Floating IP / SNAT / DNAT	✓ Basic	✓ Advanced — PAT, hairpin
Security groups	✓ Per-port stateful	<i>Per-tenant only</i>
Load balancing (Octavia)	✓ L4 / L7	—
North-bound BGP peering	H2 2026 <i>via OVN/FRR</i>	✓ Full BGP, ECMP, route policies
IPsec / WireGuard VPN	H1 / H2 2027 <i>VPNaaS</i>	✓ IPsec, WireGuard, OpenVPN
Stateful zone-based firewall	H1 2027 <i>FWaaS</i>	✓ Zone-based, deep inspection
Port forwarding (advanced)	✗ <i>Basic only — not on roadmap</i>	✓ Full DNAT / SNAT rules
OSPF / multi-WAN routing	✗ <i>Not on roadmap</i>	✓ OSPF, multi-WAN, ECMP

THE BRIDGE TODAY: VyOS is a supported per-tenant virtual border router — not a workaround. Use it where you need BGP, VPN, or stateful firewall today. OVN closes the gap on roadmap dates above.

RBAC

Today: project-scoped RBAC. Tomorrow: the vCenter role model your admins already know.

Role-Based Access Control

"You only have two roles." · We hear you. Here's where we are and where we're going.

TODAY	H2 2026 · BETA	2027 H1 · GA
3 Roles • Admin — full platform access • Self-Service — project-scoped user • Read Only • Project isolation for team separation • SAML / LDAP / SSO integrated	Granular RBAC Beta • Tenant Admin role • Domain Admin role • Per-resource permission scoping • Available to design partners	Full RBAC GA • Custom role definitions • Fine-grained resource permissions • Equivalent to vCenter role model • Full multi-tenant enforcement

TODAY'S WORKAROUND: Use separate PCD projects per team to achieve resource isolation while granular roles are on the roadmap. RBAC at the project boundary is fully enforced today.

Granular RBAC + ABAC — the target role model

OPA-backed policy engine · Custom roles as first-class citizens · LA July 2026 · GA Q1 2027

Role	Scope	Key Capabilities	Status
Administrator	Domain	Full IAM · all domains, tenants, infra, services, quotas	GA
Self-Service User	Tenant	Create/manage VMs, volumes, snapshots, console, power ops	GA
Read-Only User	Domain / Tenant	View resources and policy bindings · no writes	GA
Domain Admin	Domain	Admin access across all tenants in the domain	Phase 1 · Apr '26
	Domain	Admin for network config; self-service user elsewhere	Phase 1 · Apr '26
	Domain	Admin for storage config; self-service user elsewhere	Phase 1 · Apr '26
IAM Administrator	Domain	Manage roles & policies — no infra access · Policy Analyzer	LA · Jul '26
Tenant Administrator	Tenant	Project-scoped policy bindings · assign roles to users/groups	LA · Jul '26
Custom Roles	Domain / Tenant	Customer-defined permission bundles · ABAC conditions	LA '26 → GA '27

PHASE 1 · APRIL 2026 · CUSTOM ROLES

3 new predefined roles + domain/tenant-level custom roles · OPA on core services (Keystone, Nova, Neutron, Cinder, Glance, Resmgr)

LA JUL 2026 → GA Q1 2027 · SELF-SERVICE RBAC + ABAC

IAM Admin UI · permission catalog · policy inheritance · resource-level RBAC · ABAC conditions · platform-wide OPA

DISASTER RECOVERY

If you ran SRM, you already know this workflow — bridged today with your arrays, native by 2027 H1.

Disaster Recovery — Replacing SRM

"We rely on SRM — what replaces it?" · Transparent answer: partners today, native orchestration in 2027 H1.

BACKUP / RESTORE · NOW

- Veeam — integration GA H1 2026
- Rubrik — integration GA H1 2026
- Commvault — GA H1 2027
- Cohesity — GA H1 2027
- Per-VM backup & restore
- RPO driven by backup schedule

BRIDGE · ARRAY REPLICATION

- Storage array replication (NetApp, Pure, HPE, Dell)
- Array APIs drive replication (sync to seconds RPO)
- P9-assisted recovery orchestration (scripted)
- Networks pre-created on target site
- Cinder manage — import replicated volumes
- VM sequencing and group restart controlled

ROADMAP · 2027 H1

- Native DR Orchestrator — GA
- Recovery plans & runbooks
- Non-disruptive failover testing
- Array replication automation (framework)
- Failback workflow
- Block storage HA active-passive

The array replication bridge uses your existing storage investment — no new tools required. P9 orchestration scripts the failover sequence. Native DR Orchestrator (2027 H1) productizes this into recovery plans, runbooks, and one-click failover.

PCD DR — what we have today

Strong infrastructure primitives — but no SRM-equivalent orchestration. That's the gap we have to close.

VM snapshots

Per-VM crash-consistent snapshots via Cinder (block) or QEMU (ephemeral).

VM HA + DRR

In-cluster restart on host failure; resource rebalancing across hypervisors.

Live migration

Zero-downtime VM moves within a region for planned maintenance.

vJailbreak migration

VMware → PCD migration tool, including storage drivers for major arrays.

Multi-region PCD

Independent PCD regions can run in parallel sites; manual workload placement today.

Cinder storage backends

Pluggable drivers for NetApp, Pure, HPE, PowerMax, Ceph — replication APIs available per array.

Image library

Glance image distribution across regions; backbone for standby VMs at the recovery site.

Backup partners (early)

Veeam + Rubrik integrations land in 2026 H1; Commvault, Cohesity in 2027 H1.

WHAT WE'RE MISSING TODAY

No DR orchestrator · no recovery plans / runbooks · no non-disruptive failover testing · no array-replication automation · no failback workflow.

Use Cases


Stretch Cluster - Single PCD

- Single PCD (OpenStack) deployment
- Two Hitachi Vantara arrays connected via GAD
- Preservation of the volume and ids
- Synchronous replication
- Distance limitation - 30 miles ?


Disaster Recovery - Multiple PCD

- Multiple PCD (OpenStack) deployment
- Two Hitachi Vantara arrays connected via HUR
- Preservation of the volume and ids
- Async - Synchronous replication
- No Distance limitation

Stretch Cluster - with one PCD SaaS/Remote hosted




Disaster Recovery (Self Hosted) Dual OpenStack/PCD



Disaster Recovery — PCD-DR Appliance

Two independent PCD clusters · Async storage replication · PCD-DR Appliance orchestrates failover mapping.



DR Appliance: holds source + destination metadata · **VM metadata present**, infrastructure is not · **Infra pre-created:** ClusterBlueprint + networks must exist at target

DR and Backup — what you get, when

Two different problems, two different answers. DR is failover — backup is restore. Here's the honest map of both.

DISASTER RECOVERY · failover when a site goes down			
Capability	PCD today	PCD + Array Replication (bridge today)	Native DR Orchestrator (2027 H1)
RPO (recovery point)	<i>Manual snapshots</i>	✓ Seconds — array sync	✓ Seconds — orchestrated
RTO (recovery time)	<i>Manual</i>	✓ Minutes — P9 scripts	✓ Minutes — one-click
Site-to-site failover	—	✓ Scripted	✓ Native runbook
Non-disruptive failover testing	—	—	✓ Test bubble
Reverse replication / failback	—	✓ Scripted	✓ Native
Network re-mapping at recovery site	—	✓	✓
Array-replication integration	<i>Cinder drivers available</i>	✓ NetApp · Pure · HPE · Dell · Hitachi	✓ Automated
Stretch cluster (intra-region HA)	✓ Today	✓ Today	✓ Today
Time to deploy	<i>Day 1</i>	<i>Weeks · uses arrays you already own</i>	<i>Day 1 at GA</i>
BACKUP & RESTORE · point-in-time recovery for VMs and apps			
Veeam GA · H1 2026 Per-VM backup, VSS app-consistent, broad enterprise install base. ✓ AVAILABLE NOW	Rubrik GA · H1 2026 Per-VM backup, immutable snapshots, ransomware-resilient catalog. ✓ AVAILABLE NOW	Commvault GA · H1 2027 Per-VM backup, app-consistent across DB workloads, deep compliance tooling. → COMING	Cohesity GA · H1 2027 Per-VM backup, search and recover, data security integrations. → COMING

THE HONEST READ: For DR — use your array replication + Platform9 orchestration today; native DR Orchestrator lands 2027 H1. For backup — Veeam and Rubrik are available now; Commvault and Cohesity follow in H1 2027.

The case for PCD

Familiar for vAdmins

HA, DRS, live migration, SDN — the primitives your team already knows, on open infrastructure.

Honest Roadmap

We tell you what's GA, what's in beta, and what's coming. BGP, RBAC, DR — all on the map with dates.

Always-On Assurance

SaaS control plane. P9 monitors and maintains it. You run workloads, not control plane operations.

Full Stack: VMs + K8s

PCD-V and PCD-K on one management plane. One operational model, one support contract.

vJailbreak — Only In-Place Migration

Convert entire VMware clusters without hardware refresh. \$10–80/VM. Siemens EDA migrating 45K VMs.

Proven at Scale


24,000 hypervisors, >1M CPU cores, 10 years in production. 82% ARR growth since PCD launch.

How to start with Platform9 via Crayon

1. Find Platform9 under Programs

2. Review the Program Guide

3. Sign up to Platform9!





Q&A

Scan for Feedback &
To Book A Meeting



Crayon
A SoftwareOne Company



PLATFORM9