

BLOCK STORAGE COST OPTIMIZATION ON AZURE AND AWS

Amoeba product guide

How to onboard a VM, what Amoeba changes, and how you stay in control.

• What Amoeba does

Moves each VM's data disks onto one btrfs pool sized to your data, with room to grow. Then it grows and shrinks the pool with your data, so you pay for close to what you use.

• What Amoeba never does

Look inside a VM you haven't onboarded, remove a disk that still holds data, or delete a source disk before its snapshot is complete.

• You stay in control

You pick each VM. Per VM, you choose whether shrinks run on their own or wait for you.

Onboarding starts the move, and services on the data mounts restart at the switch. See section 6. Updated 27 October 2026.

- | | | | |
|----|-------------------------|----|---------------------------|
| 01 | How Amoeba works | 02 | Before you start |
| 03 | Onboard a VM | 04 | What gets installed |
| 05 | The console | 06 | Moving a VM onto its pool |
| 07 | Expanding and shrinking | 08 | Stay in control |
| 09 | Savings and figures | 10 | Security and data |
| 11 | Remove Amoeba from a VM | | |

SECTION 1

How Amoeba works

Amoeba works on the data disks of running VMs: Azure managed disks and Amazon EBS volumes, on Windows and Linux VMs. It never touches the OS disk or root volume.

1 Connect

Connect your Azure directory once. A script you run in Azure Cloud Shell gives Zoorik's app a custom role on the subscriptions you choose.

On AWS, enter your 12-digit account ID, then run a script in AWS CloudShell that creates the IAM role ZoorikStorageOperator. It trusts Zoorik's AWS account only with your organisation's external ID.

2 Onboard

Pick the VMs Amoeba manages. It installs its agent on each one and leaves every other VM alone.

3 Move onto a pool

Amoeba copies your data onto one btrfs pool, sized to your data with room to grow, while your services run. Then it switches each mount to the pool, with a short restart of the services that use it. Once the switch is checked, it snapshots and deletes each old disk.

4 Optimize

Amoeba grows the pool before it fills, and sooner when data arrives fast. The pool stays mounted while it changes. When your data goes down and stays down, it shrinks the pool one disk at a time. So you pay for close to what you use.

Two stages on every VM

Stage	What happens	How it starts
The move onto a pool	Copies your data onto the pool while your services run, then switches each mount to it. Once the switch is checked, it snapshots and deletes each old disk.	By itself, after the agent's first report. A move only goes forward.
Autopilot	Grows the pool before it fills. When your data goes down and stays down, it shrinks the pool one disk at a time.	The pool can grow as soon as it is built. Shrinks run on their own in "Auto", or when you start them in "Manual".

RELATED [Can you shrink an Azure managed disk?](#)

SECTION 2

Before you start

Check these before you onboard a VM.

You need	Details
A cloud account	Azure: a Microsoft Entra directory (tenant) and the subscriptions that hold your VMs. AWS: an AWS account whose EC2 instances are managed by AWS Systems Manager, with the SSM agent connected. Each must be a Nitro type that exposes EBS as NVMe, with at least one EBS data volume.
The rights to connect	In Azure, the rights to add Zoorik's app to your directory, create a custom role, and assign it on each subscription.
A running VM with data disks	Amoeba lists only running VMs with at least one data disk. The data on those disks must sit on mounted filesystems.
Data-disk slots	During the move, the pool's disks attach beside the old disks. So the VM needs a few free disk slots for the pool. If it has too few, Amoeba declines it and says why.
Linux	The kernel must be recent and support btrfs, built in or as a module, and services on data mounts must run under systemd. The agent also needs rsync, btrfs-progs (mkfs.btrfs and btrfs) and psmisc (fuser), which the installer doesn't add, and cgroup v2 for its disk-load readings. Amoeba checks the running kernel, never the OS name.
Services on data mounts	Every process that keeps a data mount open must run as a systemd service. Then the agent can stop it and start it again at the switch.
Java	The agent runs on Java 17 or later. If the VM has none, the installer adds OpenJDK 17 with apt, dnf, yum or zypper.
Network	Outbound access from each VM to Zoorik for the agent's reports and command channels, and to the hosts it downloads from.
A console account	Sign in with a Microsoft or Google account.

Onboard a VM

You connect Azure once. Then you pick the VMs Amoeba manages, in the console.

Review it first

Onboarding is your go-ahead. The move runs by itself, and services on the data mounts restart at the switch.

1 Open the console

Sign in with a Microsoft or Google account.

2 Connect Azure

Click "Connect Azure" and enter your directory (tenant) ID. The console shows a script to run in Azure Cloud Shell. It adds Zoorik's app to your directory, creates the role "Zoorik Storage Operator", and assigns it to the app on each subscription you name.

3 Verify

Click "Verify". The connection reads "Verifying" until Zoorik can read a subscription through the role, and the console keeps checking. Run Verify again later to add subscriptions you granted since.

4 Onboard VMs

Click "Onboard VMs". The list shows each VM Amoeba can onboard, with its resource group, region, size, and the number and total size of its data disks. Tick the VMs you want, then click "Onboard".

5 The agent installs

Amoeba installs its agent on each VM through Azure Run Command. The VM's agent reads "Installing", then "Connected" once its first report arrives.

6 The move starts

After the first report, Amoeba plans the pool and starts moving the VM onto it, with no further click.

Which VMs are listed

- The list comes from Azure Resource Graph only. Amoeba doesn't look inside a VM, or show its used space or savings, until the VM's agent reports.
- VMs that can't be onboarded are counted in one line under the list, with the reason. For example, the VM is stopped or has only an OS disk.
- The list can be a little behind Azure.
- A VM you onboarded before keeps its settings when you onboard it again.

Connection states

State	What it means
Verifying	Zoorik can't read a subscription yet. The console says why: the app isn't added to the directory, the role isn't assigned, or the directory wasn't found.
Connected	Zoorik reads your subscriptions through the role.
Access lost	Zoorik lost read access to a subscription. Figures that come from Azure keep their last value, dated, and onboarding waits. Zoorik checks access again regularly.

If the install fails

- The console shows why the install failed: the script failed, Azure refused the run, or the install stalled.
- It also says when Zoorik can't reach the subscription, or when the agent installed but its first report didn't arrive in time.
- Amoeba tries again, waiting longer each time, for as long as the VM stays onboarded.
- "Reinstall" in Settings starts a new install straight away, or when the current step of the move ends.

What gets installed

Amoeba adds one agent to each VM you onboard, and creates its own disks and snapshots in your subscription.

On the VM

Item	What it is
amoeba-agent systemd service amoeba-agent.service	Reports the VM's disks and mounts to Zoorik, and carries out the move, the expands and the shrinks. Runs as root from /opt/amoeba-agent, with IO accounting on.
/etc/amoeba-agent.env	The agent's settings: its own key, Zoorik's address and the VM's Azure resource ID. Readable by root only. The install command carries a single-use bootstrap token, never the key. The installer exchanges the token with Zoorik for the agent's own key. A reinstall replaces the key.
Java	OpenJDK 17, added only when the VM has no Java 17 or later.
Other packages	curl, ca-certificates and procps, added when the VM lacks them.
btrfs kernel module	The agent loads it when it isn't loaded yet.
The pool	One btrfs filesystem across the pool's disks. Each data mount becomes a subvolume on it, mounted at the same path as before.
/etc/fstab	At the switch, the agent saves a copy as /etc/fstab.amoeba-backup.<id> and points each data mount at the pool.

In your Azure subscription

Resource	What it is
Pool disks zoorik-amoeba-pool-...	The pool's managed disks. Amoeba adds, grows and removes them as the pool expands and shrinks. What Amoeba creates carries tags that start with amoeba-.
Safety snapshots zoorik-amoeba-snap-...	One incremental snapshot of each source disk, taken before the disk is deleted, in the disk's own resource group. Deleted after a few days.
Zoorik Storage Operator	A custom role in your directory, assigned to Zoorik's app on each subscription you connect. The connect script also adds the app to your directory.

In your AWS account

The IAM role ZoorikStorageOperator, the pool's gp3 volumes, and an EBS snapshot of each source volume, kept for a few days.

Updates

- Amoeba updates the agent to each new release by itself.
- If a new version fails its health check, the installer restores the previous one.

SECTION 5

The console

Overview lists every VM you onboarded. Pick a VM to see its pages.

Page	What it is for
Overview	Every onboarded VM with its status, agent, disks, capacity, spend and savings; today's disk spend, what is recoverable, and what Amoeba has saved to date. "Needs you" lists paused moves, shrinks waiting for you, and agents to reinstall. "Connect Azure" and "Onboard VMs" start here.
Virtual machine	
Dashboard	The VM's status and capacity, and what its disks cost: today and with a pool, or before and now. Also its recent used space, provisioned space and IOPS.
Disks	Every data disk with its LUN, pool and monthly cost, and how many data-disk slots are in use.
Pool	"Overview": the pool's disks, size, use, IOPS, throughput and spend, or the plan before the move. "Migration": the move's phases and copy progress, with "Resume" when a paused move can run again.
Operations	Expands and shrinks running now, and a shrink waiting for you in "Manual". A log of recent expands and shrinks, with each outcome and its reason.
Impact	Before and now, what the pool has saved since it went live, and recent capacity and spend.
Settings	"When Zoorik shrinks", the agent's status with "Reinstall", and the VM's facts.

Status words

Status	Meaning
Pool onboarding	The move is running, waiting for the agent, or paused. The line says which phase.
Ready for a pool	The agent has reported. The move starts by itself.
Expanding	An expand is running.
Shrinking	A shrink is running.
Idle	The pool is live with no expand or shrink running, or a shrink is waiting for you.

Every figure is a value, or a dash that says why it is missing. The dash reads, for example, "Available once the agent reports", "No Azure price for this disk" or "Agent not reporting". Amoeba never fills a gap with a guess.

SECTION 6

Moving a VM onto its pool

It starts after the agent's first report and runs by itself. Amoeba plans the pool, copies your data, switches over, checks the pool and cleans up.

How the pool is sized

- Amoeba sizes the pool to your data, with room to grow. It looks at your peak use and recent growth, not at the size the old disks were given.
- It picks the cheapest layout of disks that holds your data and carries your VM's disk load.
- Each VM gets one pool, on one disk type. On Azure the pool uses Premium SSD v2 where your region offers it, and Premium SSD otherwise.

- On AWS the pool uses gp3 volumes, in the instance's Availability Zone.
- Data on Standard SSD, Standard HDD or Ultra Disk also moves to Premium SSD v2 or Premium SSD. Per GiB, Premium disks can cost more than Standard ones. Zone-redundant (ZRS) disks map to a locally redundant (LRS) pool, so the data loses zone redundancy.

The phases

Phase	What happens
Check and plan	Amoeba checks the VM can take a pool, and sizes the pool to your data. Only the disks attached to this VM count as its source disks.
Build the pool	It creates the pool's disks and attaches them beside the old disks.
Copy	The agent copies your data onto the pool, inside the VM, while your services run. The copy slows down or waits when the VM's disks are busy.
Catch up	It copies what changed since, until the pool has caught up.
Switch	Each mount switches to the pool, with a short restart of the services that use it. See "At the switch" below.
Check	Amoeba checks the pool for disk errors after the switch. An error pauses the move.
Clean up	Amoeba snapshots each old disk, deletes the disk, and keeps the snapshot for a few days before deleting it too.

At the switch

- The agent stops every systemd service that holds a data mount open, with systemctl stop. It never kills a process itself.
- It checks nothing is left to copy, saves /etc/fstab, mounts the pool at each data mount's path, and starts the services again. Your mount paths stay the same.
- If any step fails, the agent tries to put the original mounts back and start the services again. The source disks are never written during the switch. Amoeba then tries the switch again, so your services can restart more than once.
- A process that holds a data mount open outside a systemd service can't be stopped cleanly, so the switch can't finish.

After the switch

- A move only goes forward. Once the switch is checked, Amoeba snapshots each old disk, deletes the disk, and keeps the snapshot for a few days before deleting it too.
- A disk is deleted only after its own snapshot succeeded, and only when its tag names this move.
- You pay for both sets of disks only while the move runs, and then for the snapshots while they're kept.
- The pool keeps growing before it fills. Shrinks wait until Amoeba has seen enough of your use after the move.

If the move pauses

- The move pauses when a phase can't go on safely. For example, the VM can't take a pool, a copy fails, or the outcome of the switch can't be known.
- Pool → Migration shows where it paused and why. "Resume" appears when the move can run again.
- A VM Amoeba declines, for example one with too few data-disk slots, shows why and stays on its own disks.

Expanding and shrinking

Amoeba grows and shrinks each pool with your data, while it stays mounted. So you pay for close to what you use.

When Amoeba expands

Amoeba grows the pool before it fills, and sooner when data arrives fast. It also adds room when btrfs needs space for its own metadata.

Expands run on their own, in "Auto" and in "Manual". The pool can grow during the move too, even while a move is paused.

How an expand runs

- The pool stays mounted while it grows, so your services keep running.
- Amoeba grows existing disks in place, or adds new ones while the VM has a free disk slot.
- When data arrives fast, it adds room sooner.

When Amoeba shrinks

When your data goes down and stays down, Amoeba shrinks the pool one disk at a time. A shrink waits until the move is done and Amoeba has seen enough of your use.

The new size follows the same rule as a new pool: sized to your data, with room to grow. Amoeba picks the cheapest layout that passes its space, performance, saving and safety checks.

In "Auto" the shrink starts by itself, when the pool is quiet. In "Manual" it waits for you on Operations.

How a shrink runs

Step	What happens
Check	Amoeba picks the disk to remove and checks the shrink is still safe.
Move the data	The agent moves the disk's data onto the other disks and takes the disk out of the btrfs pool.
Remove the disk	Amoeba checks the disk has left the pool. It then detaches and deletes the disk at once, with no snapshot. You stop paying for it.

A shrink removes one disk at a time. A large shrink is split into parts.

What stops a shrink

Reason	What it means
Room	No smaller layout would hold your data with room to grow.
Filling fast	The pool is filling too fast to move the data safely.
Disk load	The other disks couldn't keep up with the VM's disk load, judged from the peak of recent readings.
No readings	Amoeba is waiting for disk-load readings from the agent.
Kernel	The VM's kernel is too old, or doesn't support btrfs.
Disk errors	The pool reports disk errors.
VM size	The VM size isn't in Azure's size list for its region yet.

A running shrink stops when the pool needs room fast. Amoeba never detaches a disk while the pool is close to full.

Stay in control

Each VM has one setting: whether shrinks run on their own or wait for you.

Settings

Setting	Where	Default	What it does
When Zoorik shrinks	Settings → Shrinking	Auto	"Auto" runs each shrink by itself, when the pool is quiet. "Manual" holds it until you start it on Operations. Expands are automatic either way.
Reinstall	Settings → Agent		Installs the agent again. Offered when the agent stopped reporting or its install failed, while the VM is running.

Changes wait until you click "Save".

Shrinking by hand

- In "Manual", Amoeba keeps one ready shrink plan per pool, and keeps it up to date.
- Operations shows it as "Waiting for you": the disk it removes, what it saves, how full the pool is after, and how much data moves. Overview lists it under "Needs you".
- "Start shrink" checks again on fresh readings and removes exactly that disk. If the plan has changed, the card refreshes instead.
- If a check fails, the card reads "Can't shrink now" with the reason, and has no button.
- You can't stop a shrink once it starts. It stops only when the pool needs room fast.

Resume

"Resume" on Pool → Migration starts a paused move again, once it can run. If the switch's outcome can't be known, the move waits for you to click Resume. A move only goes forward.

A careful start

- 1 Run the free disk assessment first, to see what your own VMs could save.
- 2 Onboard one VM whose services run under systemd, and set "When Zoorik shrinks" to "Manual" on its Settings page.
- 3 Watch the move on Pool → Migration, and check your services after the switch.
- 4 When a shrink waits for you, start it yourself from Operations, and read its log entry.
- 5 Then switch to "Auto" and onboard more VMs.

Savings and figures

Every figure says where it comes from. Amoeba never shows a made-up number.

Figure	Where	What it means
Disk spend	Overview, Dashboard	What the VM's data disks cost a month now. During the move it counts both the source disks and the pool disks.
Recoverable	Overview, Dashboard	Before the pool goes live: what the planned pool would save a month.
Before, Now, Saving	Dashboard, Impact	After the move: what the source disks cost a month, what the pool costs, and the difference.

Figure	Where	What it means
Saved to date	Overview, Impact	What the pools have saved since they went live.
Saved by shrinks, Offset by expands	Operations	The monthly change from recent shrinks and expands, and the net.

- Prices are Azure's public list prices for the disk type and region, from the Azure Retail Prices API. A month is 730 hours.
- On AWS, prices come from the AWS Price List API. AWS charges EBS per GB-month.
- Standard SSD and Standard HDD are priced on capacity only, without transaction charges.
- Savings are measured against the source disks.
- A saving can be negative, for example when a pool has grown past the disks it replaced. It is shown as it is.
- Where Azure has no price for a disk, the figure reads "No Azure price for this disk". Until the agent reports, used space and savings read "Available once the agent reports".

RELATED [Free block storage cost assessment](#)

SECTION 10

Security and data

Zoorik acts in your Azure only through the role you assign. The agent opens every connection to Zoorik itself.

What the role allows

Action	Used for
<code>*/read</code>	Reading every resource in the subscriptions you assign it on, such as VMs, disks and sizes.
<code>Microsoft.Compute/virtualMachines/write</code> <code>Microsoft.Compute/virtualMachines/reapply/action</code>	Attaching and detaching disks, which Azure does as a VM update.
<code>Microsoft.Network/networkInterfaces/join/action</code>	The same VM update: Azure checks every resource linked to the VM, its network interface included.
<code>Microsoft.Compute/virtualMachines/runCommand/action</code>	Installing the agent, and checking the copy and the switch during the move.
<code>Microsoft.Compute/disks/write</code> <code>Microsoft.Compute/disks/delete</code>	Creating and growing pool disks, tagging source disks, and deleting pool disks and source disks.
<code>Microsoft.Compute/disks/beginGetAccess/action</code>	Needed to snapshot a disk. The download link it can create is never requested.
<code>Microsoft.Compute/snapshots/write</code> <code>Microsoft.Compute/snapshots/delete</code>	Taking safety snapshots, and deleting them after a few days.

The role is the limit on what Zoorik can do in your Azure. An action added later reaches your subscription only when you run the script again.

On AWS, Zoorik acts only through the IAM role `ZoorikStorageOperator`, which you create in AWS CloudShell. The role trusts Zoorik's AWS account only with your organisation's external ID. Zoorik's own AWS identity may only assume roles with that name and read AWS prices.

Area	The role allows
EC2 and Systems Manager	Read regions, instances, volumes, instance types and SSM status
Volumes	Create, modify, attach, detach, delete and tag
Snapshots	Create, delete and tag

Area	The role allows
Commands	Run AWS-RunShellScript through Systems Manager
KMS	Create grants, decrypt and re-encrypt, only through EC2, so encrypted volumes keep their key

Good to know

The role can create and delete disks and snapshots. Zoorik deletes a pool disk only when a shrink emptied it, or an expand that created it gave up.

It deletes a source disk only after its own snapshot succeeded. It deletes a snapshot only when its tag names its own move.

What leaves the VM

- In each report: each mount's path, filesystem type and ID, and its used and total size.
- With each mount: the devices, LUNs and cloud disk IDs behind it, and the systemd service that holds it open.
- Disk activity per device: IOPS, latency and throughput, and the 99th percentile of IOPS.
- The pool's sizes, and its btrfs error counters per disk.
- Host facts: the kernel version, whether btrfs is available, the disks' sector sizes, kernel features and the Azure Disk Encryption state.
- During a copy: its progress.
- While the VM has a pool: how busy the disks are, the copy's own reads and writes included.
- If an install fails: the agent's last log lines, with secrets removed.

The copy itself runs inside the VM, from the source disks to the pool disks.

How the agent connects

- The agent sends its reports to Zoorik, and keeps two channels open to it over WebSocket (wss). The reports and both channels use the agent's own key, from `/etc/amoeba-agent.env`.
- Zoorik signs every command with that key (HMAC-SHA256), and the agent checks the signature before it acts. A key from another VM can't act for this one.
- Zoorik reaches the VM through Azure Run Command, to install the agent and to check the move. Azure adds its Run Command extension to a VM that doesn't have it yet. It keeps a copy of each command and its output under `/var/lib/waagent`, and records each run in the activity log.
- On AWS, Zoorik reaches the instance through AWS Systems Manager (AWS-RunShellScript), which keeps each command's history in your account.

Encryption

New pool disks take the encryption key of the first encrypted source disk. That is your customer-managed KMS key on AWS, or your disk encryption set on Azure. Encryption at rest carries over.

VMs that use Azure Disk Encryption are not moved as they are.

Sign-in

Sign in with a Microsoft or Google account, as for Coral.

Remove Amoeba from a VM

A move only goes forward, so a VM stays on its pool. To stop Amoeba, wait until no VM is creating a pool. Then remove Zoorik's access and run `systemctl disable --now amoeba-agent` on each VM.

Review it first

Zoorik's access is the Zoorik Storage Operator role assignment or Zoorik's app on Azure, and the ZoorikStorageOperator role on AWS. While a pool is live, Amoeba needs its role to add disks. If you remove the role's assignment, the pool can't expand.

The source disks

A move only goes forward. Once the switch is checked, Amoeba snapshots each old disk and deletes it. It keeps each snapshot for a few days, then deletes it too.

What Amoeba added

- On the VM: the agent in `/opt/amoeba-agent`, the service `amoeba-agent.service`, `/etc/amoeba-agent.env`, the `/etc/fstab` backup, and the pool with your data on it.
- In Azure: the pool disks, the safety snapshots while they're kept, the role "Zoorik Storage Operator" and its assignments, and Zoorik's app in your directory.
- In AWS: the pool's gp3 volumes, the safety snapshots while they're kept, and the IAM role ZoorikStorageOperator.
- Your data stays on the `zoorik-amoeba-pool-...` disks, and nothing grows them once Amoeba stops, so keep them and watch their free space. Snapshots not yet deleted stay in your account.

Find your own number.

Run the free disk assessment on your own VMs.

zoorik.com/assessment →

Need help?

Pick a time at cal.com/zoorik/talk, or write to founders@zoorik.com.