

Proxmox and Ceph

Ditch the old SAN

Intelligent Systems Services

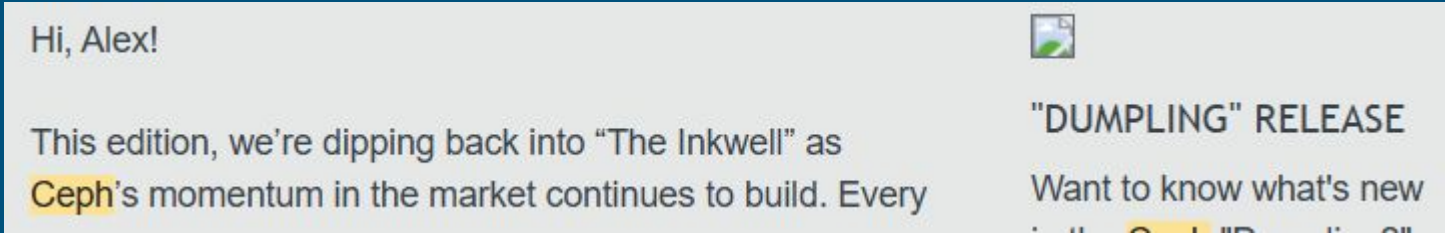


Alex Gorbachev

Intelligent Systems Services, Inc.

alex@iss-integration.com


Me




- Started Intelligent Systems Services in 1994
 - SGI, DEC Alpha, Sun, EMC, Oracle
- Sought out Ceph in 2013 (Ceph Days Santa Clara 2013)
- Built a lab in 2014 to test Ceph at production scale with SSD speeds
- Developed Storcium - "Out of the Way" Enterprise Management Integration Layer for Ceph
- Ceph Foundation Founding Member
- Proxmox Gold Partner with >200 customers as of May, 2026

Thanks Alex, you don't sleep.

My idea of a home lab in 2014...



...and today

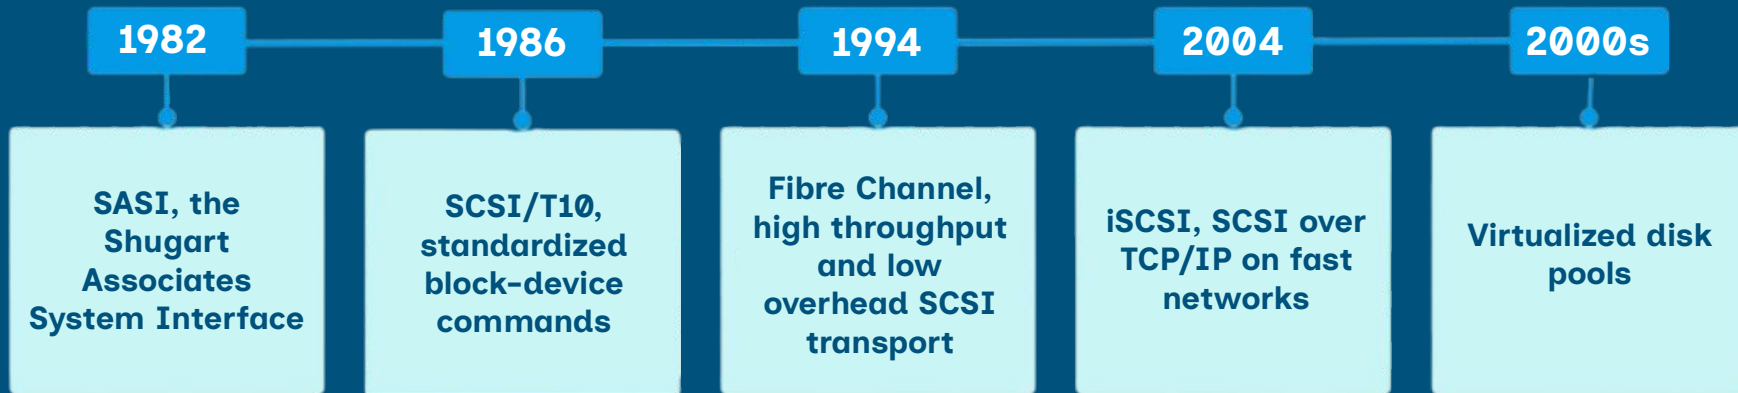


Proxmox Virtual Environment

- Integrated KVM Based Hypervisor Suite
 - Started 2005
 - AGPLv3, paid and free repos
 - Debian
 - Corosync
- Hyperconverged and External Ceph
- Also able to use iSCSI, Fibre Channel, DAS etc. as Shared Thick LVM
- Also able to use NFS as a shared file based backend



Brief History of SAN: Remote Storage as Local Devices



Designed for scarce, expensive disks, “pet” servers

Speeds have improved, delivery architecture has not changed significantly

SCSI is older than...

- World Wide Web (1989)
- Linux (1991)
- Web Browsers (1990)
- Text messages (1992)
- Smartphones (~2007)
- Wi-Fi (1997)
- **Not me**
- Ceph (2004-2005)

Introducing the Cromemco Personal Computer. Don't let the low price fool you.

Only
\$1785.



No fooling. Now you can have a complete full-featured personal computer, including software, for less than you'd expect to pay for an individual computer peripheral. And it's CP/M™ compatible.

Cromemco introduces the new C-10 SP. It's the perfect blend of an intelligent 12" CRT with a keyboard, a 5¼" disk

drive with 390K capacity, and three separate software packages. There's nothing else to add. Except, of course, a printer if you want professional word processing. And you get access to the entire line of other quality Cromemco products.

The Cromemco C-10 SP is more than a sophisticated home computer. It can serve as an executive workstation, for dis-


tributed data processing or as a front end for a mainframe computer. In fact, it fits in just about anywhere. For the C-10 SP is truly the professional personal computer for the serious user.

 **Cromemco**
A Division of Datacube, Inc.

Don't shoot SCSI...it's just a messenger*

- The real problem is the **Array!**

- Proprietary
- Constrained
- Bugs
- SPOF




vs.



* Also NVMe/TCP, NVMe-oF, Fibre Channel, iSCSI

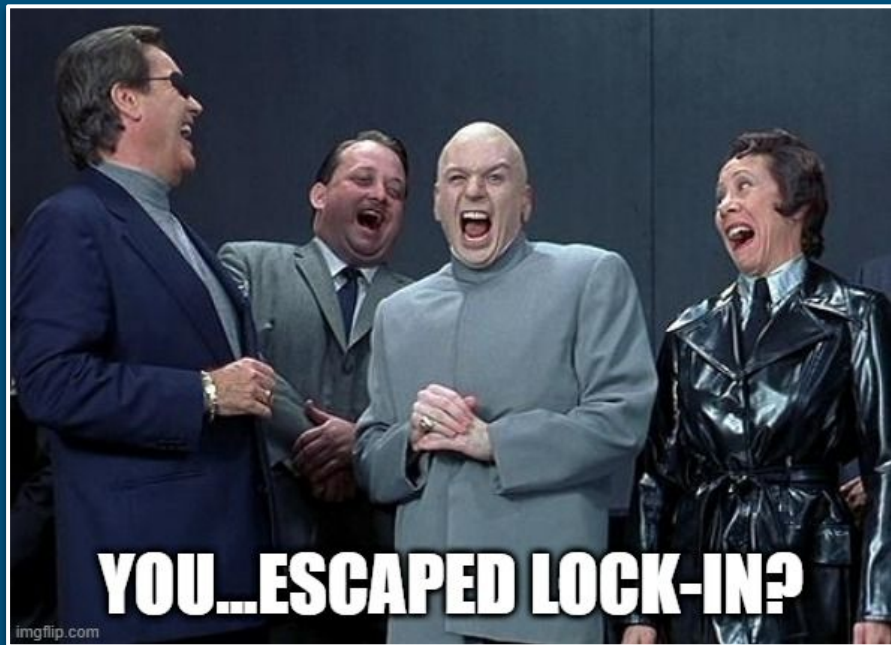
Backend: Advanced, but Opaque and Hidden

- Backend makes its own decisions
- Relocating data - limited options
- Only Vendor has Answers
- Controller is a choke point



Storage Theater: The Front End Trap

- Vendor UX Philosophy
- Proprietary Workflows
 - Obscures Storage Concepts
 - Groups Unrelated Objects
- What Standards?
 - Favors Preferred Compute

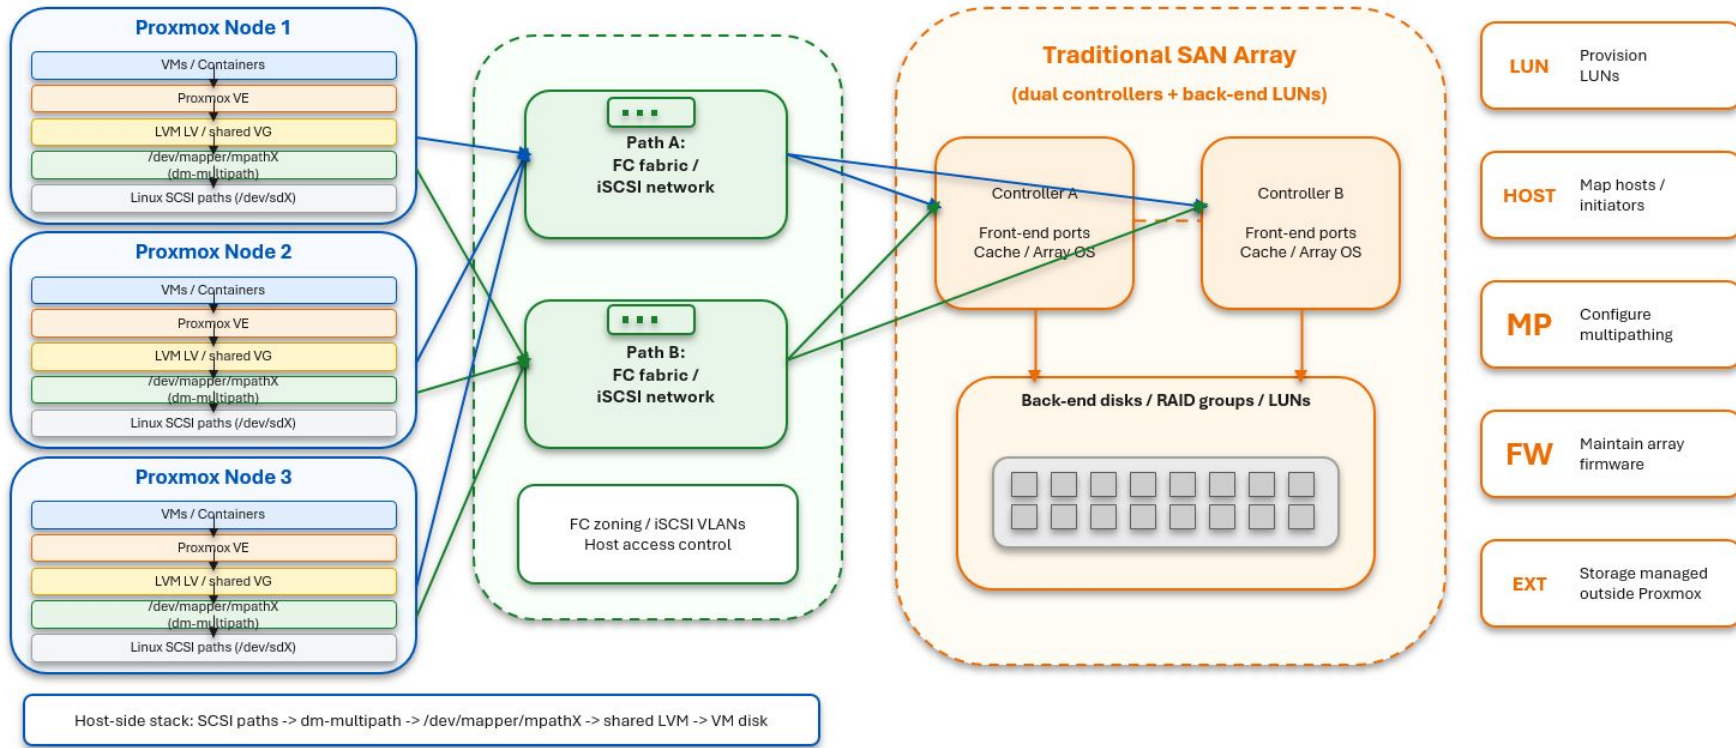


SAN Technology with Proxmox

- iSCSI, Fibre Channel, SAS devices recognized by Linux kernel
- *multipathd* set up at Linux level
- Must use shared thick LVM (reclaim is possible with TRIM/DISCARD)
- Snapshot behavior depends on storage backend and Proxmox version; LUN provisioning remains external
- Storage stack maintained by storage admins

Traditional Shared Storage for Proxmox iSCSI / Fibre Channel SAN

Shared block storage works, but the data path and control plane are outside Proxmox



Result:

LUNs can work well, but Proxmox still depends on LUN provisioning, zoning/VLANs, multipath policy, and external array management.

NFS Technology with Proxmox

- NFS provided by storage device
- Beware of CoW on CoW (ZFS)
- Thin/sparse qcow2
- Redundancy at network level
- NAS head still a Failure Domain



Ceph with Proxmox

Your old SAN was built to share disks between servers. Ceph makes distributed storage part of the Proxmox cluster itself.

Integrated & resilient

Distributed, self-healing storage native to the Proxmox cluster

Proxmox-native automation

Deploy, monitor, and manage Ceph through Proxmox VE GUI, CLI, shell, native tools

Hyperconverged or external

Run Ceph on Proxmox nodes, or connect to an existing external Ceph cluster

Transparent & extensible

Open, visible, scriptable architecture—no opaque array controller

RBD snapshots & clones

Native VM disk workflows without qcow2 snapshot overhead

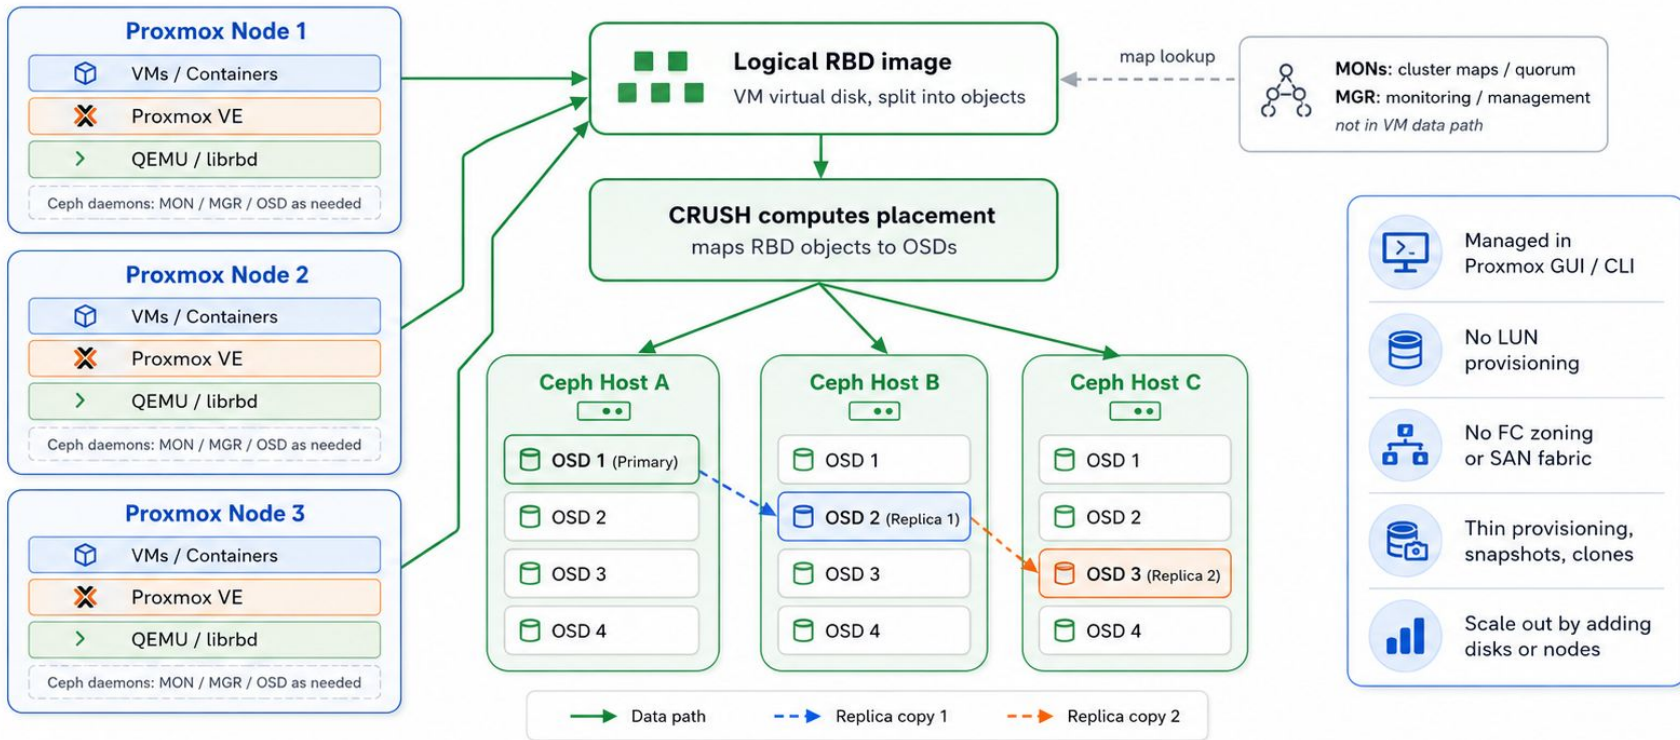
Thin provisioning + scale-out

Allocate on write and grow by adding OSDs or nodes

Ceph still requires sound engineering design, proper hardware, and networking.

Proxmox with Ceph: Integrated, Distributed Storage

VMs use RBD virtual disks; librbd and CRUSH map objects directly to distributed OSDs



Result:

Proxmox accesses a VM RBD virtual disk via librbd; CRUSH maps its objects to OSDs while replication, recovery, and scaling happen across the cluster.

Next Steps

- Evaluate:
 - Requirements
 - Costs
 - Availability
- Plan to Migrate to Ceph
- Reach out for:
 - Use Cases
 - Architecture Help
 - POC



References

https://pve.proxmox.com/wiki/Storage:_RBD

https://pve.proxmox.com/wiki/Storage:_CephFS

<https://iss-integration.com#proxmox>

<https://github.com/proxmox/pve-docs/blob/master/pveceph.adoc>

ask anything: alex@iss-integration.com