metalsoft

MetalSoft

Frequently Asked Questions (FAQ)

In Brief

MetalSoft is a complex piece of software evolved over many years and understanding how it operates can be chalanging. This document atempts to briefly answers frequent questions and point the reader to more extensive documentation.

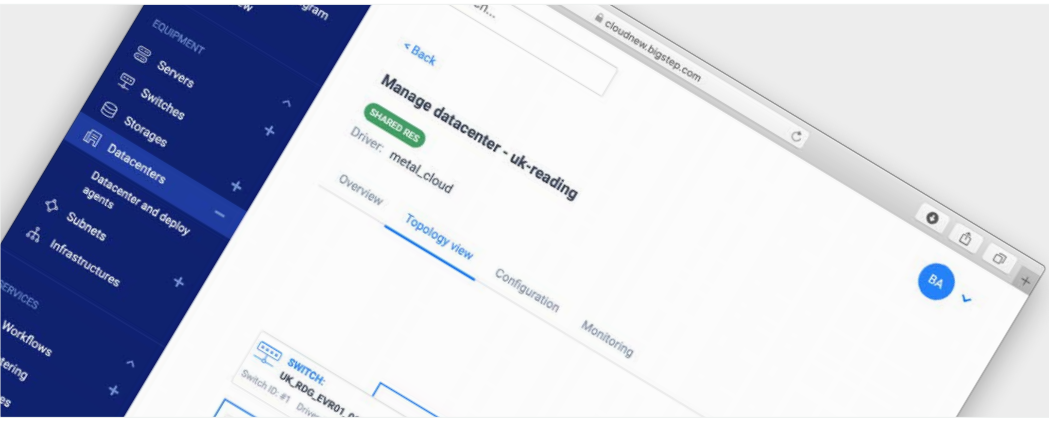

Contents

1. How is iSCSi storage managed?

The controller uses the agents as a proxy to manage iSCSI storage appliances via their APIs. We provide our own storage software but we can also interact with others.

There can be many storage appliances, they act as single pool of disk space out fof which MetalSoft creates “Drives” (iSCSI LUNs which appear in the operating system as block devices). Drives can be created starting from a template - such as CentOS in diskless mode or can simply be unformatted disk space.

Drives can be attached to any instance including the ones booted from a local drive. Drives can also be expanded while in operation up to a maximum size of 2TB. Drives use copy-on-write, compression and deduplication features of the storage pool.

2. Where are of ISO images stored?

MetalSoft requires an HTTP repository to be reachable to agents (normally located in the same datacenter as the controller).
The assets such as ISOs in this repository are accessed during deployments and cached locally in the target datacenter by the TFTP/HTTP agent. In Edge deployments, the agents are able to use other agents as source instead of the main repository in order to save bandwidth.

Admins can upload their own assets to this source repository. The assets can then be “registered” in the application and referenced in install scripts. Users can also store assets in any other repository including the official linux repositories and only register a ‘reference’ to those repositories. This will make the asset appear as if it is locally stored while being dynamically pulled from the external repository and cached in the target datacenter.

3. How is Data center agent redundancy achieved?

The Data Center Agents can be instantiated multiple times within the same datacenter (on two or more different servers). They will all receive broadcast DHCP messages from servers and they all answer with the same reply. Provisioning requests are performed by

all agents in a datacenter in a round-robin fashion. Agents register themselves as executors in the AFC queue. If an on-going task fails another agent can retry it.

4. How is controller redundancy achieved?

The controller's architecture is relatively simple. It comprises of an API server which can be instanciated any number of times across which a pair of loadbalancers will send HTTP traffic. All API servers use the same MySQL database which can be setup in master-slave or cluster mode.

Disaster recovery options are also supported by the agents as they receive from the controller (and can have hard-coded) a list of controllers to try. If the primary ones fails for a longer period of time the agents will failover to the secondary location if one is configured.

5. Does MetalSoft support VMWare ESXi?

MetalSoft has very good support for VMWare products and vSphere in particular:

1. ESXi can be installed using our custom install process
2. Nodes can be automatically enrolled in the vCenter using the Workflow
3. Nodes can access the same drives in order to enable vMotion
4. Operators can use the workflow to enable SR-IOV and other virtualization optimizations in the BIOS as a pre-deployment step.

6. Does MetalSoft support Microsoft Windows?

Windows is very well supported both for local install:
1. MetalSoft maintains the installation process and for diskless boot (for certain certified systems).
2. MetalSoft provides a custom Windows deployment utility that runs on the server after the deployment and configures users, mounts iSCSI drivers etc.
3. We also provide built-in License management

including support for on-demand (hourly) licenses.

7. How does MetalSoft perform data destruction?

Before a server is released back into the pool of servers several actions are performed:
1. Local disks are cleaned either by writing zeros or using the Secure Erase option if available.
2. BIOS and BMC configurations are reset to default
3. Switch ports are de-configured and shut down

8. How does firewall services work?

MetalSoft manages the firewall on the hosts. The default is block all and exceptions to this can be added by users. Rules are InstanceArray wide thus are automatically applied to new servers in the same cluster and modified on all servers in an InstanceArray at the same time.
Firewall management can be disabled if manual intervention directly on the host is required.

9. How is tenant segregation achieved?

MetalSoft provides its own bare metal SDN that performs dynamic network equipment reconfiguration to provision isolated L2 and L3 networks. For L2 isolation we use either VLANs, VPLS (MPLS) or eVPN (VXLAN). L3 networks are also provided within a VPC eqivalent concept called an *infrastructure* where gateways are SVIs (switch virtual interfaces) within isolated L2 networks.

10. How does traffic account work?

Traffic accounting is perfomed by the monitoring agent using a mechanism called SFLOW. Switches are configured to push samples (only the headers of packets) of certaint traffic to the agent. The agent then increments counters for certain subnets. This is done locally in the datacenter being monitored. Periodically or when a report needs to be generated the controller will interrogate the monitoring agent for the data. The data is stored in the agent.

11. How does monitoring work?

There is a special agent called the Monitoring agent. This performs SNMP monitoring of the OS using MetalSoft configured MIBs on the OS. It also monitors certain BMC metrics such as thermal data. The

agent keeps the data in-memory in a local round-robin timeseries database (similar to Cacti's RRD). The data is then queried by the controller when it needs to be presented to the user or stored for longer term.

12. How is the control plane secured?

There are my security mechanisms in place to prevent unauthorized access to the control plane. Agents use TLS for all traffic and keys generated by the controller to decrypt secrets and to authenticate themselves to the controller. Keys are unique per agent and can be "rotated" in case of a compromised agent host. Keys use 256 bit AES encryption.
We also use various mechanisms such as DHCP option 82 to protect against DHCP impersonation and other forms of attacks.

13. Can I bring my own images?

Yes. You can add any assets to the repository or host your images in your own repository and register them in our application as an external asset.

14. Can you support existing hardware or just greenfield deplopyments?

Yes- with some caveates. It is possible to enroll servers and switches manually in order to keep track of them for inventory and monitoring purposes but not for provisioning purposes. However when they are released from the current use by their users they can be put into full production mode. This is of course dependent on compatibility with MetalSoft's hardware compatibility list.

15. Can MetalSoft manage only parts of switches?

Yes- with some caveates. Our software only configures ports that are connected to servers registed in MetalSoft. However the switch must be configured in our database with apropriate exceptions (eg: VLAN ranges, subnets etc) in place in order not to create conflicts which can break new deployments and perhaps affect existing setups.

16. What is the minimum amount of hardware required?

To run the MetalSoft agents a VM or server with a minumum of 2cores 4GB of RAM, 40 GB of disk and two NICs is required. The server or VM must be connected to both the Out-of-band network (to each IPMI interfaces and switch management interfaces) and the IN-band network in order to coordinate server booting or perform installs.
To run a MetalSoft controller the minimum requirement is a server with 8 cores 64GB or RAM and 1TB of disk. MetalSoft provides an option to use the controller hosted by MetalSoft and run only the agent. The agents phone home and require no special network configuration so for POCs this option provides the fastest and simplest deployment option.

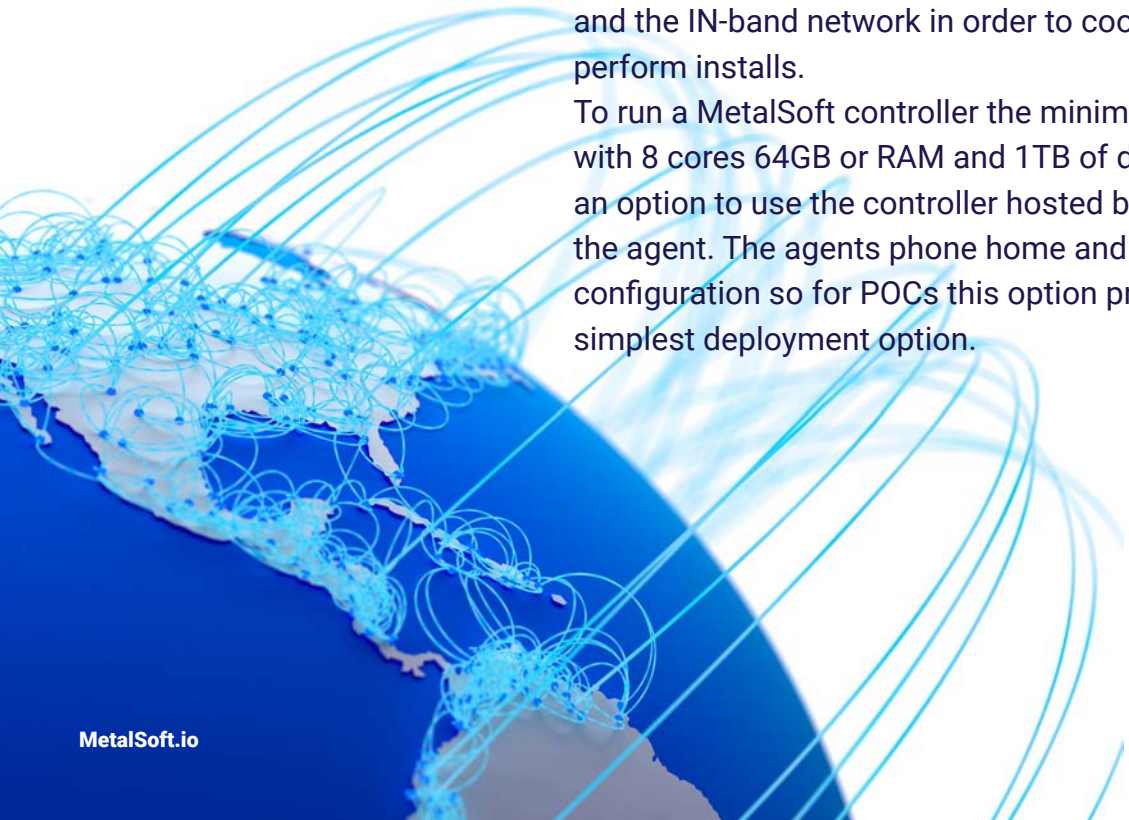

metalsoft

Let’s talk!

To find out more about our solution drop us a line at
inquiries@metalsoft.io