

# Network Management Card 3 local witness for Cisco HyperFlex Edge 2-node Clusters

By Ray Munkelwitz



## ABSTRACT

Increasingly, critical compute resources, such as hyperconverged infrastructure (HCI) are being deployed at the edge. To achieve high availability at these edge sites can be a challenge. It requires a true micro data center solution, that goes beyond just power and cooling but also takes in consideration the unique requirements of these HCI systems.

Data consistency across a hyperconverged infrastructure (HCI) clustered file system is of critical importance and is a challenge with 2-node systems. To ensure consistency, a quorum (or node majority) is required for continual operation. Traditionally, 2-node HCI clusters required a third site or dedicated infrastructure for witness VMs which can drive up deployment costs and complexity. Other solutions, such as Cisco's *Intersight Invisible Cloud Witness* are reliant on network connectivity to Intersight. A local or external witness moves this critical functionality directly to the network management card (NMC) 3 supporting the cluster, providing the highest levels of resiliency in a turnkey, validated solution.



## APPLICATIONS

Remote Branch Offices, Intelligence for the Edge and Distributed Networks.

## CUSTOMER BENEFITS

- **No dependency on the WAN.** This is critical for edge sites which may not have reliable connectivity.
- **Low cost**, especially when compared to adding a third HCI node or another dedicated infrastructure to host witness VMs,
- **Simple to deploy**, requiring only the feature to be enabled on the NMC and a simple configuration on the HyperFlex Edge cluster.
- **Easy to operate and maintain.** The NMC contains all the logic to natively serve as a HyperFlex Edge witness. The NMC is local to the rack enclosure and runs on a purpose built, hardened, and secure OS. Any updates are delivered through the native NMC update mechanisms, eliminating the need to patch and update a separate device on the network.
- **Resiliency.** This feature can be used in place of the Cisco Intersight Invisible Cloud witness and has built in battery backup, powered by the UPS hosting the NMC.

## Introduction

Along with PowerChute™ Network Shutdown, the local network management card witness function provides a highly available and resilient solution for Cisco's HyperFlex HCI clusters. This application note contains:

- The concept of a witness for 2-node HyperFlex HCI clusters and the need for a local witness.
- Basic solution components.
- Enabling the external witness on the NMC3
- Enabling the external witness on the HyperFlex Edge cluster

**Witness concept**

Figure 1 illustrates, at a high level, both a single and multi-site 2-node HyperFlex deployment. A quorum is necessary for data integrity to avoid a “split-brain” scenario. Traditionally this type of deployment would require a third site to host a VM to perform the witness function. To avoid that expense and management overhead, Cisco introduced the *Intersight Invisible Cloud Witness* to seamlessly perform this important function behind the scenes.

Unfortunately, a Cisco *Intersight* based witness can only operate when a network connection is available to *Intersight*. This is where the local witness comes in, acting on the LAN, it can serve this critical function independent from any network backhaul or WAN. This enables any Edge location to run a highly available 2-node cluster powered by the NMC local witness.

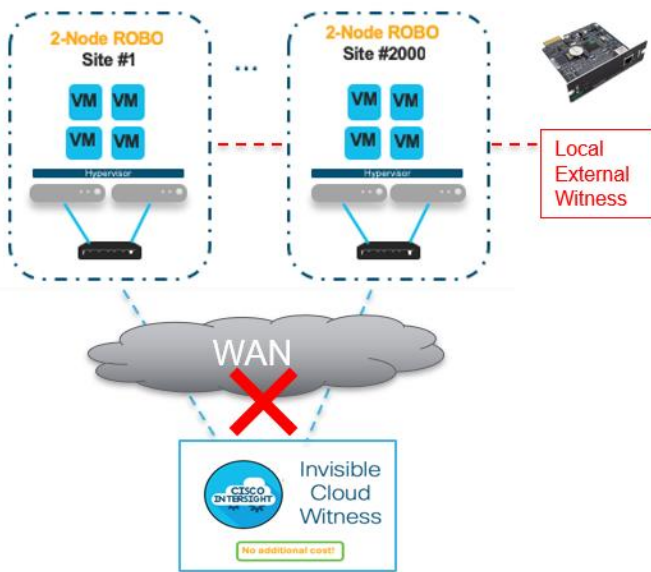


Figure 1: System level view

Note: Connectivity to Intersight is still required for operations like install, upgrade, monitoring, of the Hyperflex nodes. Only the data-plane is decoupled from the Intersight connectivity requirement. The two are designed to work together.

Figure 2: Typical solution components

The NMC communicates with the Hyperflex controller VM via a local network connection.

Cisco access switch (Catalyst 9200 shown)

Cisco Hyperflex Edge 2-Nodes Smartplay Bundle (HXAF-SP-EDGE-M5CE1(2 or 3))

APC Smart-UPS with NMC3 (SRT1500RMXLANC shown)



**Solution components**

Figure 2 illustrates the basic system components. Cisco provides several Hyperflex models and can be used with a range of switches. Similarly, APC has many UPS models compatible with the third-generation network management card (NMC3).

APC Smart-UPS single phase models with either an NMC3 (AP9640, AP9641 and AP9643) in the SmartSlot with firmware version 1.4 or later are supported. Please refer to [UPS-NMC3 Compatibility chart](#) or for UPS's with embedded or pre-populated NMC cards, please refer to: [UPS NMC3 Transition Details](#)

All HyperFlex Edge 2-Node configurations running HyperFlex Data Platform 4.0(2) and later releases are supported.

### Enabling the witness on the NMC3

The witness functionality is enabled using a standard ini file which is the method used today to mass configure the NMC. The first step is to create a new ini file or modify an existing one with the following content:

```
[HyperFlex]
Access=enabled
Username=<username>
Password=<password>
```

The witness feature requires its own username and password which differs (and is independent) from the credentials of the NMC itself. The witness uses a dedicated service account that only allows the specific witness API calls.

This document explains the methods to upload the ini file: <https://www.apc.com/us/en/faqs/FA156117/>

After the ini file is copied to the NMC, its presence can be verified by looking in the NMC UI or CLI event log.

E000: Success

--- Event Log -----

Date	Time	User	Event
------	------	------	-------

09/02/2020 12:51:41 System Configuration file upload complete, with 3 valid values.

09/02/2020	12:51:40	apc	SSH/SCP: File transfer complete.
09/02/2020	12:51:40	apc	SSH/SCP: File transfer started.

You may also test the witness server reachability using the **health** API.

```
curl -k -u <username>:<password> https://<nmc IP>/hyperflex/health
```

```
Arbitrator service is alive and healthy at 2020-11-30
15:12:24.000000000 +0000 UTC
```

If the witness does not respond please ensure that:

- The HX witness service is enabled via the ini file
- Verify you are reaching the NMC over https (if enabled) or http (if enabled)
- The request url is correct with the /hyperflex subdirectory
- The proper witness username/password is used in the request (basic auth)

### Enabling the external witness on the HyperFlex Edge Cluster

On the Cisco HyperFlex side, the witness function can be configured using the HyperFlex device connector UI in HX Connect (see figure 3) or by using the device connector API.

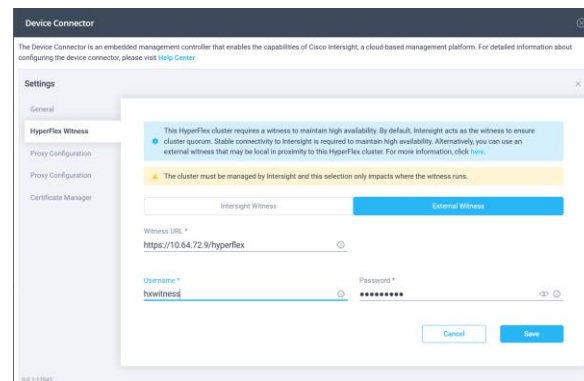


Figure 3: HyperFlex Device Connector UI

The local (or external) witness can be used with or without the Invisible Cloud Witness. All witness configuration details are encrypted and stored locally within the device connector.

Additional information can be found in the *Intersight* online help:

[https://intersight.com/help/resources/configuring\\_device\\_connector](https://intersight.com/help/resources/configuring_device_connector)

You may also learn more about the Cisco HyperFlex Invisible Cloud Witness here:

<https://www.cisco.com/c/dam/en/us/products/collaterals/hyperconverged-infrastructure/hyperflex-hx-series/whitepaper-c11-741999.pdf>

