



ObjectMatrix

Object Matrix **MatrixStore 10GigE Networking**

Version 1.4, Feb 2015

Contents

| | |
|--|----------|
| 1 Overview | 3 |
| 1.1 10GigE Cabling and Connectors | 3 |
| 1.1.1 Fiber-Optic Cables | 3 |
| 1.1.2 Copper Cables | 3 |
| 1.1.3 Connectors | 4 |
| 1.1.4 Summary | 4 |
| 1.2 10GigE Cables and Connectors in MatrixStore | 4 |
| 2 Switches | 6 |
| 2.1 SFP+ Switches | 6 |
| 2.1.1 Stacking | 6 |
| 2.1.2 External (Customer) Network | 7 |
| 2.2 10G-BASE-T Switches | 7 |
| 3 How to connect MatrixStore to customer networks | 8 |
| 3.1 Direct SFP+ | 9 |
| 3.2 Bridged SFP+ | 10 |
| 3.3 Direct RJ45 | 11 |
| 3.4 Bridged RJ45 | 12 |



1 Overview

This document describes some basic terminology around 10GigE and how MatrixStore-10GigE can be connected to 10GigE networks.

1.1 10GigE Cabling and Connectors

10GbE has been in existence since 2002. Due to the history of 10GbE, it has not been completely backward compatible with previous generations of Ethernet technology, which has contributed to some confusion about the cables and connectors required for 10GbE.

There are two basic cable types available for 10GbE applications: **copper and fiber-optic cables**.

1.1.1 Fiber-Optic Cables

There are different designations for fiber-optic cables depending on the bandwidth supported.

- **10GBASE-SR**
Currently, the most common type of fiber-optic 10GbE cable is the 10GBase-SR cable that supports an SFP+ connector with an optical transceiver rated for 10Gb transmission speed. These are also known as “short reach” fiber-optic cables.
- **10GBASE-LR**
These are the “long reach” fiber optic cables that support single-mode fiber optic cables and connectors.

1.1.2 Copper Cables

Common forms of 10GbE copper cables are as follows:

- **10GBASE-CR**
Currently, the most common type of copper 10GbE cable is the 10GBase-CR cable that uses an attached SFP+ connector, also known as a **Direct Attach Copper (DAC)**.
 - *Passive DAC*
 - *Active DAC*
Active is the trend. It includes components that boost the signal, reduce the noise and work with smaller gauge cables, improving signal distance, cable flexibility and airflow
- **10GBASE-T**
Cables are Cat6a (Category 6 augmented), also known as Class EA cables.

1.1.3 Connectors

There is a list connectors that can be combined with 10GigE cables to provide 10GigE speeds. Each combination of cable and connector determines the limits in speed and distance supported. The table below show how they can be combined.

| | Type | Lanes | Max. speed per lane (Gbps) | Max. speed total (Gbps) | Cable type | Usage |
|--|-------|-------|----------------------------|-------------------------|-----------------|--|
| Copper CX4 | CX4 | 4 | 5 | 20 | Copper | 10GbE, SDR and DDR Infiniband |
| Small Form-factor Pluggable | SFP | 1 | 4 | 4 | Copper, Optical | 1GbE, Fibre Channel: 1, 2, 4Gb |
| Small Form-factor Pluggable enhanced | SFP+ | 1 | 16 | 16 | Copper, Optical | 10GbE, 8Gb & 16Gb Fibre Channel, 10Gb FCoE |
| Quad Small Form-factor Pluggable | QSFP | 4 | 5 | 20 | Copper, Optical | Various |
| Quad Small Form-factor Pluggable enhanced | QSFP+ | 4 | 16 | 64 | Copper, Optical | 40GbE, DDR, QDR & FDR Infiniband, 64Gb Fibre Channel |

1.1.4 Summary

The most common types of 10GbE cables currently use SFP+ connectors.

- For short distances (up to 3 meters) such as within a rack or to a nearby rack, use DAC with SFP+ connectors, also known as 10GBASE-CR.
- For mid-range distances (up to 300m), use laser optimized multi-mode fiber (LOMMF) cables, either OM3 or OM4, with SFP+ connectors.
- For long-range distances, use single-mode fiber-optic cables (OS1), also known as 10GBASE-LR.
- If your Ethernet switching infrastructure and network adapters support 10GBASE-T (RJ45) cables and connectors, use Cat6 or Cat6a cables.

Remember, the cable you select is determined by the customer switch, not the adapter, in your environment.

1.2 10GigE Cables and Connectors in MatrixStore

Each MatrixStore-10GigE node contains:

- Two 10GigE ports for the internal network bonded for redundancy using the 802.3ad protocol.
- Two 10GigE ports for the external network bonded for redundancy using the 802.3ad protocol thus presenting one IP address to the client network.

The external network ports in MatrixStore can be **10GigE RJ45 or SFP+** depending on the customer preference.



Nodes with **RJ45 connectors** are compatible with:

- Copper 10GBASE-T networks directly.
- Copper 10GBASE-CR networks (AKA DAC) through a bridge switch.
- Fiber 10GBASE-SR through a bridge switch

Nodes with **SFP+ connectors** are compatible with:

- Copper 10GBASE-CR networks (AKA DAC) directly. **Up to 3 meters.**
- Fiber 10GBASE-SR directly through Intel® Ethernet SFP+ SR Optics E10GSFPSR transceiver. **Up to 300m.**
- Copper 10GBASE-T through a bridge switch.

It is important to note that MatrixStore nodes can be attached to fiber networks, however, they need a specific Intel transceiver attached to the nodes. The customer switch to which the fiber is connected at the other end may require a different model of transceivers, it is switch dependent.

Some customers may not have sufficient free ports to accommodate multiple dual connected MatrixStore nodes therefore extra switches are required to bridge MatrixStore external network to the customer's networks.

2 Switches

The switches currently used in MatrixStore configurations are:

- Allied Telesis AT-8000GS/24. 1GigE switch used for:
 - Internal network in 1GigE MatrixStore clusters
 - IMPI network connectivity
- Allied Telesis AT-GS900/8:
 - 1GigE switch used for IPMI connectivity
- Netgear Prosafe SFP+ 24 ports (XSM7224S)
 - Bridging to external 10GigE SFP+ networks (DAC or fiber) when failover is required.
- Netgear Prosafe 8/12/24 Port 10G-BASE-T Switch (XS708E/XS712T/M7100)
 - Internal network in 10 GigE MatrixStore clusters
 - Bridging to external 10G-BASE-T, DAC or fiber networks when failover is not required.

2.1 SFP+ Switches

The switches Object Matrix supports are Netgear Prosafe SFP+ 24 ports (XSM7224S). Depending on the configuration it can support up to a 19 node MatrixStore cluster. To expand beyond that requires further switches to be stacked or replacement with a 36 or 48 port switch.



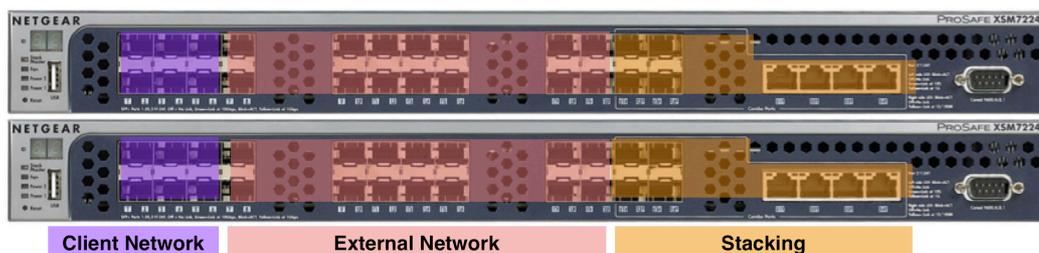
2.1.1 Stacking

The four of the 10GigE copper ports are used for the stacking function. It is possible to use only two ports for stacking thus providing more ports for MatrixStore nodes.



2.1.2 External (Customer) Network

Depending on the number of MatrixStore nodes the remainder of the ports are given over to the client network. A minimum of four is recommended thus providing over 40Gigabit/second of throughput to and from the MatrixStore Cluster. The diagram below shows 6 ports used to connect to the client network.



2.2 10G-BASE-T Switches

The 10G-BASE-T switches below are used depending on the size of the MatrixStore cluster and the amount of bandwidth required to the client network:

- Prosafe 8 Port 10G-BASE-T Switch (XS708E). This has 6 RJ45 ports for the MatrixStore External network and 1 RJ45 to connect to the client network. The remaining port is used to stack the two switches for redundancy
- Prosafe 12 Port 10G-BASE-T Switch (XS712T). This has 10 RJ45 ports for the MatrixStore External network and 2 SFP+ to connect to the client network. The remaining ports are used to stack the two switches for redundancy
- Prosafe 24 Port 10G-BASE-T Switch (M7100). This has 20 RJ45 ports for the MatrixStore External network and 4 SFP+ to connect to the client network. The remaining ports are used to stack the two switches for redundancy.

3 How to connect MatrixStore to customer networks

There are multiple options to connect MatrixStore to a 10GigE customer network. The network ports can be SFP+ or 10GigE RJ45 depending on the customer preference.

This section also contains networking options for customers that do not have sufficient free ports to accommodate multiple dual connected MatrixStore nodes. In this scenario extra switches are required to bridge the two networks.

Each MatrixStore node contains:

- Two 10GigE ports for the internal network bonded for redundancy using the 802.3ad protocol.
- Two 10GigE ports for the external network bonded for redundancy using the 802.3ad protocol thus presenting one IP address to the client network.

In all case, MatrixStore uses 10G-BASE-T internal switches.

3.1 Direct SFP+

All MatrixStore nodes have SFP+ external ports and are directly connected to the customer switch(es). Each node has two ports that connect to the client network switch. Each port should ideally connect to two different switches for redundancy. The client network switch must support the 802.3ad protocol.

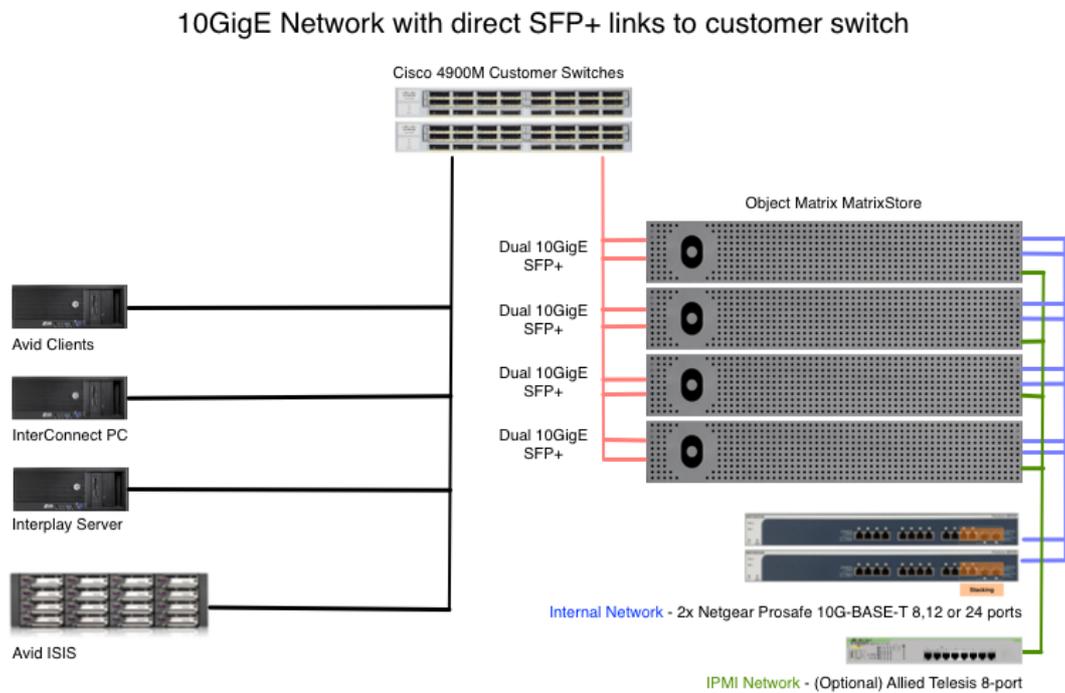
Supported customer networks:

- DAC SFP+
- Fiber: by using the previously described Intel transceiver.

Depending on the amount of MatrixStore nodes, an extra 1GigE switch may be needed to separate internal IPMI ports from the internal 10GigE switches.

NOTE:

In case the customer has Cisco switches with X2 interfaces, they can use the Cisco® OneX Converter Module to connect SFP+ MatrixStore external network.



3.2 Bridged SFP+

MatrixStore external network connects to the customer switch(es) through a set of redundant SFP+ switches bridging both networks.

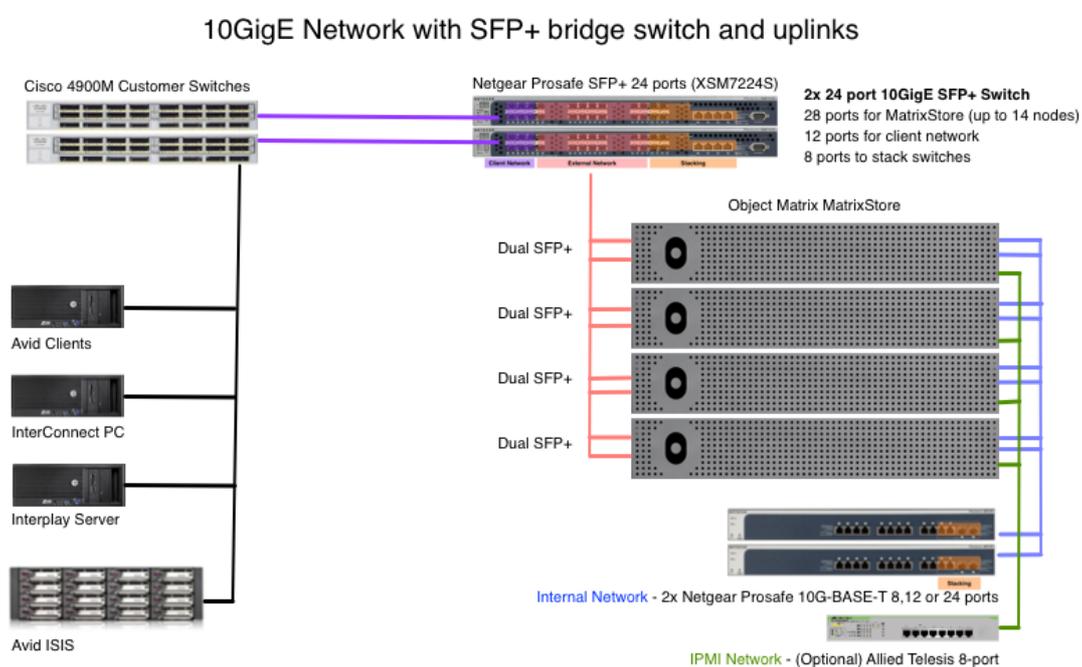
This setup reduces the amount of SFP+ connections required on the customer switch(es) for MatrixStore at the price of reducing the overall bandwidth available for MatrixStore access from client machines.

This option supports external network failover.

Supported customer networks:

- DAC SFP+
- Fiber
- 10G-BASE-T via uplinks

Depending on the amount of MatrixStore nodes, an extra 1GigE switch may be needed to separate internal IPMI ports from the internal 10GigE switches.



3.3 Direct RJ45

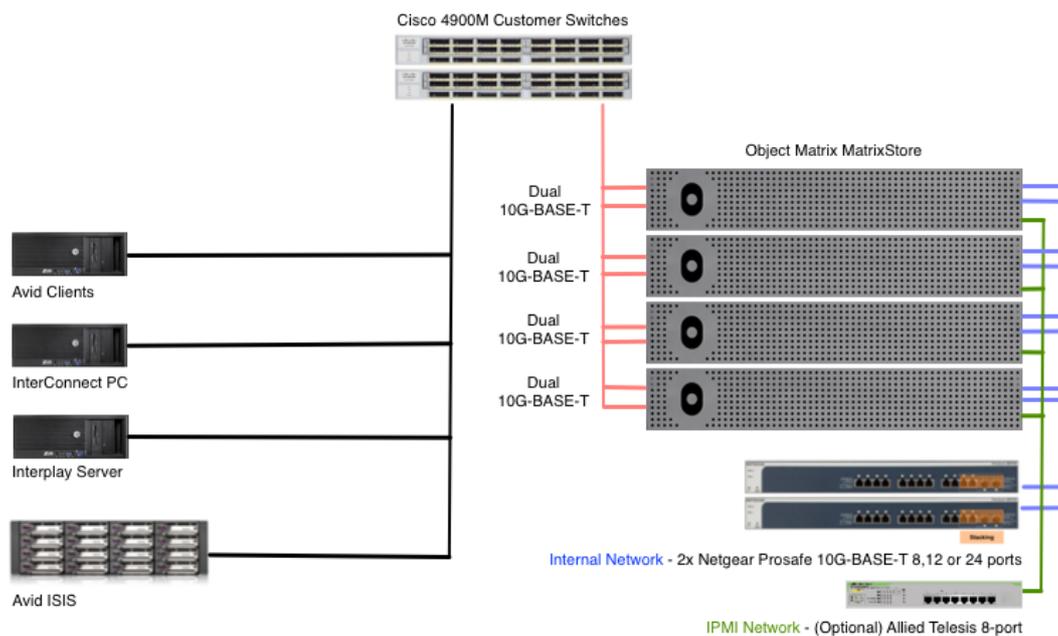
All MatrixStore nodes have RJ45 external ports and are directly connected to the customer switch(es). Each node has two ports that connect to the client network switch. Each port should ideally connect to two different switches for redundancy. The client network switch must support the 802.3ad protocol.

Supported customer networks:

- 10G-BASE-T

Depending on the amount of MatrixStore nodes, an extra 1GigE switch may be needed to separate internal IPMI ports from the internal 10GigE switches.

10GigE Network with direct 10G-BASE-T links to customer switch



3.4 Bridged RJ45

MatrixStore external network connects to the customer switch(es) through one 10G-BASE-T switch bridging both networks.

Supported customer networks:

- 10G-BASE-T
- DAC SFP+ via uplinks
- Fiber via uplinks

Depending on the amount of MatrixStore nodes, an extra 1GigE switch may be needed to separate internal IPMI ports from the internal 10GigE switches.

