

YCE Connection matrix

This article describes the IP connectivity required for the NetYCE servers. This information is essential to determine the access permissions that should be granted between the various environments (access-lists, firewalls).

The specific portnumbers that are used between the users, the servers and the network depends to a large degree on the NetYCE architecture deployed. Since the architectures can largely be divided into three basic configurations, single server, high-availability, and multi-tier, the required connectivity is presented using these setups.

Each arrow indicates the where the call setup originates, the portnumbers used in these connections are listed beside the arrow.

When the portnumber is listed between brackets, these are either optional, or only used for testing or maintenance. The purpose and usage of each portnumber is given in the legend at the bottom of this article.

Single server

The single-server setup is for small, non-critical deployments and for development / test / acceptance environments.



High-availability servers

In the High-availability setup, two servers providing all functions are working in tandem providing load-balancing and hot-switchover functionality. The databases are configured for master-master replication which is continually monitored to allow for immediate switchover to the secondary. (Fallback to primary will not take place until the databases are back in sync.)

This setup is recommended for all larger deployments that are deemed critical and where the NetYCE servers are located in relatively well protected environments (isolated from public or insecure networks).



Multi-tier servers

In the multi-tier setup, the databases are separated from the user and network facing components allowing for either better database performance and security zoning. The front-end servers are configured to connect to a primary and secondary database as is desired by the design.

Up to seven Front-end servers are supported where some can be made dedicated to a specific network or share the load of networking change jobs. Front-end servers can be configured to provide only access to the web-GUI, but are not connected to the targeted network.

Note: This functionality requiring the use of the 'distributed scheduler' is currently under development and is targeted for the 7.1 release

Up to two (master) databases can be deployed using the master-master replication, and additional slaves (read-only) are optional.



Connectivity matrix

Unless specified otherwise the protocol used is TCP.

Workstation			
	Browser to NetYCE front-end		
		80 / 443	http / https
		8080	Json TLS
		8888	test tool NetYCE API
	Terminal to NetYCE front-end		
		22	ssh / scp / sftp application mgmt
	Terminal to Network devices		
		22	ssh
		23	telnet
NetYCE front-end			
	to NetYCE database server		
		3306	odbc database API
		8888	NetYCE xml API
		8080	NetYCE Json API
		7777	NetYCE scheduler
		22	ssh / sftp
	to NetYCE front-end servers		
		8888	NetYCE xml API
		8080	NetYCE Json API
		7777	NetYCE scheduler
	to Ldap or AD servers		
		389	Ldap protocol
	to Network devices		
		22	ssh / netconf
		23	telnet
		443	Json API
		other	vendor-specific API
Network devices			
	to NetYCE front-end		
		69 (UDP)	tftp
		20 / 21	ftp
		22	sftp

NMS / OSS			
	to NetYCE db/front-end servers		
		8888	NetYCE xml API
		8080	NetYCE Json API
		80 / 443	URL based services
	from NetYCE db/front-end servers		
		8888	NetYCE xml API
		80 / 443	URL based services
		25	outgoing email
NetYCE database servers			
	to NetYCE database server		
		3306	database replication

From:
<https://wiki.netyce.com/> - **Technical documentation**

Permanent link:
https://wiki.netyce.com/doku.php?id=guides:reference:connection_matrix

Last update: **2024/07/03 12:31**

