

Introducing Tachyon

Summary

An overview of Tachyon is presented; what it does and how it can benefit your organization. An introduction to the features of Tachyon is provided, alongside a high-level explanation of the Tachyon architecture.

In this section...

Tachyon Features

An overview of all the Tachyon features and enhancements.

Tachyon Architecture

A description of Tachyon Stacks and their Tachyon components, the Tachyon Agents and how they connect to provide the Tachyon features.

On this page:

- [Introduction](#)
- [Implement Tachyon](#)
- [Add Instruction Definitions to Tachyon Server](#)
- [Assign permissions for Tachyon administrators and users](#)
- [Asking questions](#)
- [Coverage and filters](#)
- [Performing actions](#)
- [Real-time Tachyon](#)
- [Basic Instruction flow](#)
- [Content delivery](#)

Introduction

Tachyon is a systems management tool that turns the traditional systems management model on its head. It is based on the concept that a device is aware of its own current state more accurately than any centrally held repository. Also if you can query the device directly then the need to store masses of data becomes unnecessary.

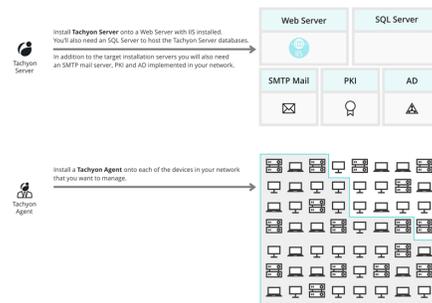
Implement Tachyon

To implement Tachyon you need to design, install and verify your installation. Full details can be found in [Implementing Tachyon](#).

Broadly speaking, installing Tachyon is done in two steps:

1. Install Tachyon Server onto an appropriate Windows server with IIS installed.
2. Install a Tachyon Agent onto each of the devices you want to manage

Install Tachyon



Add Instruction Definitions to Tachyon Server

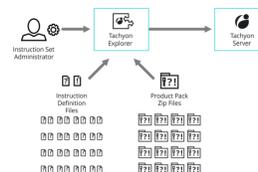
Using the Tachyon administration console, import Instruction Definitions into Tachyon. An Instruction Definition Administrator can load Instructions by dragging and dropping one or more Instruction Definition files onto the interface. The Instruction Definitions may also be grouped together in a zip file, known as a Product Pack, and added in a single operation. Tachyon provides a whole range of 1E and community maintained Product Packs, available from the Tachyon forums.

Full details of adding Instruction Definitions into Tachyon can be found in [Instruction sets page](#).

Add Instruction Definitions to Tachyon Server

Before you can use Tachyon you will need to add some Instruction Definitions. Tachyon instructions are either questions or actions and define what you can do in Tachyon. The Instruction Definitions may be grouped together into a zip file, known as a Product Pack. Tachyon provides a whole range of 1E and community maintained Product Packs, available from the Tachyon forums.

Tachyon Instruction Set Administrators are able to drag and drop one or more Instruction Definition files or Product Packs onto the Administration Instruction Sets page in the Tachyon Administration console to add them to Tachyon Server.



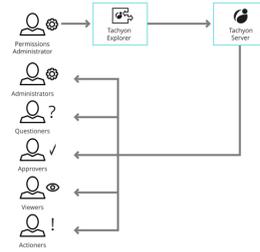
Assign permissions for Tachyon administrators and users

Following on from installation you then assign permissions to particular users corresponding to the Tachyon administration roles and product pack questioner, actioner, approver and viewer roles.

Full details of assigning permissions for Tachyon administrators and users can be found in the [Users page](#) and the [Roles page](#).

Assign permissions for Tachyon administrators and users

Using the Administration Permissions page in the Tachyon Explorer, Tachyon Permissions administrators can determine Tachyon users and define their roles.

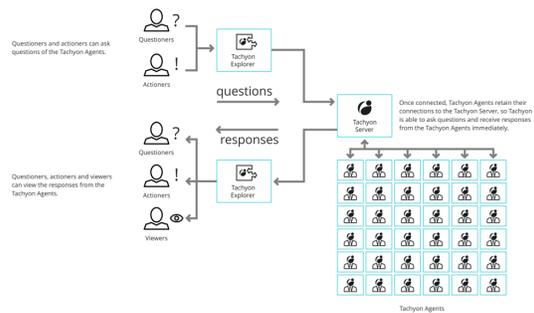


Asking questions

Using the Tachyon Explorer, questions can be asked of the Tachyon Agent devices by users with the questioner or actioner roles. The responses can be viewed by viewers, questioners and actioners.

Details on asking questions can be found in [Explorer Home page](#), [questions and responses](#).

Asking questions

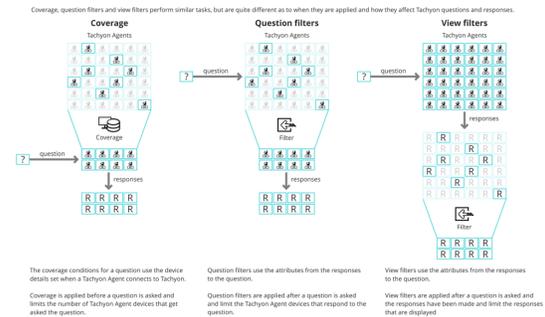


Coverage and filters

You can use coverage and filters to reduce the amount of network traffic caused by asking a question. They are also useful in reducing the amount of data that is presented to the Tachyon users. Tachyon provides coverage, question filters and view filters.

- Coverage is applied before a question is asked and reduces how many Tachyon Agent devices get asked the question. Coverage uses the device properties that are sent by a Tachyon Agent when it connects to Tachyon.
- Question filters are applied at the Tachyon Agent after the question has been asked and use the attributes from the question responses to determine whether the Tachyon Agent should send the response.
- View filters are applied in the Tachyon Explorer after the responses have been sent and reduce the responses that are displayed.

Coverage and filters



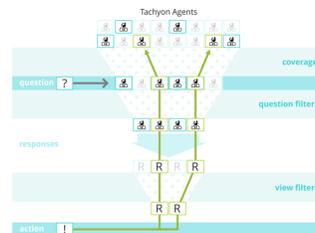
For more information see [Coverage](#), [question filters](#) and [view filters](#).

Performing actions

Before an action is performed a question must be asked. This identifies the set of devices the action will be performed on, after the coverage, question and view filters have been applied.

Performing a follow-up action

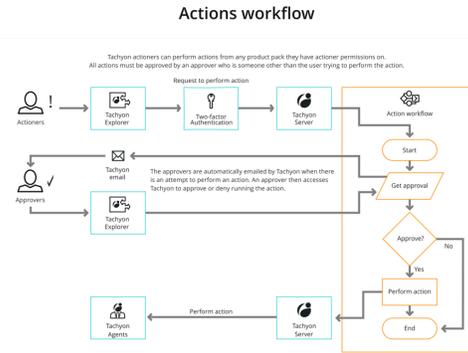
Actions are performed as the follow-up to questions. They are targeted at the devices that pass the coverage and question filters and whose responses pass the view filters.



Using the Tachyon Explorer, actions can be performed on the Tachyon Agent devices by users with the actioner role on a product pack subject to the following:

- they will need to enter their password
- if the action has been configured to use two-factor authentication, they will need to enter the authentication code
- they will need the approval of someone other than themselves with an approver role.

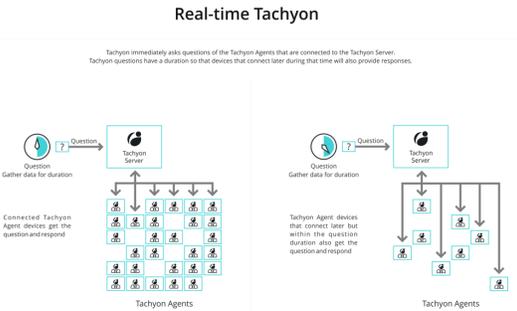
Details on performing actions and the approval workflow can be found in [The action approval workflow](#).



Real-time Tachyon

Tachyon works in real-time, so not all of the devices may be connected at the time a question is asked. Tachyon questions have a configurable duration that they will gather data for - allowing devices that connect later, within the **gather data for** duration, to respond.

Questions also have a configurable **keep answers for** duration - given that the answers from a live network can potentially get stale quickly and Tachyon questions are fast enough that you can simply ask the question again rather than relying on old answers.



Basic Instruction flow

The picture opposite shows the basic flow of Tachyon Instructions between the Tachyon components. For simplicity the workflow section, where the Instruction is authenticated and verified, has been omitted.

The basic Instruction flow is:

1. A Tachyon Consumer initiates an Instruction via a request to the Consumer API
2. The Consumer API stores the Instruction in the Tachyon Master database
3. The Instruction is authenticated and verified by the Workflow and becomes ready for processing by the Core
4. The Core relays the Instruction to the Switch
5. The Switch transmits the Instruction to the appropriate connected Tachyon Agents
6. The Tachyon Agents process the Instruction on their devices and produce responses in real-time
7. The Tachyon Agents pass the responses to the Switch, intelligently minimizing bandwidth utilization
8. The Switch relays the responses to the Core
9. The Core stores the responses in the Tachyon Responses database while intelligently minimizing the amount of data and maximizing speed of retrieval
10. The Consumer API fetches the responses from the Tachyon Responses database and makes them available to the Consumer

The Tachyon Architecture has been designed to the sending of Instructions and gathering of responses for thousands of devices in seconds.

Content delivery

To maximize the speed of Instruction delivery and gathering responses from the Tachyon Agents, Tachyon keeps the size of Instructions to a minimum fixed size. If you want to run an Instruction that requires a script or other sizable resource in order to carry out its functionality, this is made available via an alternative mechanism using the Tachyon Background Channel.

The first step in the content delivery process actually occurs before the Instruction is used, when the Instruction containing the resource is first loaded into Tachyon:

1. A Tachyon Instruction Set Administrator loads an Instruction definition with an associated script, or other sizable resource, into Tachyon using Tachyon Explorer
2. The Consumer API copies the script or resource to the Background Channel
3. The Consumer API stores the Instruction and resource in the Tachyon Master database

After the Instruction has been loaded into Tachyon:

1. A suitably permissioned Tachyon Questioner or Actioner requests the Instruction
2. The Instruction request flows through the Tachyon components as normal
3. This time, when the Instruction gets to the Tachyon Agents, they know that they need to download the Instruction's script or resource in order to process the Instruction
4. The script or resource is downloaded from the Background Channel using HTTPS
5. If 1E Nomad is installed on the Tachyon Agent device it will be used for the download, providing all its download optimization features to speed up the process
6. Once the script or resource file has been downloaded to the Tachyon Agent device it processes the Instruction and returns the responses made as a result of using the script or resource
7. The responses flow back through the Tachyon components as normal