



# 128T Release Notes Version 4.2.6

**Abstract**

128T Release Notes

**8 April 2020**

# Contents

128T 4.2.6 Release Notes .....	3
Special Considerations .....	3
Issues Fixed .....	4
Caveats .....	5

## 128T 4.2.6 Release Notes

*Note:* The 4.2.6 release is a superset of the 4.2.5 release. Features and corrections in the 4.2.5 release are not provided in these release notes. Please refer to the 4.2.5 Release notes for further information

### Special Considerations

- Validation has been added to prevent issues when destination service address and port ranges overlap with interface (waypoint) address and wayport ranges. If your configuration has an overlapping port range prior to the upgrade, you will now receive a validation error. The configuration must be remediated in order to make any additional changes. (I95-33556).
- Support has been added to provision management interfaces. Typically, these are Linux interfaces that are not part of the 128T configuration. These interfaces are called non-forwarding interfaces. It is strongly recommended that non-forwarding interfaces now be configured within the 128T. When the Conductors are upgraded to 4.2.X, the configuration validation will provide a warning when non-forwarding interfaces are not configured for 128T routers configured in a HA pair. If non-forwarding interfaces are configured prior to upgrading the Conductors to 4.2.X, the shared or fabric device interface type must be explicitly set. DHCP however is not supported. Note: 4.1.5 does not raise a validation error if the interface had DHCP. (I95-30831)
- When creating non-forwarding interfaces on pre 4.2.0 routers, these interfaces should not be configured with default routes until after upgrading the routers to 4.2.0 or greater. (I95-30940)
- Support has been added to allow the 128T Software to automatically determine the number of CPU forwarding cores that the router will use. In previous releases the "automatic/manual" parameter was not available and the core count was defined in environment config or within the UI. If set within the environment config (local.init), 4.2.X will set core count to automatic and determine the number of cores to use. If set in the UI, this parameter will be set to manual and the existing UI value will be used. Note that if automatic is set and the core count is changed in the PCLI, the change will be accepted as valid configuration but the automatic parameter will override as the automatic parameter takes precedence over the value in the forwarding core count. (I95-30884)
- Prior to 4.1.5, peering with routers that have the same IP addresses was not supported. In 4.1.5 or greater remote peers having the same IP address is now supported in one direction, with the use of "outbound-only". 128T currently does not support ALL routers having the same IP address, the same IP address support is unidirectional. For example:

The following is supported:

```
R1 (172.16.1.1) --- peered --- R2 (192.168.1.1)
|----- peered --- R3 (192.168.1.1)
```

The following is not supported:

```
R1 (192.168.1.1) --- peered --- R2 (192.168.1.1)
|----- peered --- R3 (192.168.1.1)
```

- On 128T nodes installed after July 1, 2019 on pre 4.2.X versions, Chrome running on MAC OSX Catalina does not consider the 128T self-signed web certificates valid. Apple has put out an advisory: <https://support.apple.com/en-us/HT210176> . Once the node is upgraded to 4.2.X, create a new self-signed web cert in the PCLI `create certificate self-signed webserver (I95-31672)`

## Issues Fixed

- **I95-34068** SVR sessions fail to establish due to waypoint allocation failures after HA node failover.

*Symptom:* The following warning log is generated:

```
Mar 03 09:25:10.813 [HWMC| - ] WARN (icmpManager ) Base Exception:
failed to allocate ports for WayPoint; intf=5.0; local=172.27.233.47;
remote=10.61.55.109
```

Until the system is upgraded to 4.2.6, this issue can be mitigated by removing the corresponding adjacency configuration and adding it back.

- **I95-34164** Load balancer occasionally returns standby paths during packet duplication flow setup
- **I95-34371** Large FIBs can take excess of 30 seconds to load, causing the highway process to fault and restart  
*Symptom:* System would appear to take a long time to load its FIB  
*Conditions:* System running at TRACE log level may induce this problem Until the system is upgraded to 4.2.6, this issue can be mitigated by ensuring the system is configured at log level INFO, and if necessary, reduce the size of the service configuration.
- **I95-34645** Swagger API for "clone" and "move" operations are incorrect. They are `/config/{configStore}/authority/district/{district}/clone` when they should be `/config/{configStore}/authority/district/clone`
- **I95-34577** Interface never becomes active when `shared-physical-address` is configured to be the same as the physical MAC

## Caveats

- **I95-29592** Conductor UI and/or PCLI may not update the asset software version correctly

**Symptom:** The Conductor UI and/or the PCLI may not correctly reflect the software version running on the asset

**Conditions:** After the asset has been upgraded

**Corrective Action:** If the asset is not updated after ~5 minutes after an upgrade is performed, the salt-minion will need a restart on the asset node that does not update the version. This is done with the following command on the node as the root Linux user:

```
"systemctl restart salt-minion"
```

- **I95-27808** `sync peer addresses router force` from conductor may not trigger router to send address information from peer

**Symptom:** When performing the following command on the Conductor PCLI, `sync peer addresses router force` the router may not provide the peer address information

**Conditions:** Unknown

**Corrective Action:** Perform the PCLI command on the router to update the information on the conductor.