


1 Virtual Devices

The virtual devices implemented as part of the VIO infrastructure will be represented in the guest's machine description. The devices will be represented as nodes in the MD along with their properties. This section provides an overview of the virtual device nodes, along with the device hierarchy and their properties.

5



1.1 MD information for virtual devices

All virtual devices will be represented as a node in the guest MD along with its sub-nodes as children of the virtual-devices node. All virtual devices nodes are of the name *virtual-device*. The name and compatible property will identify the name of the specific device and the driver associated with the device. There are two types of virtual device nodes and can be grouped into two separate classes. The first class of device nodes are ones that do not use logical domain channels (LDC) like console, and the existing platform service nodes. These will continue to appear as children of the *virtual-devices* node in the MD. All virtual-device nodes that use LDCs will belong to a class called channel devices and will be grouped under a new node called *channel-devices*. The channel-devices node will be a child of the the virtual-devices node. Some of the virtual-device nodes under the channel-devices node will have one or more child port nodes of type *virtual-device-port*. Each virtual-device-port node can point to one or more channel-endpoint nodes corresponding to the channels within that port.

10

15

1.1.1 Virtual devices node

Name virtual-devices

20

Category optionally required by root

Required subordinates -

Optional subordinates channel-devices, virtual-device

25

Description

This construction node leads directly to all the virtual devices supported within this virtual machine. The number of instances for each device can be derived by counting the number of nodes for each device.

| Name | Tag | Required | Description |
|-------------|-----------|----------|--|
| name | PROP_STR | yes | A string name for this node. This value is currently defined as "virtual-devices". |
| device-type | PROP_STR | yes | A string type for this node. This value is currently defined as "virtual-devices". |
| compatible | PROP_DATA | yes | An array of string names for this node. This value is currently defined as "SUNW,sun4v-virtual-devices". |
| cfg-handle | PROP_VAL | yes | A 64-bit unsigned integer identifying this device uniquely. |

30

1.1.2 Channel devices node

Name channel-devices

Category optionally required by virtual-devices

Required subordinates -

Optional subordinates virtual-device

35

Description

This construction node leads directly to all the channel based virtual devices supported within this virtual machine. The number of instances for each device can be derived by counting the number of nodes for each device.

| Name | Tag | Required | Description |
|-------------|-----------|----------|--|
| name | PROP_STR | yes | A string name for this node. This value is currently defined as "channel-devices". |
| device-type | PROP_STR | yes | A string type for this node. This value is currently defined as "channel-devices". |
| compatible | PROP_DATA | yes | An array of string names for this node. This value is currently defined as "SUNW,sun4v-channel-devices". |
| cfg-handle | PROP_VAL | yes | A 64-bit unsigned integer identifying this device uniquely. |

40 **1.1.3 Virtual device node**

| | |
|-----------------------|---|
| Name | virtual-device |
| Category | optionally required by virtual-devices, channel-devices |
| Required subordinates | - |
| Optional subordinates | virtual-device-port |

45 Description

This node uniquely represents an instance of a virtual device. The properties listed here are applicable to all virtual devices. Each of the virtual devices may specify additional properties that are device class specific.

Common properties

| Name | Tag | Required | Description |
|-------------|-----------|----------|---|
| name | PROP_STR | yes | A string name for this node. (see virtual-device class table). |
| device-type | PROP_STR | yes | A string type for this node. (see virtual-device class table). |
| compatible | PROP_DATA | yes | An array of string names for this node. (see virtual-device class table). |
| cfg-handle | PROP_VAL | yes | A 64-bit unsigned integer identifying this device uniquely. |

50

Device class specific properties

| Name | Tag | Required | Description |
|----------------------|-----------|----------|---|
| vsw-phys-dev | PROP_DATA | no | An array of string names identifying the physical network devices available locally for use by a virtual switch device. |
| vsw-switch-mode | PROP_DATA | no | An array of string names identifying the order of the preferred switching mode(s) for this switch device. Current valid values are <i>switched</i> , <i>promiscuous</i> , and <i>routed</i> . |
| local-mac-address | PROP_VAL | no | A 64-bit unsigned integer in which the lower 48-bits holds the mac address assigned to a virtual network or switch device. The upper 16-bits must be zero. |
| priority-ether-types | PROP_DATA | no | An array of 64-bit integers where the lower 16-bits of each element holds a high priority ethernet type. The upper 48-bits of each element must be zero. The ethernet type corresponds to the <i>Type</i> field in the Ethernet frame as defined by the Ethernet v2 / DIX standard. The virtual network and switch device should prioritize frames with these types over other frames, and ensure that these frames are not dropped, under congestion. |

virtual-device class table (non-channel devices)

| Service Group | Class | compatible | device-type | name |
|---------------|--------|---------------------|-------------|---------|
| Console | Client | SUNW, sun4v-console | serial | console |

virtual-device class table (channel devices)

| Service Group | Class | compatible | device-type | name |
|---------------|--------|----------------------------------|-------------|-----------------------------------|
| Network | Client | SUNW, sun4v-network | network | network |
| Network | Server | SUNW, sun4v-network-switch | vsw | virtual-network-switch |
| Block | Client | SUNW, sun4v-disk | block | disk |
| Block | Server | SUNW, sun4v-disk-server | vds | virtual-disk-server |
| Console | Server | SUNW, sun4v-console-concentrator | vcc | virtual-console-concentrator |
| Serial | Server | SUNW, sun4v-channel | serial | virtual-channel |
| Serial | Client | SUNW, sun4v-channel | serial | virtual-channel-client |
| Serial | Server | SUNW, sun4v-data-plane-channel | serial | virtual-data-plane-channel |
| Serial | Client | SUNW, sun4v-data-plane-channel | serial | virtual-data-plane-channel-client |

1.1.4 Virtual device port node

| | | |
|----|-----------------------|--|
| 55 | Name | virtual-device-port |
| | Category | optionally required by virtual-device node |
| | Required subordinates | - |
| | Optional subordinates | channel-endpoint |

Description

60 This node uniquely represents an instance of a virtual port device. All virtual-device channels connected to the same client are grouped under a single port device. Every virtual-device will have zero or more virtual-device-port nodes.

Common properties

| Name | Tag | Required | Description |
|------|----------|----------|---|
| name | PROP_STR | yes | A string name for the device. (See virtual-device-port class table) |
| id | PROP_VAL | yes | A 64-bit unsigned integer identifying this port uniquely within the virtual-device. |

Device class specific properties

| Name | Tag | Required | Description |
|-----------------------|-----------|----------|--|
| vds-block-device | PROP_STR | no | A string name identifying the block device used by a port in a SUNW,sun4v-disk-server device . |
| vds-block-device-opts | PROP_DATA | no | An array of string names identifying the options for the device used by a vds-port in SUNW,sun4v-disk-server device. Current valid options are: “ro” - The device is used and exported by vds as a read-only device “slice” - The device is exported by vds as a disk slice. “exclusive” - The device is opened for exclusive use by this vds instance only. The device cannot be used by another client or vds instance on the guest. “shared” - The device is exported by the virtual disk server instance to one or more clients connected to it. |
| vdc-timeout | PROP_VAL | no | A 64-bit integer identifying a block device's connection timeout. The value specified in seconds determines the period after which a SUNW,sun4v-disk device will timeout submitting requests if it cannot establish a connection with the virtual disk server. If the property is either not specified or set to 0, the block device will wait indefinitely to establish a connection with the virtual disk server. |
| vcc-tcp-port | PROP_VAL | no | A 64-bit unsigned integer identifying the TCP port assigned to a console group. Provided to vnts daemon via the vcc driver. |
| vcc-group-name | PROP_STR | no | A string name identifying the console group for a domain. Provided to the vnts daemon via the vcc driver. |
| vcc-domain-name | PROP_STR | no | A string name identifying the a domain's console uniquely. Provided to the vnts daemon via the vcc driver. |
| remote-mac-address | PROP_DATA | no | Array of 64-bit unsigned integers where the lower 48-bits of each element holds the mac address assigned to the virtual network or switch device. The upper 16-bits of each element must be zero. This array is a list of mac addresses that are known to be accessible via this port. This is not a complete and comprehensive list. |

| Name | Tag | Required | Description |
|---------------|----------|----------|---|
| switch-port | PROP_VAL | no | Identifies this port as being associated with a SUNW, network-switch device. Property value must be zero. Other values are reserved. <i>Programming note: When using a distributed switch model, this property assists a simple guest in finding a switch port rather than querying every port directly.</i> |
| vldc-svc-name | PROP_STR | no | A string name identifying the service a SUNW, sun4v-channel device is providing over this port. |
| vdpc-svc-name | PROP_STR | no | A string name specifying the service a SUNW, sun4v-data-plane-channel device is providing over this port |

65

virtual-device-port class table

| Service Group | Class | name | name of parent virtual-device node |
|---------------|--------|-----------|------------------------------------|
| Network | Client | vnet-port | network |
| Network | Server | vsw-port | virtual-network-switch |
| Block | Client | vdc-port | disk |
| Block | Server | vds-port | virtual-disk-server |
| Console | Client | vcc-port | virtual-console-concentrator |
| Serial | Server | vldc-port | virtual-channel |
| Serial | Client | vldc-port | virtual-channel-client |
| Serial | Server | vdpc-port | virtual-data-plane-channel |
| Serial | Client | vdpc-port | virtual-data-plane-channel-client |

1.1.5 Channel endpoints node

Name channel-endpoints
 Category optionally required by root node
 Required subordinates -
 Optional subordinates channel-endpoint
 Description

This node uniquely represents a collection of channel endpoint nodes being used by this guest. There should be only one channel-endpoints node. The single channel-endpoints node will have zero or more channel-endpoint nodes as subordinates.

70

75

1.1.6 Channel endpoint node

Name channel-endpoint
 Category optionally required by channel-endpoints node
 optionally required by virtual-device-port nodes
 Required subordinates -
 Optional subordinates -
 Description

80

This node uniquely represents an instance of a channel endpoint available to this guest. Every virtual-device-port node will have zero or more channel-endpoint nodes.

Properties

| Name | Tag | Required | Description |
|--------|----------|----------|--|
| id | PROP_VAL | yes | A 64-bit unsigned integer identifying this endpoint uniquely within the guest. |
| tx-ino | PROP_VAL | yes | A 64-bit unsigned integer identifying the interrupt number assigned to the transmit interrupt for this endpoint. |
| rx-ino | PROP_VAL | yes | A 64-bit unsigned integer identifying the interrupt number assigned to the receive interrupt for this endpoint. |

85