Scalable Fabric Playbook for IXPs : Automation Options for Interconnecting Internet, Content and Cloud

Mikael Holmberg
Distinguished Systems Engineer

sensu
moog
kubernetes
amazon web services
heroku
Jenkins
CROWBAR
puppet labs
C
SALTSTACK
openstack
CLOUD FOUNDRY
git
New Relic
splunk>
docker
CoreOS
mesosphere
JIRA
PAGERDUTY
VictorOps
vmware vSphere
Microsoft

Manual operations
Custom scripts
….

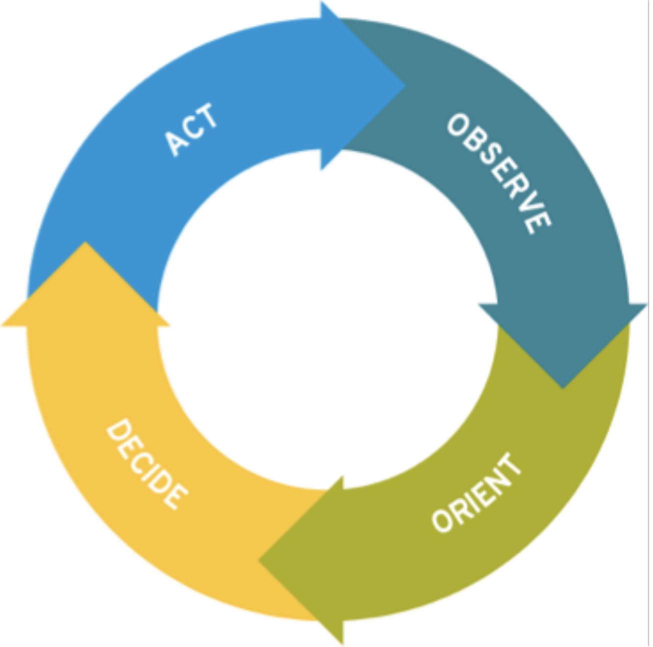

Event Driven Automation

StackStorm

OpenSource Event Automation Platform

Event Driven Automation

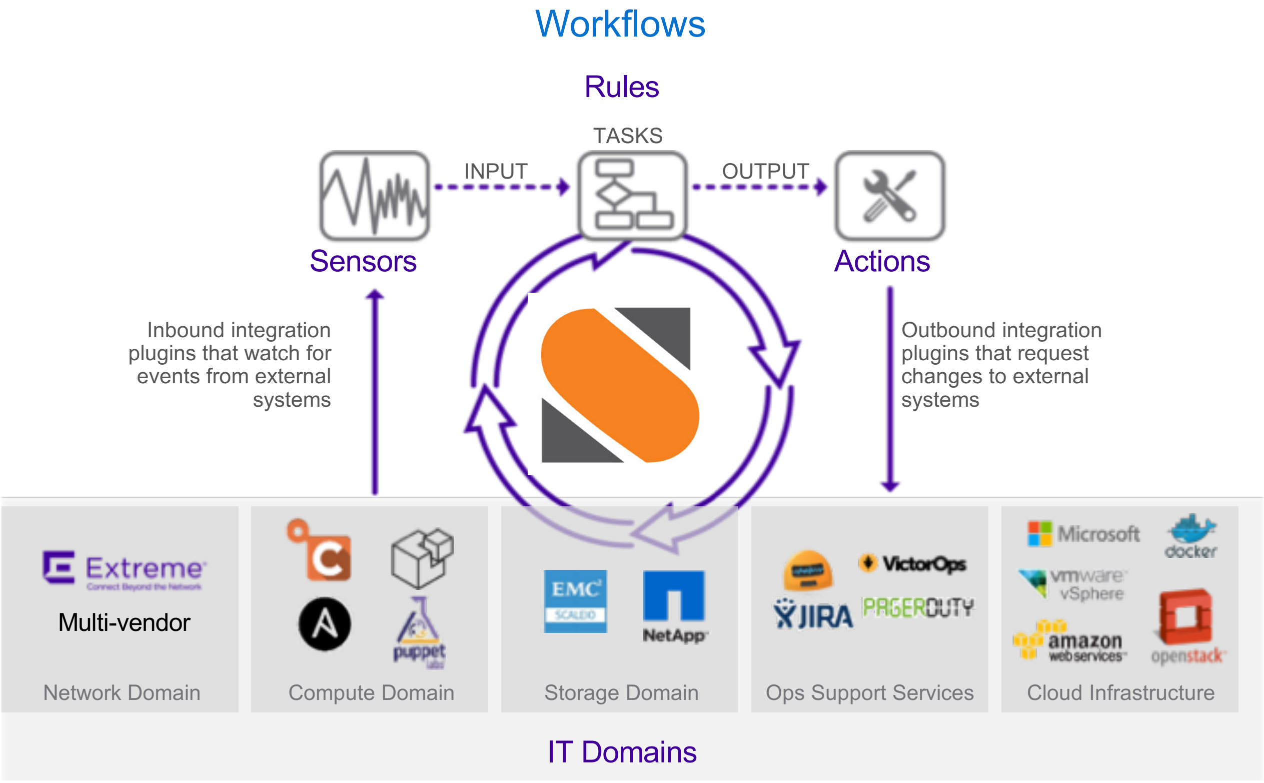

Sensors : Listen for Events like outage in an area

Actions: How to make the change via tools or Stacktorm

Workflow anatomy

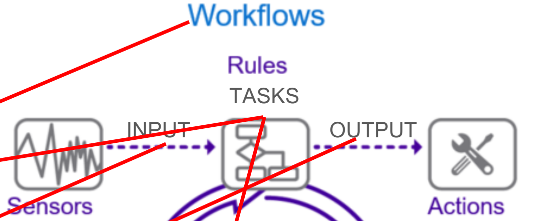

Workflow : Cassandra.replace_host

Tasks

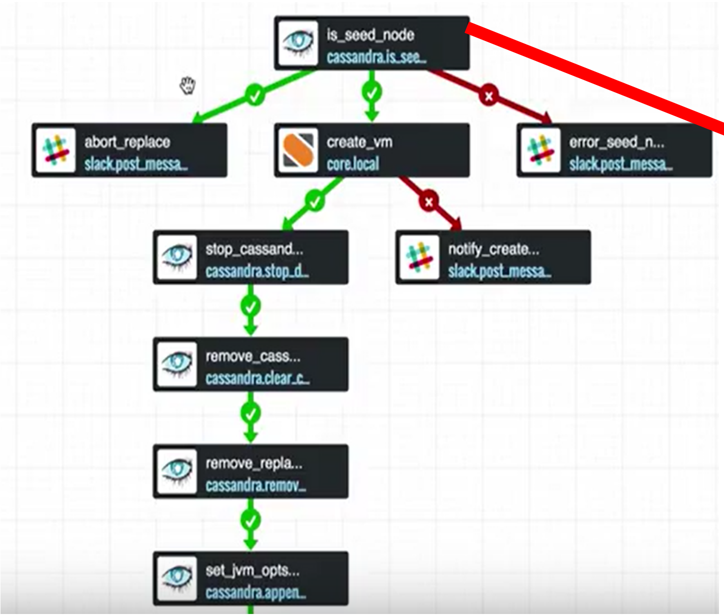

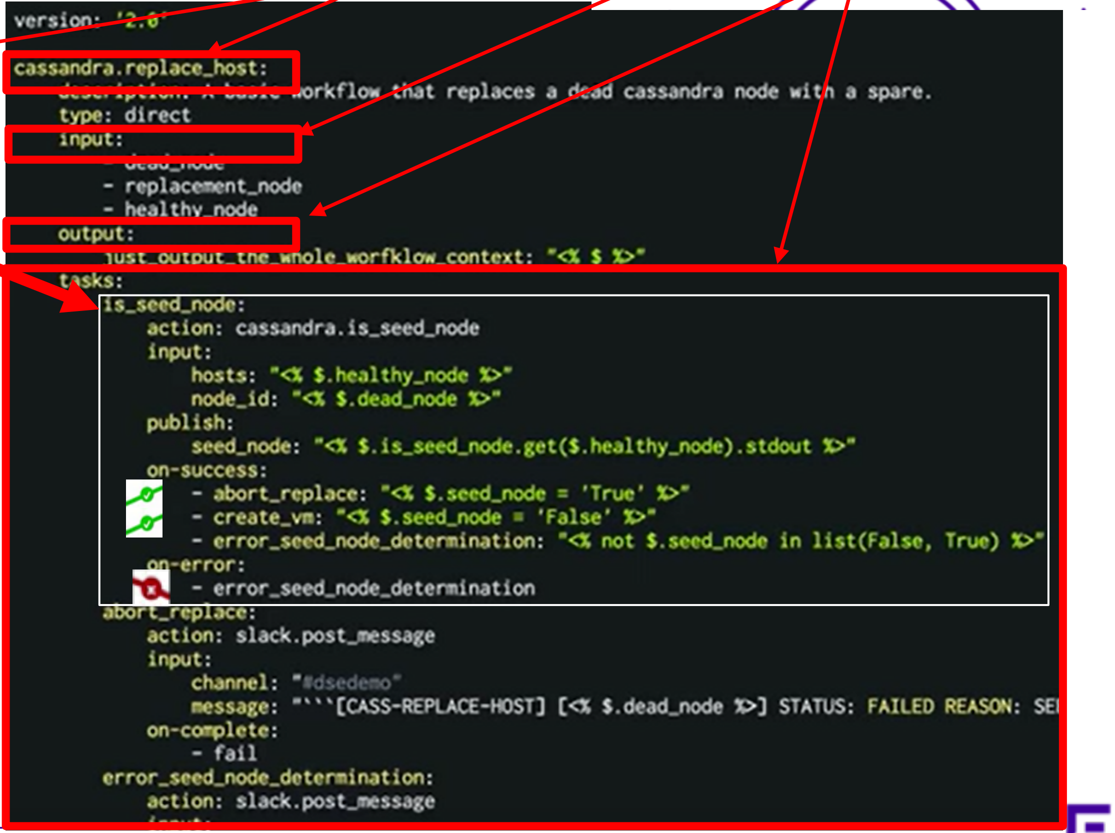

```
version: '2.0'

cassandra.replace_host:
    description: A basic workflow that replaces a dead cassandra node with a spare.
    type: direct
    input:
        - dead_node
        - replacement_node
        - healthy_node
    output:
        just_output_the_whole_worfklow_context: "<% $ %>"
    tasks:
        is_seed_node:
            action: cassandra.is_seed_node
            input:
                hosts: "<% $.healthy_node %>"
                node_id: "<% $.dead_node %>"
            publish:
                seed_node: "<% $.is_seed_node.get($.healthy_node).stdout %>"
            on-success:
                - abort_replace: "<% $.seed_node = 'True' %>"
                - create_vm: "<% $.seed_node = 'False' %>"
                - error_seed_node_determination: "<% not $.seed_node in list(False, True) %>"
            on-error:
                - error_seed_node_determination
        abort_replace:
            action: slack.post_message
            input:
                channel: "#dsedemo"
                message: "```[CASS-REPLACE-HOST] [<% $.dead_node %>] STATUS: FAILED REASON: SE
            on-complete:
                - fail
        error_seed_node_determination:
            action: slack.post_message
            input:
```

Workflow representation and code side by side

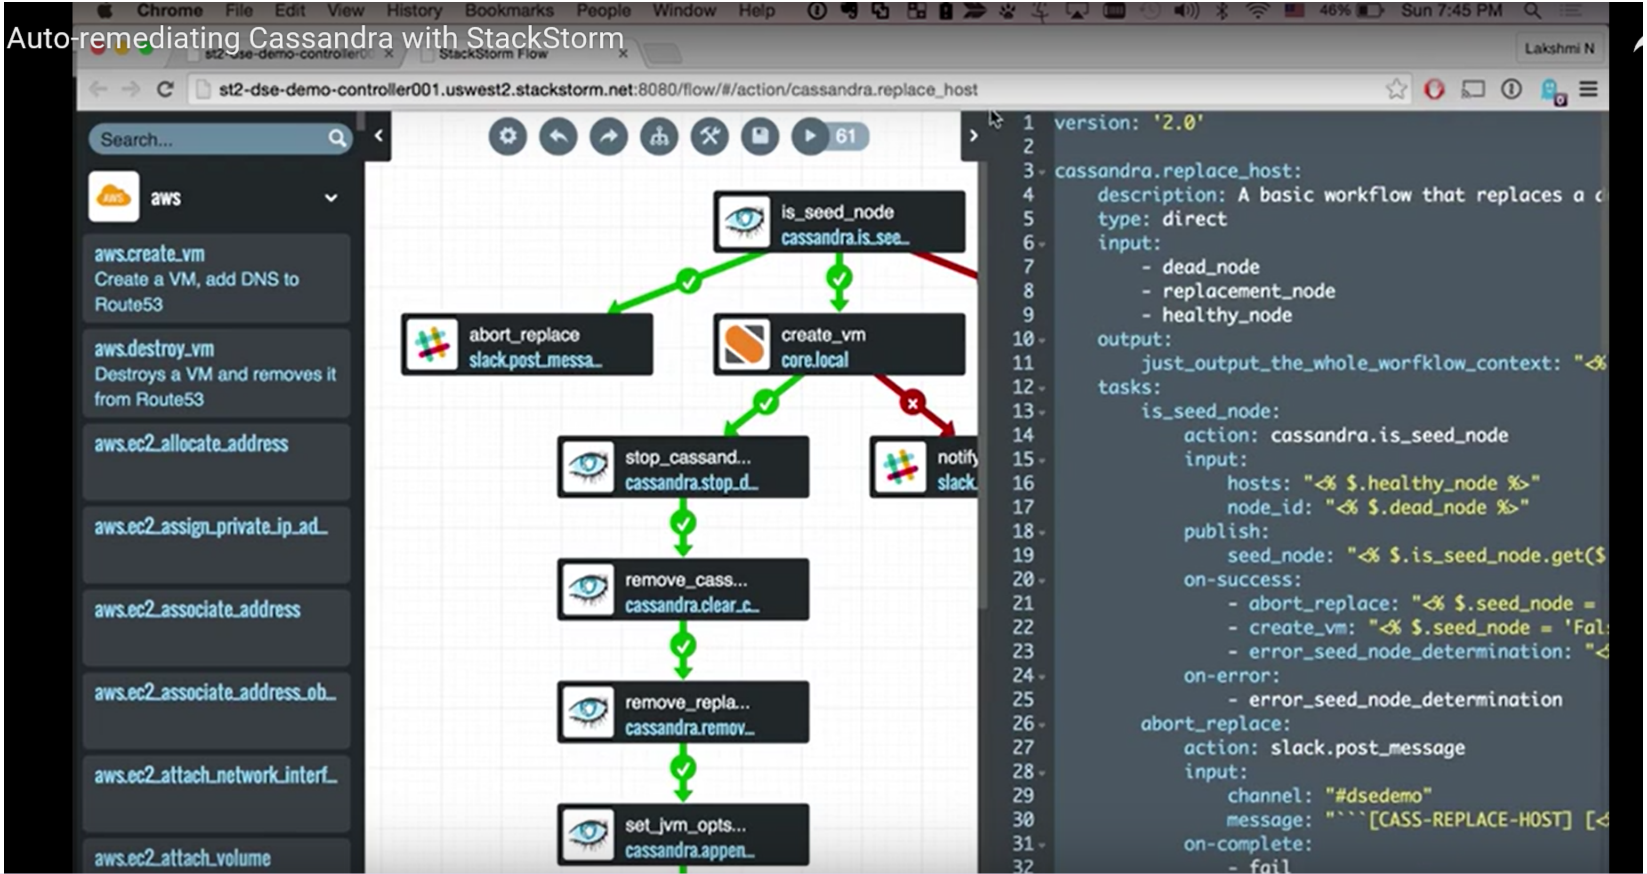
Auto-remediating Cassandra with StackStorm
Chrome File Edit View History Bookmarks People Window Help
Sun 7:45 PM
st2-dse-demo-controller0
StackStorm Flow
Lakshmi N
st2-dse-demo-controller001.uswest2.stackstorm.net:8080/flow/#/action/cassandra.replace_host
Search...
aws
aws.create_vm
Create a VM, add DNS to Route53
aws.destroy_vm
Destroys a VM and removes it from Route53
aws.ec2_allocate_address
aws.ec2_assign_private_ip_ad...
aws.ec2_associate_address
aws.ec2_associate_address_ob...
aws.ec2_attach_network_interf...
aws.ec2_attach_volume
61
is_seed_node
cassandra.is_see...
abort_replace
slack.post_messa...
create_vm
core.local
stop_cassand...
cassandra.stop_d...
notify
slack
remove_cass...
cassandra.clear_c...
remove_repla...
cassandra.remov...
set_jvm_opts...
cassandra.appen...

```
version: '2.0'

cassandra.replace_host:
    description: A basic workflow that replaces a d
    type: direct
    input:
        - dead_node
        - replacement_node
        - healthy_node
    output:
        just_output_the_whole_worfklow_context: "<%
    tasks:
        is_seed_node:
            action: cassandra.is_seed_node
            input:
                hosts: "<% $.healthy_node %>"
                node_id: "<% $.dead_node %>"
            publish:
                seed_node: "<% $.is_seed_node.get($
            on-success:
                - abort_replace: "<% $.seed_node =
                - create_vm: "<% $.seed_node = 'Fal
                - error_seed_node_determination: "<
            on-error:
                - error_seed_node_determination
        abort_replace:
            action: slack.post_message
            input:
                channel: "#dsedemo"
                message: "```[CASS-REPLACE-HOST] [<
            on-complete:
                - fail
```

Tasks in Workflow: Linear , semi parallel or parallel

Linear

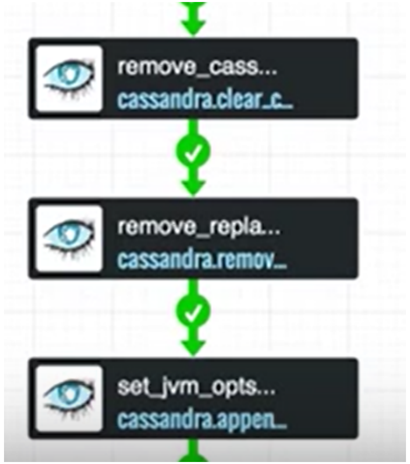

Parallel

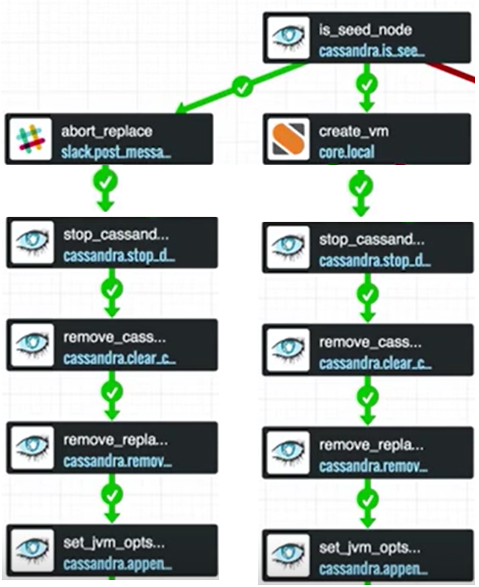

Semi-parallel

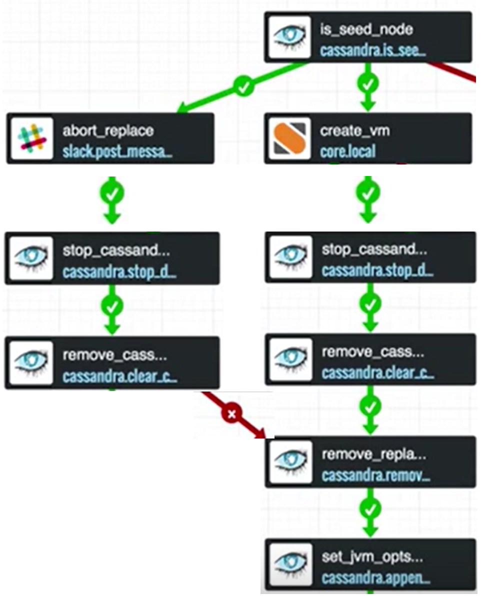

Tasks can collapse with “Joins”

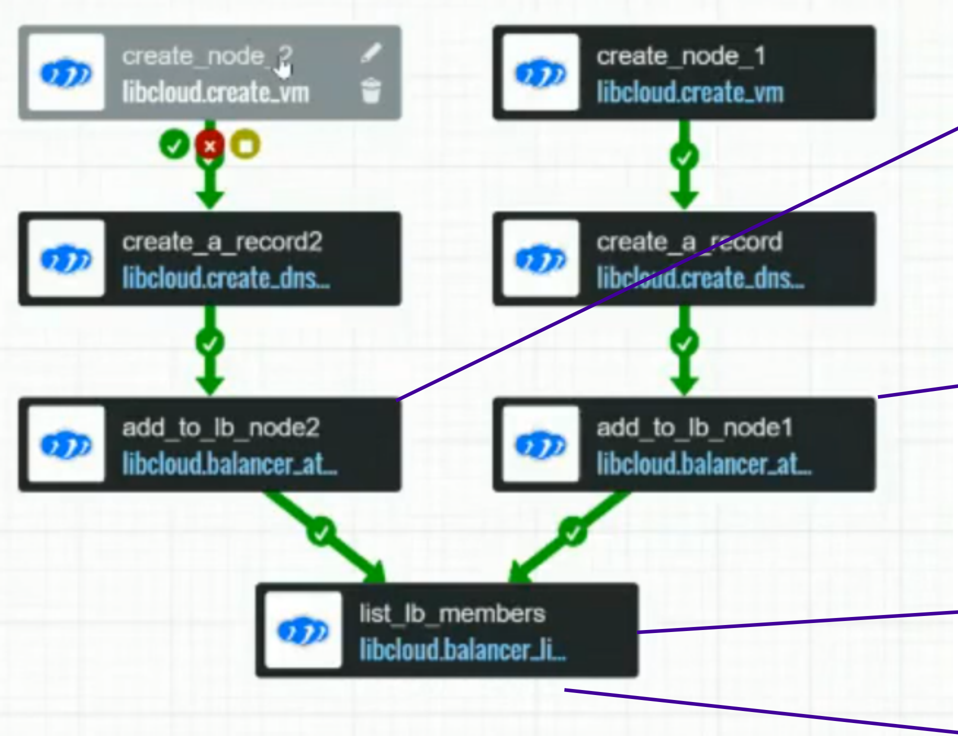

```
tasks:
  .....
  add_to_lb_node2
        action: ….
        input: ….
        publish:…
        on-success:
              - list_lb_members
              - notify

  add_to_lb_node1
        action: ….
        input: ….
        publish:…
        on-success:
              - list_lb_members
              - notify

  list_lb_members
        join: all
        input: …
        publish:…
        on-success:
              -close_request
              - notify
```

Retry policies : for example reboot vm and wait for node to reboot

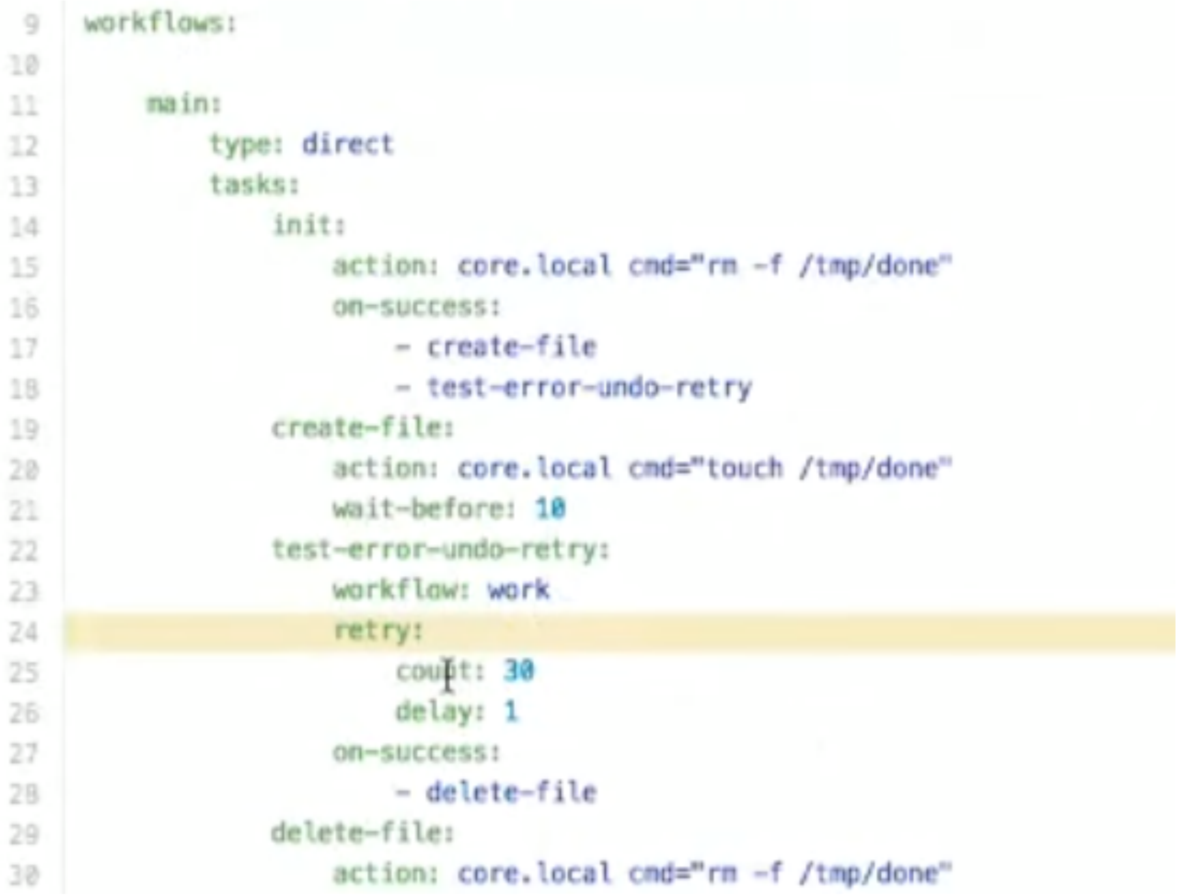

```
workflows:

    main:
        type: direct
        tasks:
            init:
                action: core.local cmd="rm -f /tmp/done"
                on-success:
                    - create-file
                    - test-error-undo-retry
            create-file:
                action: core.local cmd="touch /tmp/done"
                wait-before: 10
            test-error-undo-retry:
                workflow: work
                retry:
                    count: 30
                    delay: 1
                on-success:
                    - delete-file
            delete-file:
                action: core.local cmd="rm -f /tmp/done"
```

If you like writing visual code…..

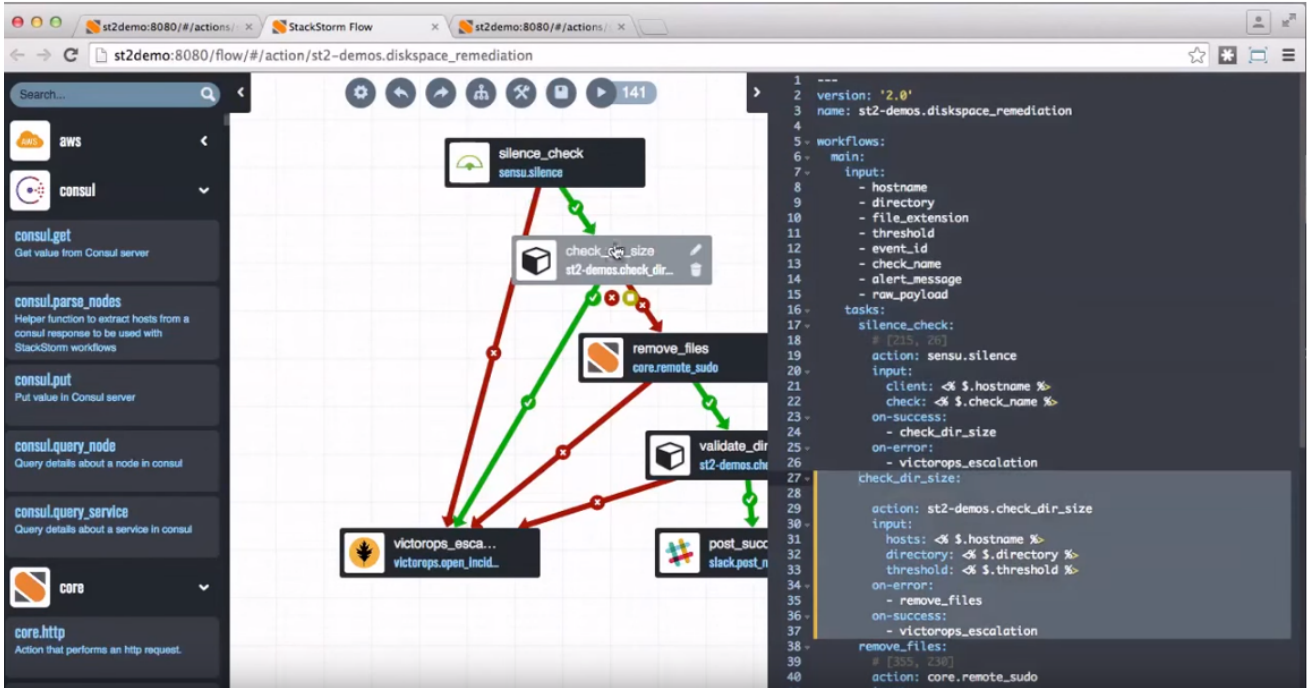

```
---
version: '2.0'
name: st2-demos.diskspace_remediation

workflows:
  main:
    input:
      - hostname
      - directory
      - file_extension
      - threshold
      - event_id
      - check_name
      - alert_message
      - raw_payload
    tasks:
      silence_check:
        # [215, 26]
        action: sensu.silence
        input:
          client: <% $.hostname %>
          check: <% $.check_name %>
        on-success:
          - check_dir_size
        on-error:
          - victorops_escalation
      check_dir_size:

        action: st2-demos.check_dir_size
        input:
          hosts: <% $.hostname %>
          directory: <% $.directory %>
          threshold: <% $.threshold %>
        on-error:
          - remove_files
        on-success:
          - victorops_escalation
      remove_files:
        # [355, 230]
        action: core.remote_sudo
```

Select the workflow

Click on the task you want to edit

Write only that portion of the highlighted code

Stackstorm Integration Packs ….

Integration packs:
https://exchange.stackstorm.org/

Cloud Providers

- aws
- azure
- dimensiondata
- libcloud
- rackspace

Automations and Monitoring

- EXOS
- Network Essentials
- servicenow
- datadog
- sensu
- newrelic
- mmonit
- icinga2
- dripstat

Essentials

- ansible
- napalm
- slack
- chef
- splunk
- cloudflare
- email
- elasticsearch
- docker
- excel

Curiosities

- astral
- cubesensors
- hue
- nest
- powerpoint
- urbandict
- save_kittens
- tesla

Working with Integration Packs

Managing Packs

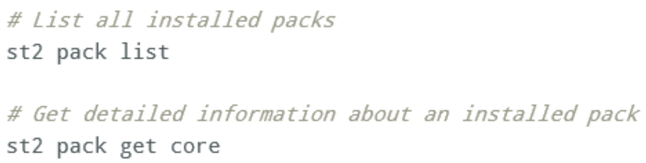

```
# List all installed packs
st2 pack list

# Get detailed information about an installed pack
st2 pack get core
```

Discovering Packs

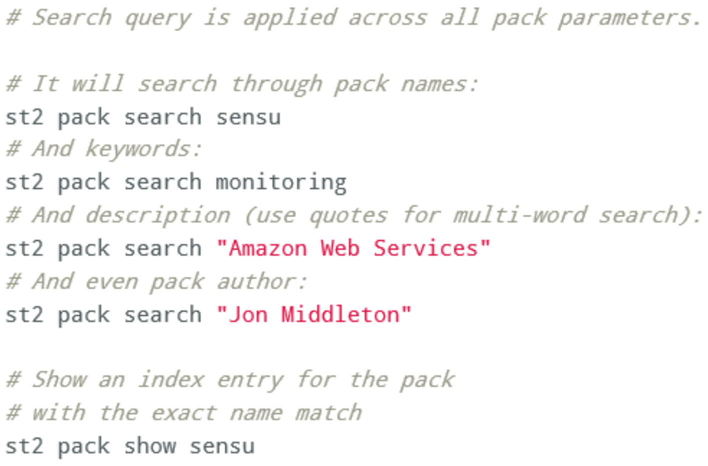

```
# Search query is applied across all pack parameters.

# It will search through pack names:
st2 pack search sensu
# And keywords:
st2 pack search monitoring
# And description (use quotes for multi-word search):
st2 pack search "Amazon Web Services"
# And even pack author:
st2 pack search "Jon Middleton"

# Show an index entry for the pack
# with the exact name match
st2 pack show sensu
```

Installing a Pack

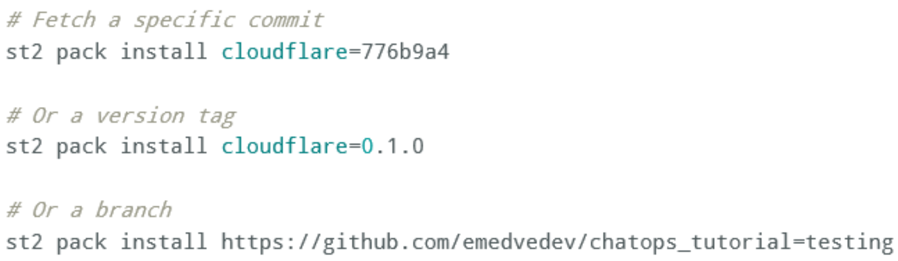

```
# Fetch a specific commit
st2 pack install cloudflare=776b9a4

# Or a version tag
st2 pack install cloudflare=0.1.0

# Or a branch
st2 pack install https://github.com/emedvedev/chatops_tutorial=testing
```

Configuring a Pack

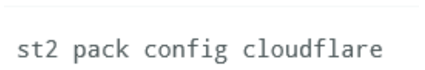

```
st2 pack config cloudflare
```

Example #1 ChatOps Pack : Notify others on job status

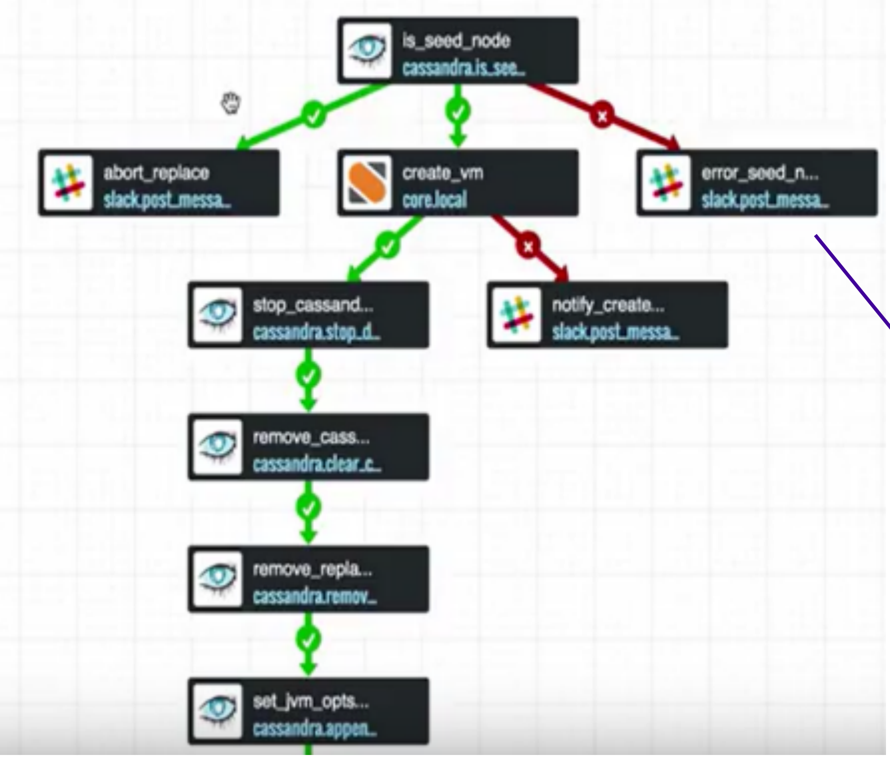

```
cassandra.replace.host:
  type:direct
  input:…
  output:…
  tasks:
         is_seed_node:
                  action: ….
                  input: ….
                  publish:…
                  on-success:
                           - error_seed_notify: …..
                           - create_vm: …..
                  on-error:
                           - error_seed_node_determination:…
```

```
error_seed_notify:
         action: slack.post.message
         input:
               channel: "#IXPOperations"
               message: "…Error: Cassandra replace host...
```

Example #2 Excel Pack : Load Information in workflow

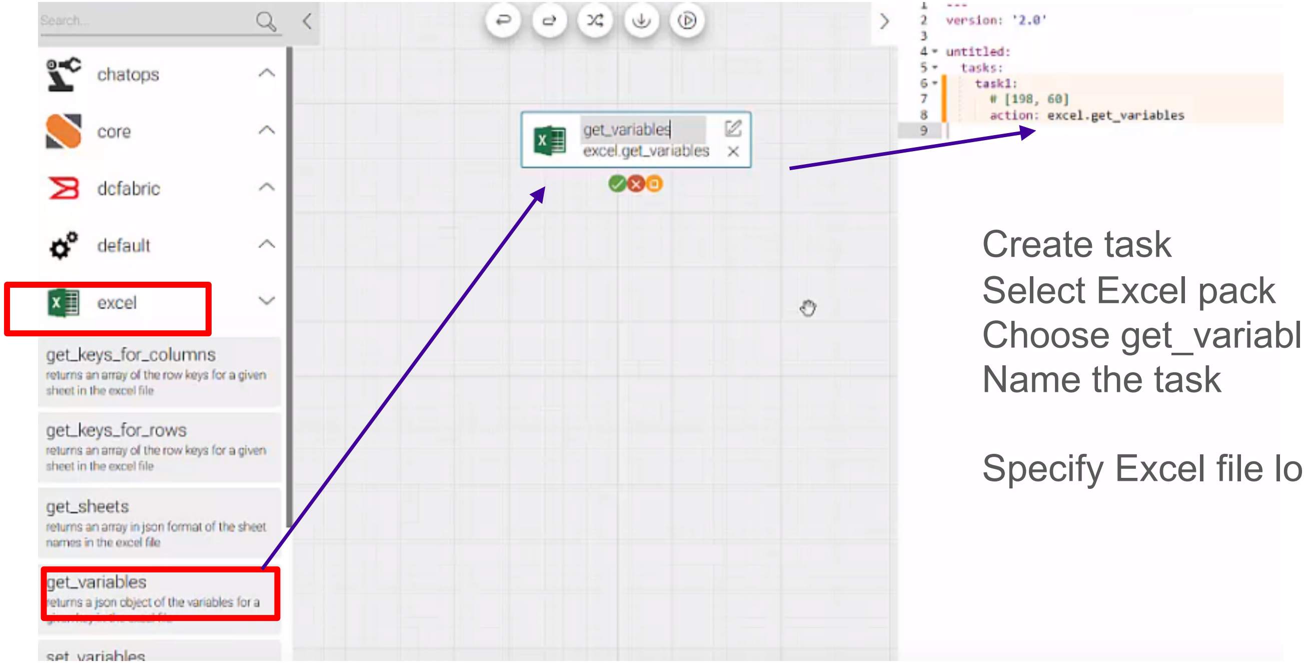

get_keys_for_columns: returns an array of the row keys for a given sheet in the excel file

get_keys_for_rows: returns an array of the row keys for a given sheet in the excel file

get_sheets: returns an array in json format of the sheet names in the excel file

get_variables: returns a json object of the variables for a

```
---
version: '2.0'

untitled:
  tasks:
    task1:
      # [198, 60]
      action: excel.get_variables
```

Create task
Select Excel pack
Choose get_variables
Name the task

Specify Excel file location in Add Metadata

Example #2 Excel Pack : Load Information in workflow

Add Excel Parameters for **inputs**

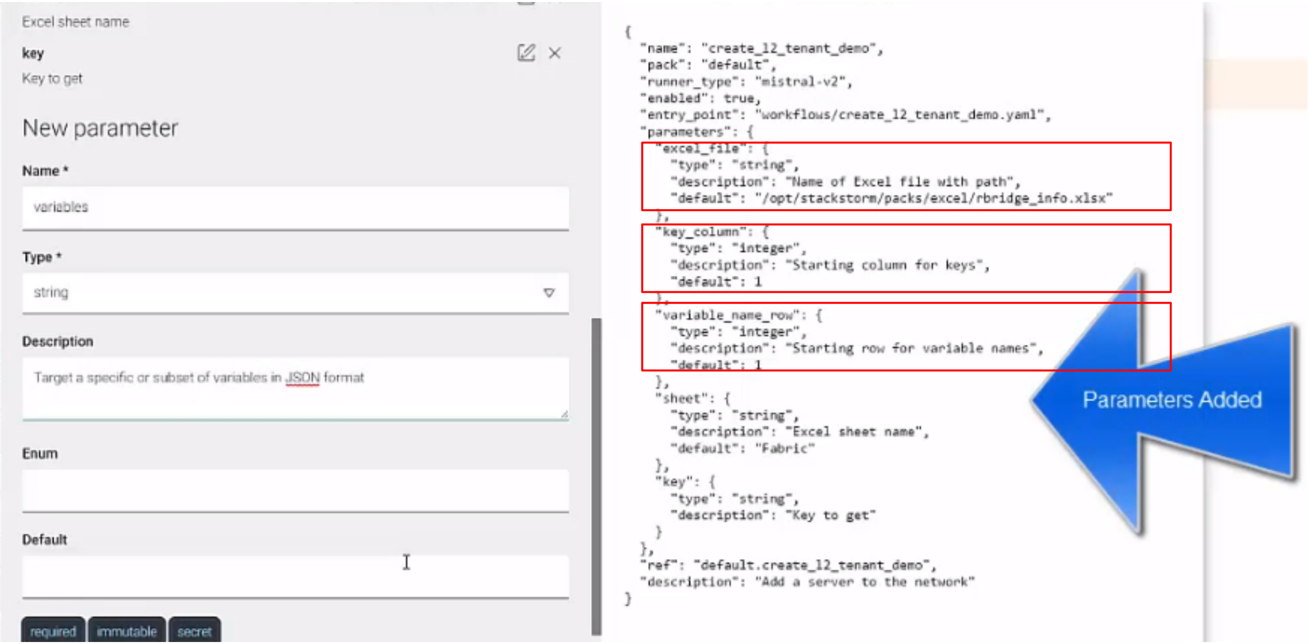

Excel sheet name

key

Key to get

New parameter

Name *

variables

Type *

string

Description

Target a specific or subset of variables in JSON format

Enum

Default

```
{
  "name": "create_l2_tenant_demo",
  "pack": "default",
  "runner_type": "mistral-v2",
  "enabled": true,
  "entry_point": "workflows/create_l2_tenant_demo.yaml",
  "parameters": {
    "excel_file": {
      "type": "string",
      "description": "Name of Excel file with path",
      "default": "/opt/stackstorm/packs/excel/rbridge_info.xlsx"
    },
    "key_column": {
      "type": "integer",
      "description": "Starting column for keys",
      "default": 1
    },
    "variable_name_row": {
      "type": "integer",
      "description": "Starting row for variable names",
      "default": 1
    },
    "sheet": {
      "type": "string",
      "description": "Excel sheet name",
      "default": "Fabric"
    },
    "key": {
      "type": "string",
      "description": "Key to get"
    }
  },
  "ref": "default.create_l2_tenant_demo",
  "description": "Add a server to the network"
}
```

Parameters Added

Example #2 Excel Pack : Create Workflow

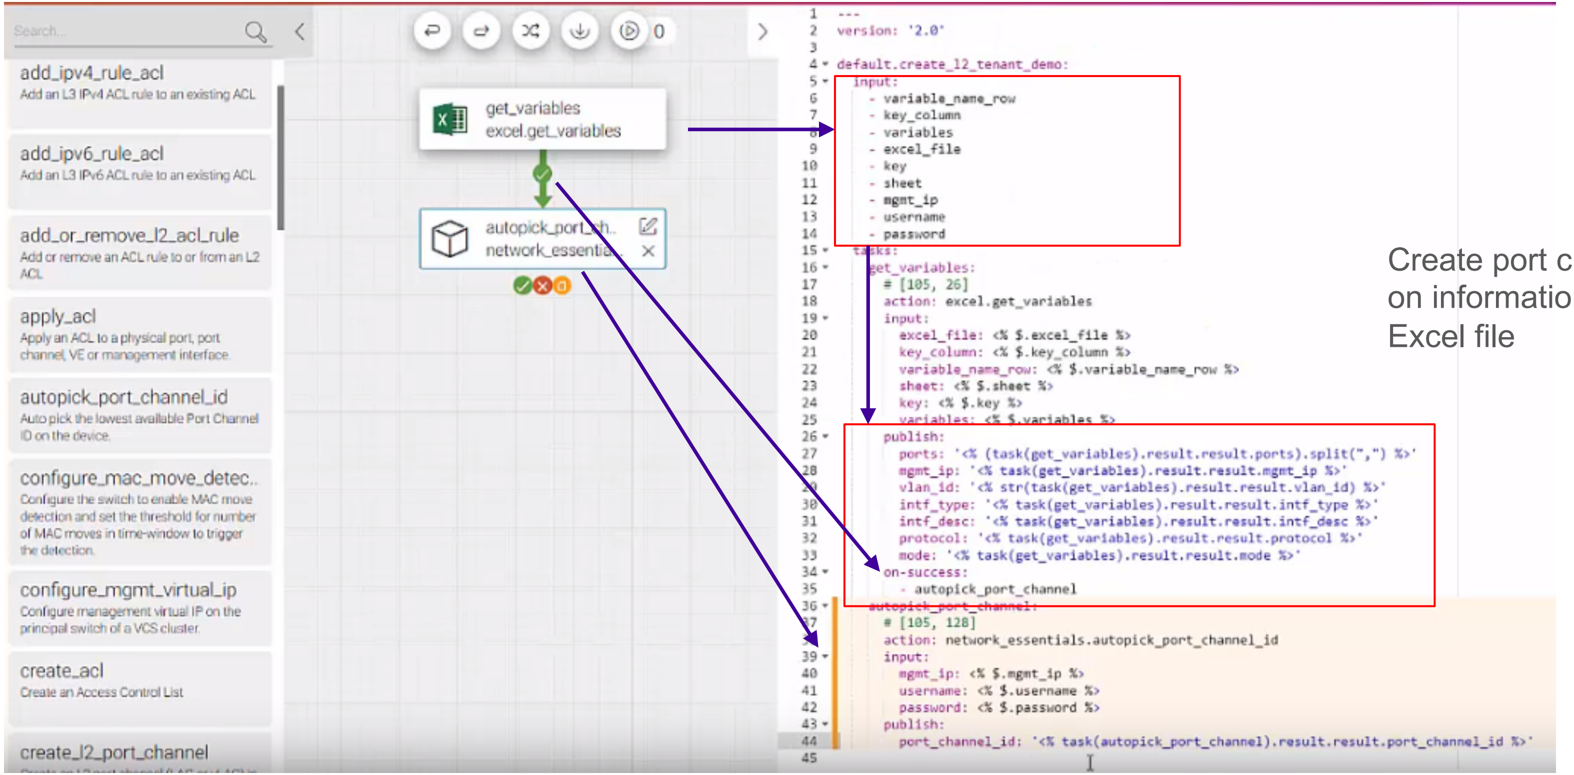

add_ipv4_rule_acl
Add an L3 IPv4 ACL rule to an existing ACL

add_ipv6_rule_acl
Add an L3 IPv6 ACL rule to an existing ACL

add_or_remove_l2_acl_rule
Add or remove an ACL rule to or from an L2 ACL

apply_acl
Apply an ACL to a physical port, port channel, VE or management interface.

autopick_port_channel_id
Auto pick the lowest available Port Channel ID on the device.

configure_mac_move_detec...
Configure the switch to enable MAC move detection and set the threshold for number of MAC moves in time-window to trigger the detection.

configure_mgmt_virtual_ip
Configure management virtual IP on the principal switch of a VCS cluster.

create_acl
Create an Access Control List

create_l2_port_channel

```yaml
---
version: '2.0'

default.create_l2_tenant_demo:
  input:
    - variable_name_row
    - key_column
    - variables
    - excel_file
    - key
    - sheet
    - mgmt_ip
    - username
    - password
  tasks:
    get_variables:
      # [105, 26]
      action: excel.get_variables
      input:
        excel_file: <% $.excel_file %>
        key_column: <% $.key_column %>
        variable_name_row: <% $.variable_name_row %>
        sheet: <% $.sheet %>
        key: <% $.key %>
        variables: <% $.variables %>
      publish:
        ports: '<% (task(get_variables).result.result.ports).split(",") %>'
        mgmt_ip: '<% task(get_variables).result.result.mgmt_ip %>'
        vlan_id: '<% str(task(get_variables).result.result.vlan_id) %>'
        intf_type: '<% task(get_variables).result.result.intf_type %>'
        intf_desc: '<% task(get_variables).result.result.intf_desc %>'
        protocol: '<% task(get_variables).result.result.protocol %>'
        mode: '<% task(get_variables).result.result.mode %>'
      on-success:
        - autopick_port_channel
    autopick_port_channel:
      # [105, 128]
      action: network_essentials.autopick_port_channel_id
      input:
        mgmt_ip: <% $.mgmt_ip %>
        username: <% $.username %>
        password: <% $.password %>
      publish:
        port_channel_id: '<% task(autopick_port_channel).result.result.port_channel_id %>'

```

Create port channel based on information of ports from Excel file

Agility through Automation and Visibility

Accelerating mean-time-to-innocence through automation for IXPs

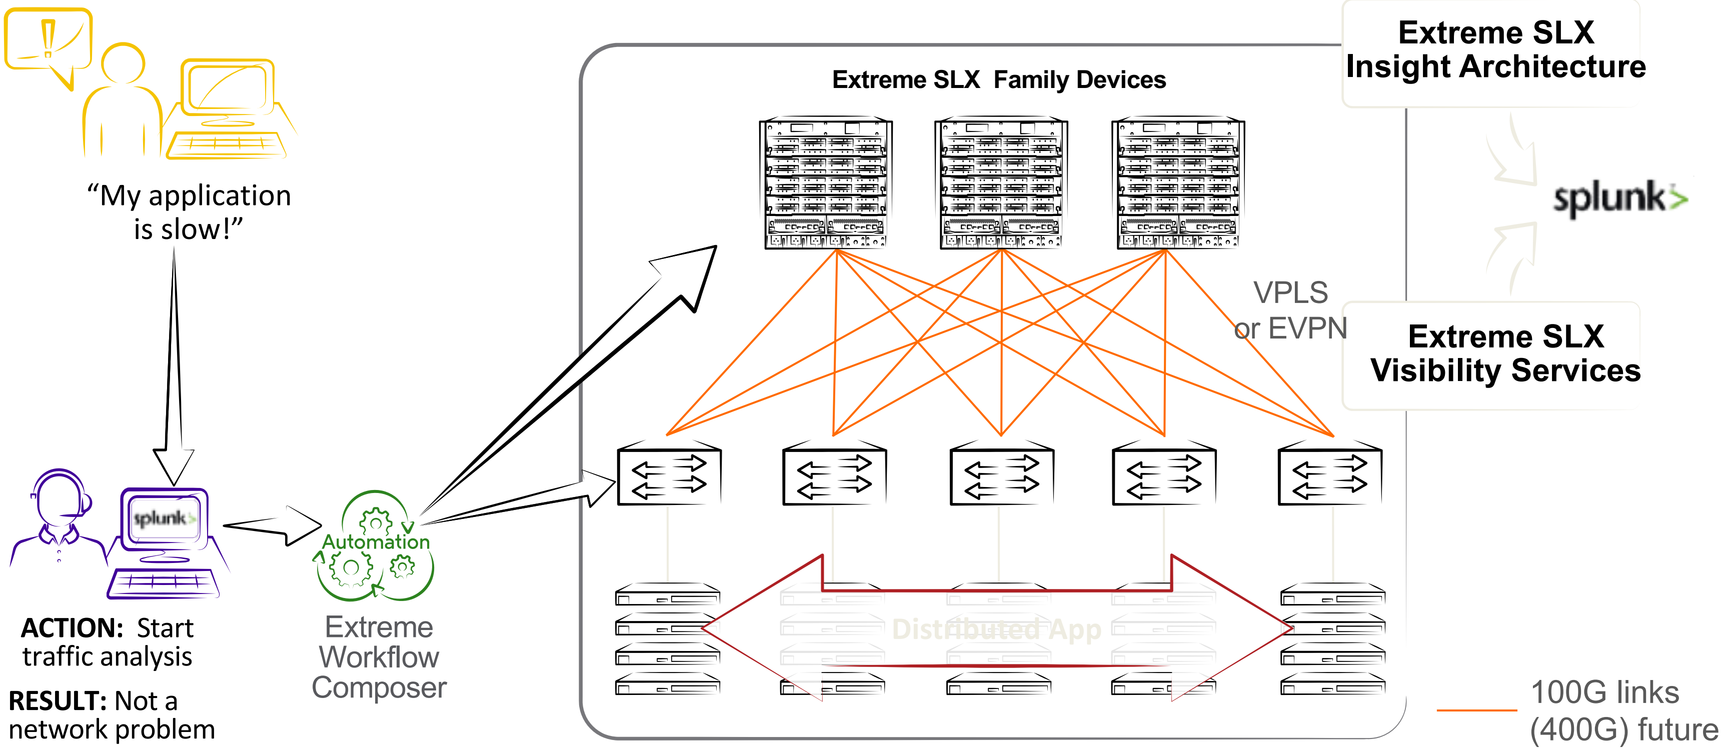

“My application is slow!”

ACTION: Start traffic analysis

RESULT: Not a network problem

Extreme Workflow Composer

Extreme SLX Family Devices

VPLS or EVPN

Extreme SLX Insight Architecture

Extreme SLX Visibility Services

100G links (400G) future

Agility through Automation and Visibility

Accelerating mean-time-to-innocence through automation for IXPs

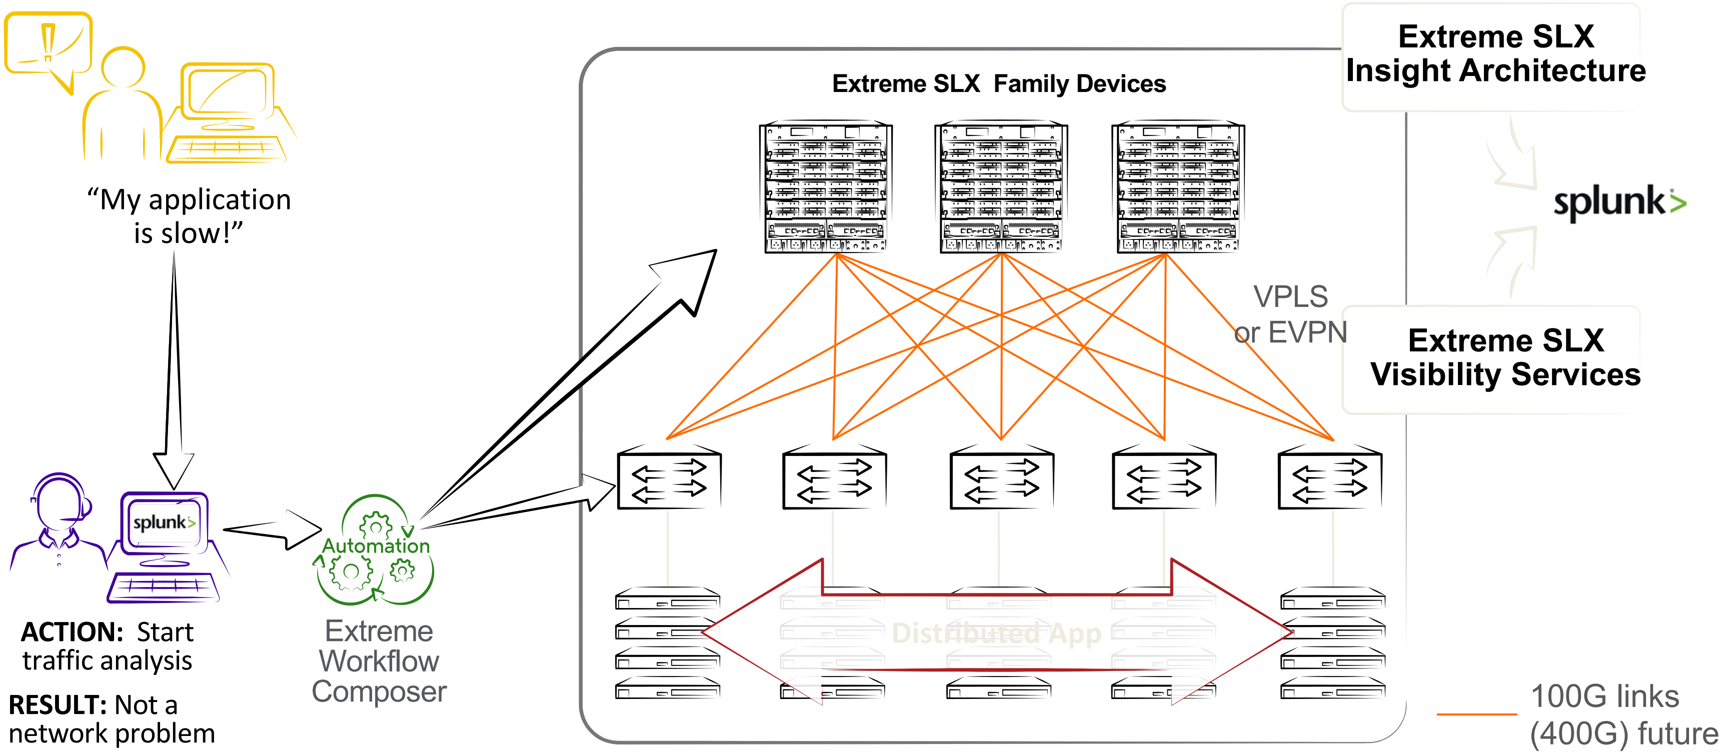

"My application is slow!"
splunk>
ACTION: Start traffic analysis
RESULT: Not a network problem
Automation
Extreme Workflow Composer
Extreme SLX Family Devices
VPLS or EVPN
Distributed App
Extreme SLX Insight Architecture
splunk>
Extreme SLX Visibility Services
100G links (400G) future

Extreme®

Connect Beyond the Network

Thank You

WWW.EXTREMENETWORKS.COM