Drain a Linecard for a <redacted>

Use this procedure to drain a linecard for a <redacted>.

Determine the Linecard Type

Verify if the card you are draining is an <redacted>.
* Typically linecards 15-18 are the core linecards and the rest are RSW facing.
You can verify the linecard type by logging in and running this command. It will show you what the Linecard is attached to: <code>show interface description | grep Eth<Linecard #></code>

Draining a Core-facing Linecard

Verifying Available Capacity

To determine if there is enough capacity to drain the linecard, use Metroid Prime to check bandwidth.

1. Fill out the Device Name text box with the CSW name.
2. Fill out the Interface Name text box with the following – “Ethernet<module#>.
3. Select [7d] from the Time Period selector. Doing so ensures that we check for trends and make sure enough capacity will be available under normal circumstances to complete the drain, replacement, and undrain process.

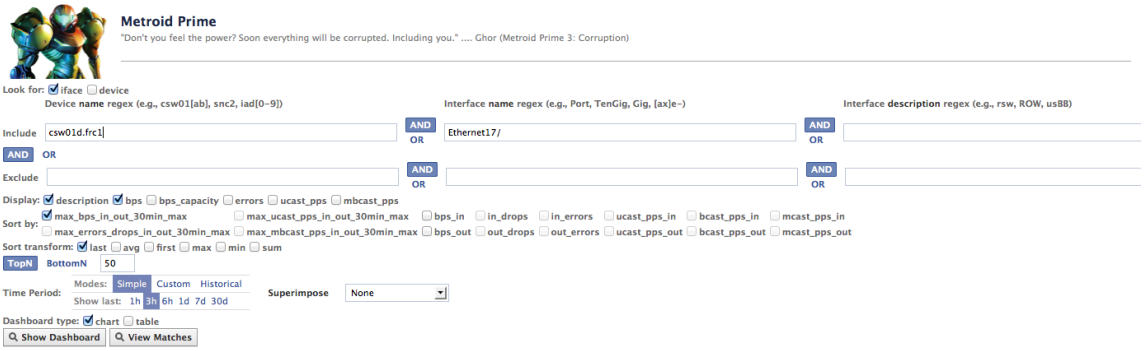

4. Click the Show Dashboard button. Doing so will take you to a display of each physical port on the switch and the amount of traffic traversing that link. Keep this dashboard showing as you will use it for reference.

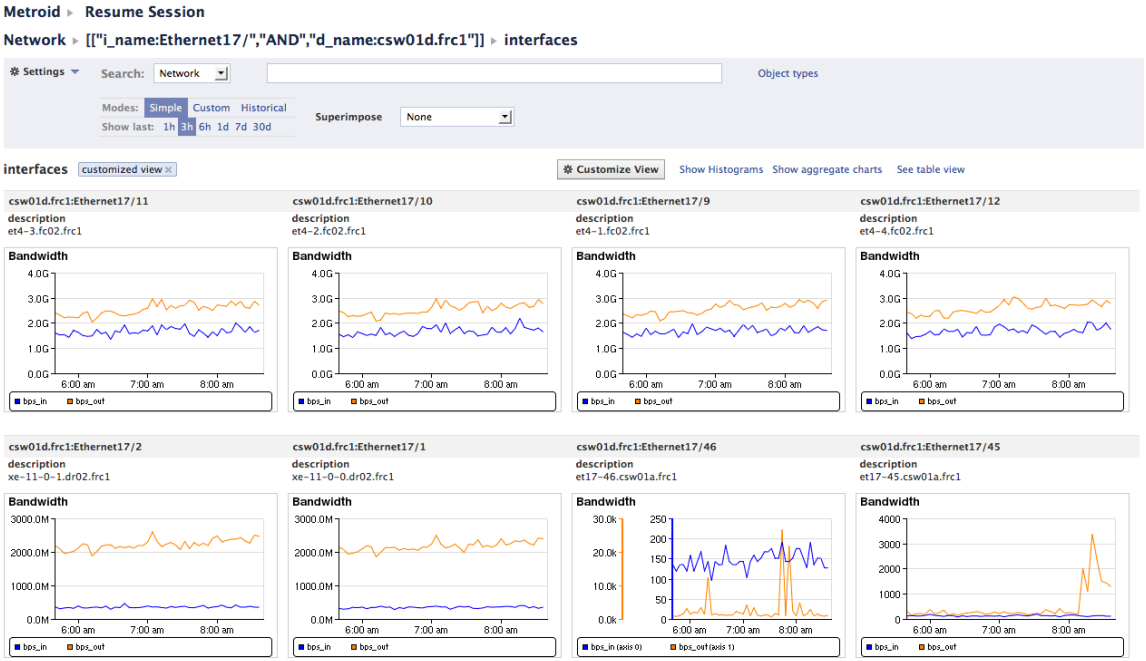

5. Calculate the maximum bandwidth for the core-facing CSW linecard link in the cluster. Ensure the amount of capacity currently in use is not over 65% of the total capacity for a four-post cluster and 40% for a two-post cluster.
    a. To calculate maximum bandwidth for a port-channel within the cluster, first find the port-channel number associated with those interfaces using the following command:

```
show port-channel summary
```

```
csw09a.frc1# sh port-channel summary
Flags:  D - Down        P - Up in port-channel (members)
        I - Individual  H - Hot-standby (LACP only)
        s - Suspended   r - Module-removed
        S - Switched    R - Routed
        U - Up (port-channel)
        M - Not in use. Min-links not met
--------------------------------------------------------------------------------
Group Port-       Type     Protocol  Member Ports
      Channel
--------------------------------------------------------------------------------
911   Po911(RU)   Eth      LACP      Eth18/9(P)   Eth18/10(P)  Eth18/11(P)
                                     Eth18/12(P)
912   Po912(RU)   Eth      LACP      Eth17/9(P)   Eth17/10(P)  Eth17/11(P)
                                     Eth17/12(P)
913   Po913(RU)   Eth      LACP      Eth16/9(P)   Eth16/10(P)  Eth16/11(P)
                                     Eth16/12(P)
914   Po914(RU)   Eth      LACP      Eth15/9(P)   Eth15/10(P)  Eth15/11(P)
                                     Eth15/12(P)
1019  Po1019(RU)  Eth      LACP      Eth15/1(P)   Eth15/2(P)
1020  Po1020(RU)  Eth      LACP      Eth16/1(P)   Eth16/2(P)
1021  Po1021(RU)  Eth      LACP      Eth17/1(P)   Eth17/2(P)
1022  Po1022(RU)  Eth      LACP      Eth18/1(P)   Eth18/2(P)
1023  Po1023(RU)  Eth      LACP      Eth17/45(P)  Eth17/46(P)  Eth17/47(P)
                                     Eth17/48(P)  Eth18/45(P)  Eth18/46(P)
                                     Eth18/47(P)  Eth18/48(P)
1024  Po1024(RU)  Eth      LACP      Eth15/45(P)  Eth15/46(P)  Eth15/47(P)
                                     Eth15/48(P)  Eth16/45(P)  Eth16/46(P)
                                     Eth16/47(P)  Eth16/48(P)
3000  Po3000(SU)  Eth      LACP      Eth17/41(P)  Eth18/41(P)
3001  Po3001(SU)  Eth      LACP      Eth18/21(P)  Eth18/22(P)
```

6. Find any ports that are also tied to the port-channel you identified in the last step.
7. Multiply the total links found by 10. That is the capacity on the CSW.
8. Multiply the amount of capacity on each CSW by how many CSWs are in the cluster (typically 4). That is the total cluster capacity for that port-channel type.

   NOTE: This is not true for the LB0[1-9][AB] linecard. The LB0[1-9]A links only to CSW-A. The LB0[1-9]B linecard links only to CSW-B. Multiply by 2.

9. Use Metroid Prime to find the maximum amount of bandwidth usage on each port-channel member over the past seven days. Add the bandwidth for each channel together to get a grand total. Typically, the amount on one member is very similar to the next, but you should verify each one.
10. Multiply the total of the maximum usage for each port-channel member by the number of CSWs in that cluster. This calculation will give you the total maximum cluster capacity usage.

    NOTE: Multiply by 2 when you calculate this value.

11. If the total cluster capacity usage max is less than 70-percent of the total cluster capacity, it should be safe to drain that port-channel.
12. Repeat these steps for each port-channel type.

    Example –

    Port-channel911 has 4 members.
    - 4 x 10 = 40G capacity on the CSW
    - 4 CSWs in this cluster – 40G x 4 =160G total cluster capacity
    - 3G traffic on each member at peak= 12G total for max usage
    - 12G x 4 = 48G total cluster max usage
    - 48G/160G = 30% usage max

    "'You can drain"'.

If all port-channels pass the capacity check, you can start the drain.

Squelching the Alarms

1. Use the *fbnet_suppressor* tool to prevent alarms to NetOps once maintenance begins.
2. From your shell, use the following command:
   ```
   fbnet_supressor --device=<CSW name> --time=<time in minutes that the full drain should take> --reason=<why?(typically task number will suffice)>
   ```

   Example
   ```
   fbnet_suppressor --device=csw09a.frc1 --time=60 --reason=2456129
   ```

Draining the Card

1. Notify NetOps (#netops) via IRC of your intent to drain the linecard. Reference the task that prompted the drain. Unless there is a major ongoing issue, they should approve your request.
2. Return to the port-channel summary screen.

```
csw09a.frc1# sh port-channel summary
Flags:  D - Down        P - Up in port-channel (members)
        I - Individual  H - Hot-standby (LACP only)
        s - Suspended   r - Module-removed
        S - Switched    R - Routed
        U - Up (port-channel)
        M - Not in use. Min-links not met
--------------------------------------------------------------------------------
Group Port-       Type     Protocol  Member Ports
      Channel
--------------------------------------------------------------------------------
911   Po911(RU)   Eth      LACP      Eth18/9(P)   Eth18/10(P)  Eth18/11(P)
                                     Eth18/12(P)
912   Po912(RU)   Eth      LACP      Eth17/9(P)   Eth17/10(P)  Eth17/11(P)
                                     Eth17/12(P)
913   Po913(RU)   Eth      LACP      Eth16/9(P)   Eth16/10(P)  Eth16/11(P)
                                     Eth16/12(P)
914   Po914(RU)   Eth      LACP      Eth15/9(P)   Eth15/10(P)  Eth15/11(P)
                                     Eth15/12(P)
1019  Po1019(RU)  Eth      LACP      Eth15/1(P)   Eth15/2(P)
1020  Po1020(RU)  Eth      LACP      Eth16/1(P)   Eth16/2(P)
1021  Po1021(RU)  Eth      LACP      Eth17/1(P)   Eth17/2(P)
1022  Po1022(RU)  Eth      LACP      Eth18/1(P)   Eth18/2(P)
1023  Po1023(RU)  Eth      LACP      Eth17/45(P)  Eth17/46(P)  Eth17/47(P)
                                     Eth17/48(P)  Eth18/45(P)  Eth18/46(P)
                                     Eth18/47(P)  Eth18/48(P)
1024  Po1024(RU)  Eth      LACP      Eth15/45(P)  Eth15/46(P)  Eth15/47(P)
                                     Eth15/48(P)  Eth16/45(P)  Eth16/46(P)
                                     Eth16/47(P)  Eth16/48(P)
3000  Po3000(SU)  Eth      LACP      Eth17/41(P)  Eth18/41(P)
3001  Po3001(SU)  Eth      LACP      Eth18/21(P)  Eth18/22(P)
```

If a port-channel's members are all on the same card, you cannot use BGP draining. Otherwise, shut down the members.

BGP Draining

For any BGP draining, you will first need the BGP autonomous system (AS) number of the CSW we are on.

1. Use the following command to find the AS number:

   `show run bgp | grep "router bgp"`

   ```
   csw09a.frc1# sh run bgp | grep "router bgp"
   router bgp 65249
   ```

   NOTE: Keep this code sample. You will use it for each BGP drain.

2. For any port-channel where all the members are on the same linecard (except <redacted>), you will need to find the IP address of the other end. To find the IP address, run the following command:

   ```
   show run interface port-channel<number>
   ```

   ```
   csw09a.frc1# sh run int po911

   !Command: show running-config interface port-channel911
   !Time: Tue Jul  9 13:10:12 2013

   version 6.1(2)

   interface port-channel911
     description po1090.fc01.frc1
     load-interval counter 2 30
     no ip redirects
     ip address 10.46.152.64/31
     ipv6 address 2401:db00:e211:4001:1090::/127
   ```

   Note the corresponding interface/bundle from the description. Check the BGP config:

   ```
   csw09a.frc1# sh run bgp | sec nei | inc po1090.fc01.frc1 prev 2
     neighbor 10.46.152.65
       inherit peer FC
       description po1090.fc01.frc1
   ```

   The IP address after the neighbor statement is the IP address of the corresponding bundle.

3. Repeat this process for each port-channel that you need to BGP drain.
4. Now that we have the AS number and BGP neighbor IP addresses, we can apply the route maps to drain the traffic. These commands will need to be executed for each neighbor address.
5. To drain the port-channel, run the following command:

   ```
   configure terminal
      router bgp <AS Number>
      neighbor <neighbor IP address>
      address-family ipv4 unicast
      route-map NOTHING in
   ```

```
     route-map NOTHING out
address-family ipv6 unicast
     route-map NOTHING-V6 in
     route-map NOTHING-V6 out
      end
```

6. The ports in that bundle should begin to lose traffic. You can check via refreshing Metroid or by running the following code:

   ```
   show int port-channel<port-channel #> | grep rate |
   grep bps
   ```

   ```
   csw09a.frc1(config-router-neighbor)# sh int po911 | grep rate | grep bps
      input rate 16.49 Gbps, 3.81 Mpps; output rate 11.11 Gbps, 4.70 Mpps
   ```

Once the traffic is completely gone (minus-maybe-a couple bps) you are done with BGP drains.

Draining LB0[1-9][AB]

For now, work with TI and NetOps for this drain.

Non-BGP Draining

This draining is for all other port-channels that have links on more than one linecard, which is typically CSW to CSW.

1. To drain non-BGP cards, shut down the entire port-channel
2. Configure the terminal
   ```
   interface port-channel<port-channel number>
               shutdown
               exit
   ```
3. The ports in that bundle should begin to lose traffic. You can check via refreshing Metroid or by running the following code:

   ```
   show int port-channel<port-channel #> | grep rate |
   grep bps
   ```

   ```
   csw09a.frc1(config-router-neighbor)# sh int po911 | grep rate | grep bps
      input rate 16.49 Gbps, 3.81 Mpps; output rate 11.11 Gbps, 4.70 Mpps
   ```

All draining should be complete after running the command.