Information Note

Resilient setup of digital Vega gateways

This information note provides details on:

- The physical configuration required for resilient installation of Vega gateways
- Information available from Vega gateways that indicate potential problems.

Introduction

In larger systems, especially call centre solutions where connectivity is directly related to financial success, it is advisable to use a resilient infrastructure, such that if a cable is damaged, a trunk goes down, power is lost or an IP element is lost, calls can continue to be handled at full capacity.

Hardware architecture

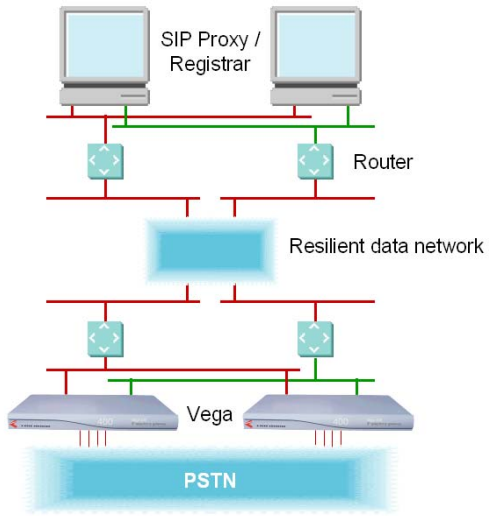

For a resilient hardware architecture, each element in the system must be duplicated, and connectivity replicated so that loss of any single component does not affect the capability of the whole system.

On the telephony interfaces at least 1 extra gateway beyond the number of gateways required for full capacity must be installed in order to allow a single gateway to be taken out of service.

Robust supply of power to devices is required, and this power should be driven from multiple UPSs.

Sending calls to the Vega gateways

In order to minimise problems with undetected / uncorrected errors, calls should be sent in a round-robin order to the Vega gateways. This also has the benefit that all gateways and trunks will be exercised so that problems will be detected as they occur, not just when a further element fails.

Calls which clear due to cause codes 34 or 41 should be re-presented by the SIP proxy to a different Vega gateway, as the gateway is reporting that it is incapable of presenting the call to the telephone network.

Similarly, of course, if a Vega is unreachable for any reason the call should be re-presented to a different Vega gateway.

On the LAN side, some calls should be sent over one LAN and others over the other LAN so that both LAN interfaces are exercised.

Configuring the Vega

When sending calls to the telephony network the Vega gateways should be configured to use Call Presentation Groups with their Round Robin setting, so that calls are routed sequentially through all interfaces to keep them tested. Dial plans should be written such that calls will roll from one trunk to another if the trunk is not in a state to make calls.

On the IP side, if possible DNS SRV should be used to define the registrar and proxy to use. A 25% weighting on each of the two interfaces on each of the two proxy servers should be used so that on average an equal number of calls are sent across each of the 4 LAN interfaces on the registrar / proxy servers.

If DNS SRV is not possible, then the 4 registrar / proxy interfaces should be defined as registrars and proxy servers in the Vega and the Vega configured to use the proxy servers in cyclic mode.

Static routes need to be configured on the Vega gateways so that 1 LAN interface on each proxy / registrar is accessed via Vega LAN interface 1 and the second LAN interface on each proxy / registrar is accessed via Vega LAN interface 2.

In order to allow timely switching if elements fail, set the following Vega parameters:

- sip.registrar.timeout_ms=5000
- sip.proxy.timeout_ms=5000
- sip.outgoing_call_setup_to=20000 (all 4 proxies can be tried before filing call with reason code 3)

N.B. In each Vega gateway its two LAN interfaces must be on separate subnets.

Determining problems on the Vega

If an element fails, then because other elements are there to take its place the challenge is to determine what has failed.

Vega gateways have a number of ways to allow problems to be detected. These include:

- SNMP traps
- Log display on output
- Call Detail Records

These should be permanently monitored to check from problem indications.

The system monitoring device should receive traps from all attached Vega gateways. These will report conditions like:

- Vega cold boot
- Vega warm boot
- Telephony trunk down / up
- LAN link down / up
- SIP Registration failed / succeeded[1]

Log display output can be used as an alternate source for Vega alarm messages (details are available in the Vega Primer). Log display output may be collected via telnet, serial or syslog.

Call related information can be obtained through the processing of call detail records. They should be analysed on a per trunk basis in order to identify problems with specific trunks.

Changes, or uneven-ness across the system of the following per trunk values may indicate a problem:

- Average call duration
- Answered calls / presented calls ratio
- Average time to answer
- Number of hours of calls in last hour

Call cleardown codes that are not 16, 17, 18, or 19 may also indicate a problem.

Call detail records may be obtained from the Vega using telnet, serial, syslog or Radius.

Note where predictive diallers are used, calls cleared down by the predictive dialler should be given average call values to avoid skewing the statistics.

Contact Details
Email: support@vegastream.com
Web: www.vegastream.com
www.vegaassist.com

EMEA Office
VegaStream Limited
Western Centre
Western Road
Bracknell
Berks RG12 1RW
UK

+44 (0) 1344 784900

USA Office
VegaStream Inc.
6200 Stoneridge Mall Road
3rd Floor
Pleasanton
California 94588
USA

+1 925 399 6428

[1] To be available in the near future