

# A New VRM Architecture for N+1 Redundancy at the Point of Load

*Support output currents above 300A by adding the required number of phases*

*The preferred power distribution topology for servers and switches confines redundancy provision to card- or rack-level. A new POL topology supporting hot-swapping moves N+1 redundant design closer to individual ICs.*

*By George Schuellein, Director, DC/DC Applications Engineering, International Rectifier*

A distributed power architecture incorporating non-isolated point-of-load (POL) regulators is the solution of choice for systems containing modern ICs with multiple power domains, low operating voltages and high current requirements. But conventional regulators are not able to support N+1 redundant power design at the point of load.

For telecom, internet and server equipment, where 99.999% (or five-9s) availability is required, this has prevented designers from implementing redundancy on a per-rail basis

and has restricted redundant design to the board or rack level. This results in at least two disadvantages. Firstly, the entire board or rack must be replaced if an individual component fails. The second disadvantage is the loss of customer data being processed at the time of the failure.

The key deficiencies in a conventional POL converter, from the redundancy point of view, concern the absence of hot-swap capability, including the need to isolate the input and output.

## N+1 Redundancy at Point of Load

To address these limitations of conventional POLs, International Rectifier has developed an N+1 redundant POL Voltage Regulator Module (VRM) by combining a new, dedicated control IC and companion phase IC with input MOSFETs allowing hot-swapping and output MOSFETs for ORing.

Figure 1 shows how the IR3510 control IC and IR3088 phase IC create an efficient synchronous buck converter providing system protection against failures such as short circuits. The chipset provides input isolation, which allows hot-swapping of power modules while enabling 100% availability of power to the system.

The topology is scalable to support output currents above 300A by adding the required number of phases. The IR3510 provides overall system control and interfaces with multiple IR3088A phase ICs through a five-wire analogue bus. For its part, the IR3088 implements the functions required by the converter of each phase. These include the gate drivers, PWM comparator and latch, over-voltage protection, and current sensing and sharing.

From the system perspective, this architecture reduces redundant silicon in multiphase systems compared to other solutions such as a four-phase controller that can be configured for two-, three- or four-phase operation. PCB layout is also easier, since short gate-drive and current-sense connections can be implemented locally to the phase ICs.




Figure 1. N+1 redundant power at point of load.

### Chipset Operation

The IR3510 controls the input and output FETs, using an on-chip charge pump supplying 13V of gate drive to ensure full MOSFET enhancement. The output transistors can be paralleled as required to support the desired input and output current rating.

The input MOSFET provides both hot-swap capability and input overcurrent protection. Converter input current is accurately sensed across a high-side current-sensing resistor. A programmable linear current limit controls the input current during inrush but will latch off the input MOSFET if the overcurrent persists for longer than a programmable delay period.

Protection for a short-circuit across the input MOSFET is implemented by detecting whether the input voltage supplied to the Synchronous Buck power stage rises prior to turn-on of the charge pump. If this occurs operation of the Synchronous Buck is inhibited and the error condition communicated to the host system.

At the output of the system, the body diode of the ORing MOSFET allows current to flow to the load. The IR3510's active ORing control function senses the voltage drop across the Source-Drain terminals and will apply gate drive to turn the device on if a programmable threshold is exceeded. Reverse current, which may be caused by a synchro-

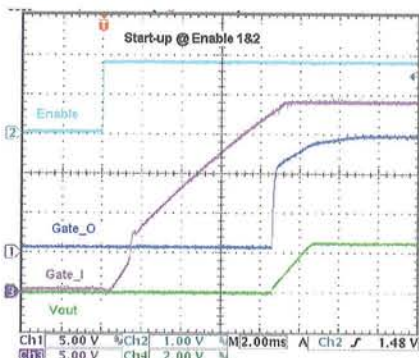
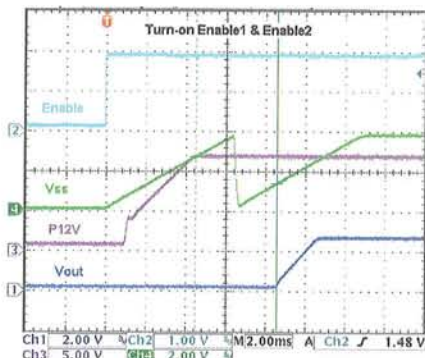


Figure 2a and 2b: Protection and Failover Capabilities

nous MOSFET short-circuit failure or other converter malfunction, will cause the controller to quickly turn off the ORing MOSFET. The output voltage will then be supplied by one or more other redundant POLs.

To maximise MTBF, the IR3510 uses average current mode control to implement active droop sharing between converters. This eliminates the need for any control wires connecting between two or more current sharing POLs and prevents a potential single-point-failure mode from disabling the system. An analogue current monitor signal is provided and allows the system microcontroller to monitor the current being delivered by each POL. A digital interface is also included to transmit data back to the host system indicating status of the POL's output voltage, input MOSFETs, and ORing MOSFETs.

### Soft-Start Sequence

The IR3510 implements a programmable two-stage soft-start with timing programmed by an external capacitor. Typical start-up waveforms are shown in Figures 2a and 2b. Following receipt of an Enable signal from the host, the soft start capacitor begins linearly charging from ground. The voltage on the soft start capacitor ( $V_{SS}$ ) is used to control the gate of the input MOSFET ( $Gate_I$ ). This provides closed-loop control, ensuring a controlled ramp of the Synchronous Buck input voltage ( $P12V$ ).

After the input voltage has fully ramped the soft-start capacitor is discharged and a 2nd soft start is initiated to ramp the POL output voltage. The charge pump continues to boost the voltage on the input MOSFETs gate up to  $V_{in}+13V$  to ensure it is fully enhanced. The rising output voltage causes current to flow out of the POL and turns on the gate of the ORing MOSFET ( $Gate_O$ ). The charge pump continues to boost this gate voltage up to  $V_{out} + 13V$ .

A number of faults can give rise to an output over voltage condition, including a shorted control MOSFET or open-circuit control loop. The IR3510 provides an independent input for an over-voltage comparator whose threshold is programmable. Since the remote sense lines cannot be relied upon to provide valid sense data in the event of a fault, over-voltage is typically sensed locally within the POL.

When operating in N+1 mode the goal is to identify and disable the failed POL while keeping "good" POL(s) operating and maintaining a stable output voltage.

The "good" POLs will act to maintain regulation of the output voltage. To this end, they will sink current from the output as they try to clamp the over voltage. This causes the ORing MOSFETs to turn-off, which effectively disconnects them from the load and prevents damage to the "good" POLs. The Input MOSFET(s) of the "good" POL remain on.

"Bad" POLs will source current such that the associated their ORing MOSFETs will be turned on when the over-voltage threshold is exceeded. This provides the indication they are "bad". In this case the input MOSFETs are quickly turned-off, removing the source of the overvoltage event. Once the output voltage has dropped down to normal levels the remaining "good" POLs will begin switching and providing positive current to the load. The "bad" POL will remain latched off. The host system can then be notified and a service call scheduled to replace the "bad" POL.

### N + 1 Transition

POLs are enabled and disabled from an operating N+1 redundant system. To replace a POL it is first disabled and then physically removed from the equipment, typically with a connector designed to facilitate this. The replacement POL is then inserted into the connector and enabled. The ORing MOSFETs will block negative current in an enabled POL until its output voltage exceeds the already regulated output voltage.

During the transition from N to N+1, the output voltage will overshoot during the period when POLs that are already operating reduce their output current to match the newly enabled POL. Conversely, a small undershoot occurs during N+1 to N transitions while the remaining POL(s) increase their output current to the level required by the load.

### Conclusion

While infrastructure owners count the total cost of system downtime in thousands or millions of dollars, wholesale replacement of cards or racks is also expensive. Support for N+1 redundancy at the POL level now combines the advantages of high system reliability with economic replacement of individual components.

[www.irf.com](http://www.irf.com)