Using Local Energy Storage For Organized System Shutdown Simplifies Power Supply Hold-Up Time Requirements

by Viktor Vogman, Power Conversion Consulting, Olympia, Wash.

Many instrument and computer systems require early warning of an imminent power failure. Such a warning is required to provide sufficient time for safe data transfer in an organized system shutdown process and subsequent recovery without critical data loss. During this self-deactivation process, the supply voltages must be maintained above their minimum specified values, which requires that sufficient energy be stored in the power supply unit (PSU) over the specified time interval.

The power fault warning is usually provided by power supply control circuitry that monitors ac line and/or PSU bulk capacitor voltage and then generates different types of alerts, such as ACOK and DCOK signals. Power Good, or PWROK (which is essentially the terminology used for DCOK in computers and servers), is the most commonly used alert signal.

The Power Good warning time represents a delay starting from this signal de-assertion to the time at which supply voltages drop out of their regulation limits. By changing its logic state, the alert signal indicates that energy stored in the PSU is sufficient only to keep supplied voltages above minimum for a certain time interval. This time interval essentially represents a portion of the PSU hold-up time which is needed for an organized system shutdown and which is usually specified between 1 and 10 ms.[1, 2] Typical PSU major voltage and signal waveforms in this mode are shown in Fig. 1.

Traditionally, the energy needed for organized system shutdown is stored in the PSU bulk capacitor, which supplies power to the entire system during an ac fault event. In high-density PSUs—even with the SmaRT technology feature[3] easing the hold-up time requirements—this capacitor can occupy anywhere from 5% to 20% of the PSU module space.

For a few-millisecond increase in hold-up time duration, hold-up time extension circuits on the primary side can be used. Such circuits are able to extract more energy in the shutdown mode by using an auxiliary cap that can be discharged to lower voltage levels.[4] However, the use of such circuits offers very limited hold-up time extension, while also decreasing efficiency in the normal operating state.

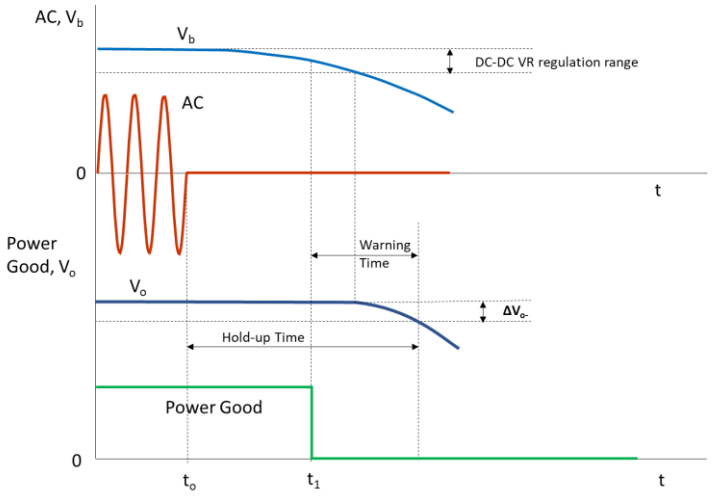

Fig. 1 Typical voltage and signal waveforms in server system shutdown mode. The Power Good warning time represents a delay from time t_1 when bulk capacitor voltage V_b reaches some critical level, and Power Good de-assertion to the time at which output voltage drops out of regulation limits, i.e. below V_o-ΔV_o..

When warning and/or hold-up time needs to be significantly extended (5 to 10 ms) to complete a large number of operations in an organized shutdown mode, a straightforward solution is to proportionally increase the PSU bulk cap size. But this greatly increases power supply size and cost, especially in server redundant subsystems using two or more PSU modules.

Another issue associated with power faults is insufficient warning time intervals under PSU hard failure situations, e.g. in the case of a malicious attack on the power subsystem, which results in an abrupt power-off condition. In such fault states the system power hold-up and warning time requirements cannot be met due to the PSU's unmanageable behavior.

This article studies an opportunity for easing the PSU hold-up time requirements for applications that require significantly extended warning time intervals. It presents the conditions for keeping the PSU bulk cap and PSU design unchanged and for effective reallocation of stored energy between the bulk cap and a local system buffer capacitor. The article shows how the local buffer cap can supply power just to the components critical to the organized system shutdown and increase system immunity to indiscriminate power faults.

Capacitor Core Volume

Since bulk capacitor size is considered the major factor in selecting the energy source during hold-up time, let's determine the capacitor core volume as a function of the energy stored in it. The capacitance *C* of a parallel-plate capacitor can be determined by the formula:

$$C = \frac{\varepsilon \varepsilon_o A}{d}$$

where ε is dielectric permittivity, ε_0 is vacuum permittivity, *A* is plate area, and *d* is dielectric thickness. Solving for plate area *A* and substituting the result into the equation for capacitor core volume $Vol = d \cdot A$, we can determine *Vol* as follows:

$$Vol = d \cdot A = \frac{d \cdot d \cdot C}{\varepsilon \varepsilon_o} = \frac{C d^2}{\varepsilon \varepsilon_o}$$

Dielectric thickness *d* is directly proportional to the breakdown voltage V_{bd}, *such that* $d = V_{bd}/\sigma$, where σ is dielectric strength of the insulator material, which characterizes the maximum voltage required to produce a dielectric breakdown through the material. Dielectric strength is expressed in terms of volts per unit thickness.

Rated capacitor voltage V_{nom} is related to the breakdown level with a given derating factor k_d, *such that* $V_{nom} = k_d V_{bd}$. Combining these two relations and multiplying the numerator and denominator of the above expression for *Vol* by two, we can determine that capacitor core volume is directly proportional to the amount of energy stored between the plates at its rated voltage:

$$Vol = \frac{C \cdot V_{nom}^2}{\sigma^2 k_d^2 \varepsilon \varepsilon_o} \cdot \frac{2}{2} = \frac{E}{k_0} \qquad (1)$$

where $E = CV_{nom}^2/2$ is the energy stored in the cap at its rated voltage, and $k_0 = \sigma^2 k_d^2 \varepsilon \varepsilon_o / 2$ is a constant factor that characterizes the capacitor core energy density. The k_0 factor depends on cap dielectric material and capacitor manufacturing technology factors.

Advantages Of A High-Voltage Cap

As it is common practice to use a high-voltage-rated bulk capacitor in the PSU to store energy,[5] let's examine the justification for this traditional selection. Let's assume that during ac dropout the voltage across the PSU bulk cap C1 decreases from its nominal operating level V_1 by some value ΔV_1 while the PSU is supplying constant power P to the load. Using this assumption, we can write the following energy-balance equation:

$$\Delta E_1 = \frac{C_1 {V_1}^2}{2} - \frac{C_1 (V_1 - \Delta V_1)^2}{2} = \frac{C_1 \cdot (2V_1 - \Delta V_1) \cdot \Delta V_1}{2} = \frac{Pt}{Eff}$$

where Eff is the power supply efficiency at power level P and t is the time interval during which V_1 decreases by ΔV_1. Similarly, assuming that energy supplied to the load is entirely taken from the secondary-side cap C2:

$$\Delta E_2 = \frac{C_2 {V_2}^2}{2} - \frac{C_2 (V_2 - \Delta V_2)^2}{2} = \frac{C_2 \cdot (2V_2 - \Delta V_2) \cdot \Delta V_2}{2} = Pt$$

Dividing the numerator and denominator in these equations by V_1^2 and V_2^2, respectively, we find:

$$\frac{C_2 \cdot {V_2}^2 \left(2 - \widehat{\Delta V_2}\right) \cdot \widehat{\Delta V_2}}{2} = Pt \qquad (2)$$

$$\frac{C_1 \cdot {V_1}^2 \left(2 - \widehat{\Delta V_1}\right) \cdot \widehat{\Delta V_1}}{2} = \frac{Pt}{Eff} \qquad (3)$$

where $\widehat{\Delta V_1}, \widehat{\Delta V_2}$ are permissible cap voltage droops normalized to their nominal voltage levels.

Designating primary and secondary cap stored energies as $E_1 = CV_1^2/2$ and $E_2 = CV_2^2/2$, respectively and dividing (2) by (3) we find:

$$E_1 / E_2 = \frac{\left(2 - \widehat{\Delta V_2}\right) \cdot \widehat{\Delta V_2}}{Eff\left(2 - \widehat{\Delta V_1}\right) \cdot \widehat{\Delta V_1}}$$

Assuming that primary (high voltage) and secondary (low voltage) caps have different energy densities (k_{o1} and k_{o2}, respectively) and using equation (1), we can now determine the cap volume ratio as:

$$Vol_1 / Vol_2 = \frac{k_{02} E_1}{k_{01} E_2} = \frac{k_{02}\left(2 - \widehat{\Delta V_2}\right) \cdot \widehat{\Delta V_2}}{k_{01} \cdot Eff\left(2 - \widehat{\Delta V_1}\right) \cdot \widehat{\Delta V_1}} \qquad (4)$$

This equation shows that besides the allowed voltage droops $\widehat{\Delta V_1}, \widehat{\Delta V_2}$, the cap volume ratio also depends on the constant k_{02}/k_{01}, which characterizes capacitor core energy densities and is related to cap manufacturing technologies.

Let's evaluate the ($k_r = k_{02}/k_{01}$) ratio. As it can be seen from the left portion of equation (4), $k_{02}/k_{01} = E_2Vol_1/E_1Vol_2$ is actually a ratio of energy densities of the low-voltage (E_2/Vol_2) and high-voltage (E_1/Vol_1) caps. Assuming standard primary (450-V) and secondary (16-V) voltage ratings, let's evaluate this ratio for three major cap suppliers—Panasonic, Rubicon and Nichicon. These values are shown in a bar chart in Fig. 2.

Although these data points represent the overall (packaged) cap energy density, for our purposes we can assume that this energy density is related to core energy density via a constant factor, determined by the capacitor package technology, so that ratio of the packaged cap energy densities for a given manufacturer is essentially equal to their core energy densities. From this chart we can conclude that the actual k_{02}/k_{01} constant factors ratio can vary between 0.205 and 0.35. In other words, the high-voltage bulk cap energy density is at least 2.86 times that of the low-voltage rated caps.

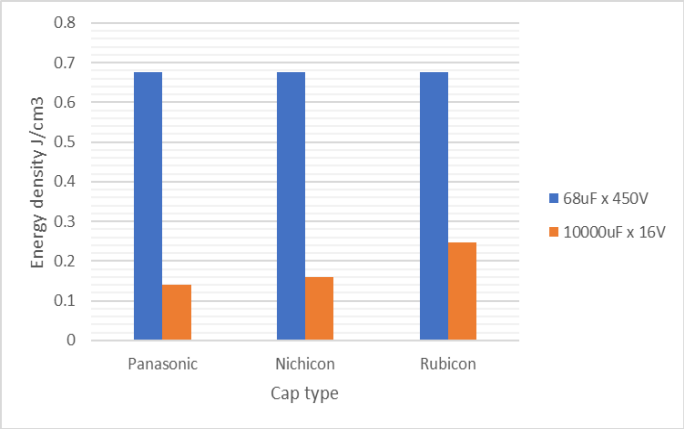

Fig. 2. High-voltage and low-voltage electrolytic cap energy densities from three major cap suppliers. High-voltage bulk cap energy density is at least 2.86 times that of the low-voltage rated caps.

The standard design guideline value for primary-side normalized voltage droop $\Delta\widehat{V_1}$ is 0.25, while the standard $\Delta\widehat{V_2}$ spec value is 0.05. Plugging in these numbers into equation (4) along with a typical PSU efficiency Eff = 0.95 we find that $Vol_1/Vol_2 > 0.082$. This means that in a conventional case for a given power level the required primary-side energy storage cap is at least 12 times smaller in size than the secondary one. This constitutes the main reason for using the PSU bulk cap as the hold-up time energy storage component.

Can A Secondary-Side Energy Storage Solution Be Comparable?

The analysis above shows that in a conventional case using a high-voltage cap for energy storage has significant size advantages over secondary-side storage. Does this mean that a secondary-side energy storage option has no chance to compete?

Let's take a closer look at the above equations and identify the conditions under which a secondary-side energy storage solution could become comparable to the primary side. Equations (2), (3) and (4) show that the requirement to support the same full load power and much smaller specified voltage droop (i.e., 5% on the secondary side versus 25% on the primary side) makes the secondary-side option appear unfavorable—even if energy densities of high- and low-voltage caps are identical, i.e. when $k_{02}/k_{01} = 1$.

So what can be done to reverse this relationship? There are system components, such as memory, whose operation is critical in the shutdown mode, whereas other components like cooling fans are not critical. Based on equation (4), if circuitry with critical components can be separated from the power delivery path and can be made to operate with a supply voltage droop that is much greater than the supply voltage droop required for the entire system, then the ratio of primary and secondary cap volumes Vol_1/Vol_2 can drastically change.

Let's suppose that the power required for critical circuitry operation P_{cr} represents only a portion of total system power: $P_{cr} = aP$, where a characterizes a secondary-side buffered power share $(a < 1)$. Let's also assume that the allowed supply voltage droop $\Delta\widehat{V_{2CR}}$ can be increased over the original $\Delta\widehat{V_{2CR}} = b\Delta\widehat{V_2}$ ($b > 1$). If in the shutdown mode such circuitry is separated from the rest of the system power distribution, then for this critical circuitry, equation (4) can be rewritten as follows:

$$Vol_1/Vol_2 = \frac{k_r\left(2 - b\Delta\widehat{V_2}\right)\cdot b\Delta\widehat{V_2}}{a\cdot Eff\left(2 - \Delta\widehat{V_1}\right)\cdot\Delta\widehat{V_1}} \qquad (5)$$

Using ($k_r = k_{02}/k_{01}$), $\Delta\widehat{V_1}$, and $\Delta\widehat{V_2}$ values from the previous analysis we can now determine a and b ranges that could make primary and secondary energy storage components comparable in size. A graph illustrating the dependency of the cap volume ratio on the buffered power share a and relative increase in the secondary supply voltage droop b is shown in Fig. 3a. The shaded area in the graph identifies a and b ranges over which the secondary-side energy storage option has size advantages versus the primary side.

Equating the right side of equation (5) to unity allows us to establish the relationship between a and b to determine conditions under which primary and secondary storage cap volumes are equal. Or, in other words, the conditions under which the storage cap placed on the secondary side will be as effective as the primary-side cap in the PSU.

A graph illustrating this relationship is shown in Fig. 3b for the lowest and highest k_r limits. This chart demonstrates that the greater the allowed secondary-supply-voltage droop $b\Delta\widehat{V_2}$, the larger the share of total system power a that can be effectively buffered with a local energy storing cap on the secondary side.

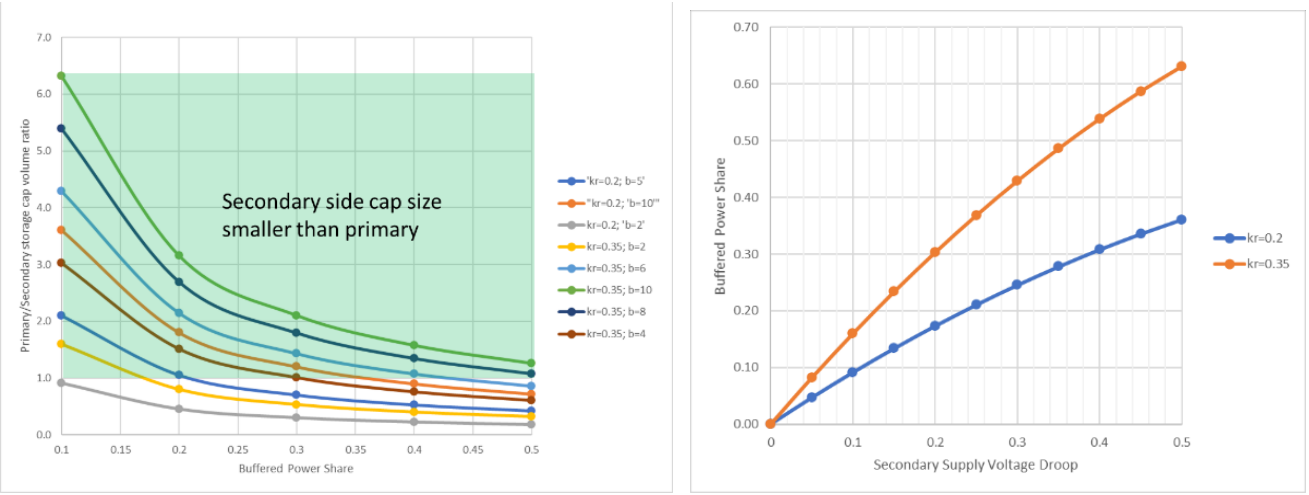

(a) (b)

Fig. 3. Graphs showing primary/secondary storage as a function of buffered power share (a) and the buffered power share as a function of allowed secondary-supply voltage droop (b). The shaded area in the first graph identifies a and b ranges for which the secondary-side energy storage option provides component size reduction versus the primary-side case (a). The larger the secondary voltage droop that is allowed, the larger buffered power share that can be effectively supported with secondary-side energy storage (b).

The curves in Fig. 3a demonstrate that benefits of the secondary energy storage increase as buffered power share reduces. An obvious explanation for this phenomenon is that with the same stored energy the time interval over which the lower power level can be supplied gets naturally extended. Increased voltage droop has

a similar effect on the hold-up time because it allows us to extract more energy from the same cap. Let's quantify the impact of the increased voltage droop factor.

An increase in the normalized delivered energy can be obtained from the left side of equation (3):

$$\widehat{\Delta E_2}(b\Delta V_2)/\widehat{\Delta E_2}(\Delta V_2) = \left(2 - \widehat{b\Delta V_2}\right) \cdot \widehat{b\Delta V_2}/[\left(2 - \widehat{\Delta V_2}\right) \cdot \widehat{\Delta V_2}]$$

Using this expression, we can plot the graph quantifying the increase of delivered energy over the conventional case $\left(\widehat{\Delta V_2} = 0.05\right)$. This graph is given in Fig. 4. As it can be seen from this plot, the delivered energy increase can be quite significant—3.7 to 8.6 times for a practically usable range of voltage droop $\left(\widehat{\mathrm{b}\Delta V_2} = 0.2 \; to \; 0.6\right)$.

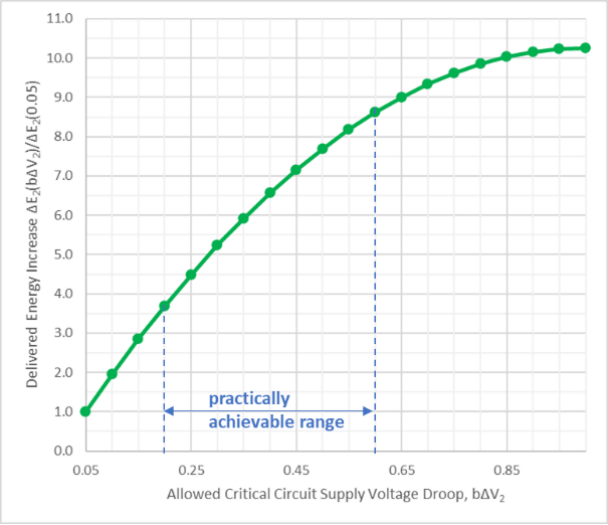

Fig. 4. An increase in the allowed critical circuit supply voltage droop provides a significant increase in the delivered energy over the conventional case ($\widehat{\Delta V_2} = 0.05$). The delivered energy can increase by a factor of 3.7 to 8.6 for the practically achievable range of voltage droop ($b\Delta V_2 = 0.2 \ldots 0.6$).

Arrangements For Buffering Critical Circuitry

Separating and buffering of power delivery to critical circuits can be provided with networks containing active and passive components. When buffered power share is relatively small, a buffering voltage source selector can be implemented with a passive network, such as the one shown in Fig. 5a.[6]

With this arrangement, in normal operating mode the power to the critical circuit is delivered through diode D1, while energy storing cap C_{CR} gets charged through the same diode and resistor R1. In the shutdown mode, the power to the critical circuit is supplied through diode D2. Charge and discharge current paths for the secondary storage cap are shown in the green and blue dashed lines, respectively.

For the portion of the circuitry that requires tight supply voltage, power delivery can be provided with a small boost regulator that tolerates larger voltage sags on its input (Fig. 5b). This voltage regulator (VR) operates only for a few milliseconds and its size can be very small, similar to an auxiliary boost VR used in high-voltage applications.[4]

For comparatively large buffered-power levels, for which the efficiency drop must be minimized under normal operating conditions, an active circuit using a low $R_{DS(ON)}$ p-channel power MOSFET controlled by a Power Good signal would have advantages over a passive component option due to its lower power dissipation. This circuit is shown in Fig. 5c.

During a shutdown caused by ac or PSU faults, components D1 (Fig. 5a, b) and Q1 (Fig. 5c) "isolate" the critical circuit energy storing cap C_{CR} from the rest of the system, which reserves its energy solely for operation of the critical circuit.

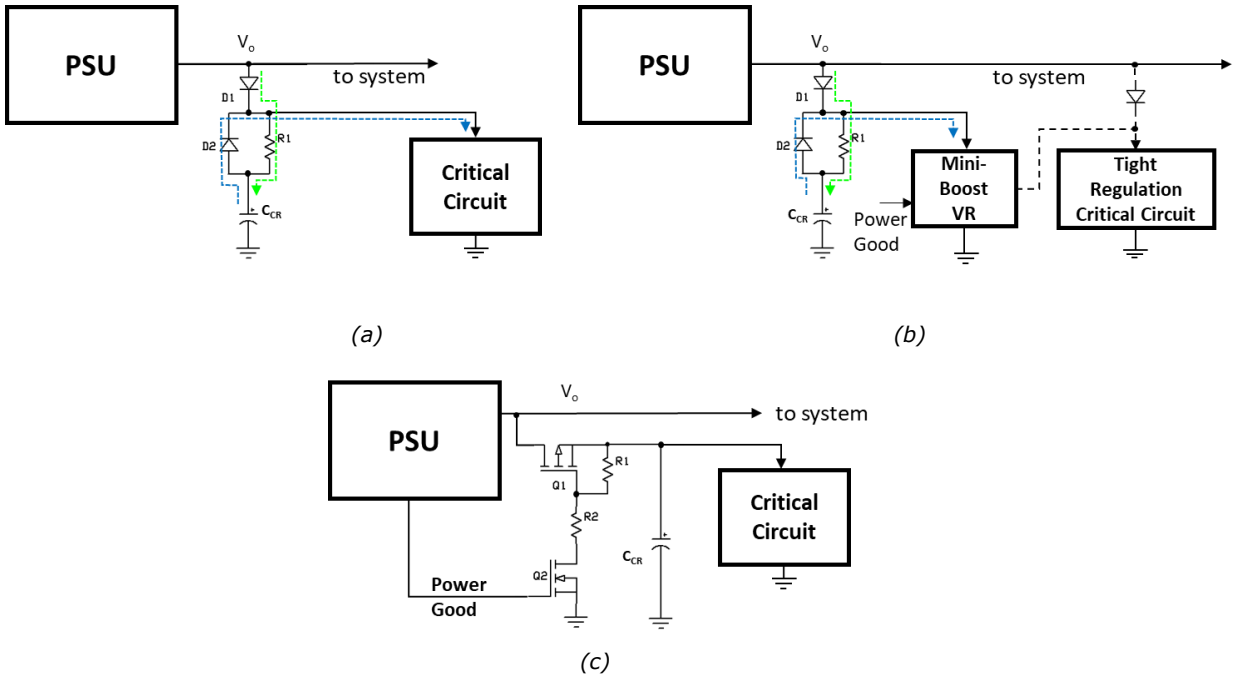

Fig. 5. Buffering voltage source selector arrangements. A passive diode selector network can be used for small buffered power cases (a). A mini boost regulator that tolerates larger voltage swings on its input can be used for critical circuitry that requires tightly regulated supply voltage (b). And for higher buffered-power levels, an active network (c) minimizes power dissipation in normal operating mode.

Simulation timing diagrams illustrating the critical circuit buffering function in Fig. 5c are shown in Fig. 6a and b. In these diagrams the yellow waveforms represent power supply output voltage. They also correspond to the supply voltage waveforms in the conventional case with the original warning time. The blue waveforms represent critical circuit supply voltages with secondary-side capacitor buffering.

In Fig. 6a the warning time is extended by 12+ms when a standard 5% secondary voltage droop is specified. If the allowed voltage droop is increased to 33% (Fig. 6b), a similar warning time extension can be achieved with a much smaller buffer cap. The simulation in both cases was conducted for 12-V dc output in constant power mode.

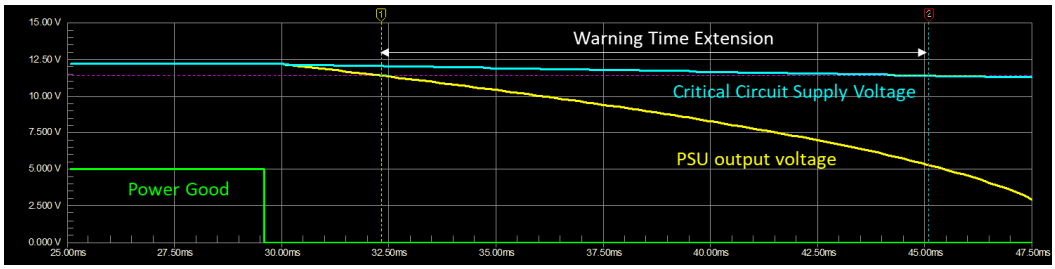

(a)

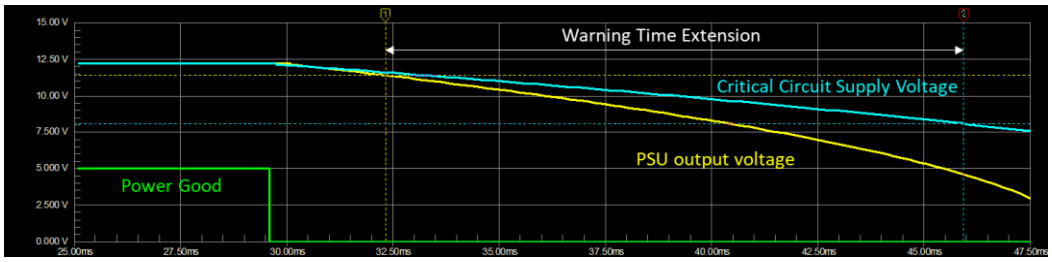

(b)

Fig. 6. Simulation waveforms in constant power mode for the buffering circuit arrangement shown in Fig. 5c. Isolating a critical circuit that consumes a small portion of system power and using a local energy storage cap provides a 12-ms warning time extension while using the original PSU and sticking with a 5% secondary voltage droop (a). But if the specified droop for the critical circuit supply voltage can be increased from 5% to 33%, a similar warning time extension can be achieved with a smaller local energy storage cap (b).

This analysis demonstrates that when the warning time needs to be drastically extended, separating the critical circuit power delivery path from the rest of the power delivery network and using local energy buffering can provide cost and size reduction for the energy storage component. Furthermore, this setup can avoid the need for power supply redesign when the warning time needs to be significantly extended.

In the PSU hard failure mode, including its output short circuit, with the traditional power delivery architecture, system power hold-up and warning time requirements cannot be met. In the secondary-side buffering case, such as the buffering voltage source selector arrangements shown in Fig. 5, the critical circuitry power hold-up/ warning time requirements are supported. This provides much better system immunity to a significantly broader variety of power fault conditions.

The proposed method calls for optimizing the circuitry incorporating components critical to organized system shutdown, for expanding the range of supplied voltages required for this circuit's normal operation, and for minimizing the portion of system power consumed by it. Under these conditions, the method discussed here becomes the most effective.

Conclusions

In many cases separating the critical-components power delivery path from the rest of the system power distribution network and using a local energy storage component represents a better option for extending the fault warning time than increasing the size of the PSU bulk cap. Using local energy storage on the system side allows us to significantly extend critical circuitry hold-up and/or warning time while using conventional power supplies with standard hold-up time specifications and avoids PSU redesign.

In this article we have shown that requirements for buffered power share and allowed secondary voltage swing influence requirements for secondary-side versus primary-side energy storage. The smaller the buffered power share the more advantageous the use of secondary-side energy storage. Meanwhile, the larger the allowed droop in the secondary voltage, the larger the buffered power share that can be supported with a given secondary-side energy storage component.

These relationships have been quantified in this article. By using equation 5 or the graphs in Fig. 3, designers can determine whether a given portion of buffered power and allowed droop provide size and cost advantages for secondary-side energy storage over the conventional case of primary-side energy storage.

Finally, when weighing these options there's another consideration. Local secondary-side energy buffering is capable of providing sufficient warning time for safe data transfer and facilitates an organized system shutdown under a much broader variety of power fault conditions including indiscriminate PSU fault cases.

References

1. Functions / Requirements of Direct-Off-Line SMPS -- POWER FAILURE WARNING CIRCUITS.
2. "Surviving SSD sudden power loss" by Zsolt Kerekes.
3. "Closed Loop System Throttling (CLST): A New Power Technology for Improving Server Efficiency" by Brian Griffith, Viktor Vogman and Andrew Watts, IDF 2010, TMTS001
4. "Hold-Up time extension circuit with integrated magnetics" by Yungtaek Jang, Milan M. Jovanovic, and Dave L. Dillman, Proceedings of Applied Power Electronics Conference and Exposition (APEC 2005), pp. 219-225.
5. "High-Voltage Energy Storage: The Key to Efficient Holdup" by Jean Picard.
6. "Power-loss protection" by Byron A Alcorn et al., US Patent 10,198,320, February 5, 2019.

About The Author

Viktor Vogman currently works at Power Conversion Consulting as an analog design engineer, specializing in the design of various power test tools for ac and dc power delivery applications. Prior to this, he spent over 20 years at Intel, focused on hardware engineering and power delivery architectures. Viktor obtained an MS degree in Radio Communication, Television and Multimedia Technology and a PhD in Power Electronics from the Saint Petersburg University of Telecommunications, Russia. Vogman holds over 50 U.S. and foreign patents and has authored over 20 articles on various aspects of power delivery and analog design.

For articles relating to powering servers, see How2Power's Design Guide, locate the "Application" category and select "Data Centers". For more on power backup, see the "Power Supply Function" category and select "UPSs".