

# Overcoming Rack Power Limits with VPS Dynamic Redundancy and Intel® RSD

## Summary

This paper describes how SourceMix, a dynamic redundancy technology from VPS, allows Intel® Rack Scale Design (Intel® RSD) customers to take full advantage of system composability and module upgradeability by extending the existing data center power infrastructure.

## 1. Introduction to Intel® RSD

Intel® RSD is a new data center architecture that physically separates or “disaggregates” compute, storage and network resources into unallocated groups called pools, which can be efficiently assembled or “composed” on demand to create a precise hardware configuration to match a specific software workload requirement. In addition, Intel® RSD provides a comprehensive management architecture for all resources in the data center, enabling a true software defined infrastructure based on open APIs. Intel® RSD concepts of disaggregation, pooling and composition are shown in Figure 1.

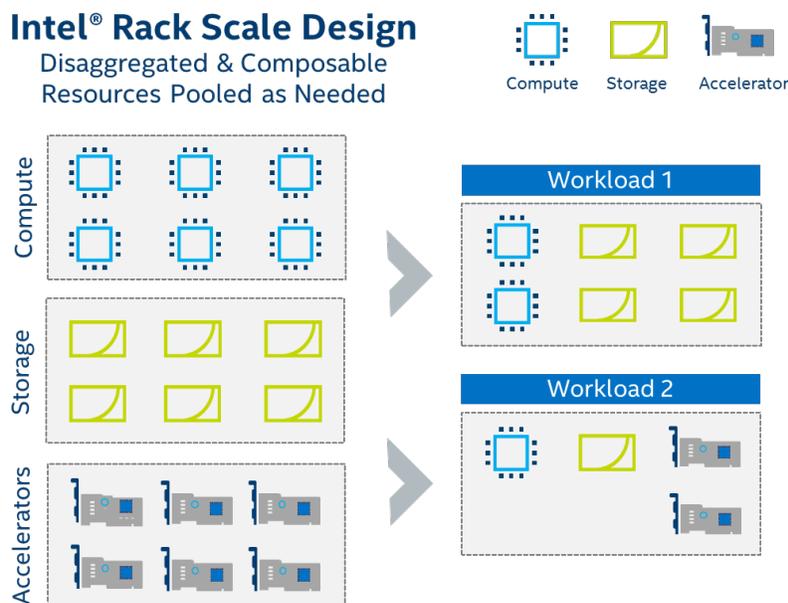


Figure 1 Intel® RSD Overview

Intel® RSD supports dynamic composition of platforms for specific workloads, and the ability to easily upgrade modules to gain additional capacity. But the flexibility to compose and re-compose platforms can be limited by the data center power infrastructure if a new configuration exceeds the power available in a rack. A more dynamic power infrastructure complements and extends the flexibility offered by Intel® RSD.

## 2. Dynamic Redundancy from VPS

Virtual Power Systems, a Silicon Valley-based startup, has developed the Software-Defined Power® (SDP) product suite to unlock capacity and remove power constraints by implementing the principles of disaggregation and composition within the power infrastructure. Using SDP, data center suppliers can provide additional flexibility in their Intel® RSD solutions. End-customers can deploy SDP on their own in their existing infrastructure.

A key component of SDP is called SourceMix, which addresses the need for *Dynamic Redundancy*. This technology complements the static power infrastructure in a typical data center by providing intelligent, software-defined power redundancy through optimized control of power flow switches. SDP's in-rack switches allow customers using dual-corded power distribution to dynamically assign power redundancy while staying within the overall capacity of the power infrastructure. SourceMix dynamic power redundancy allocation can optimize the mix of single and dual corded power distribution to unlock unused power capacity as workload requirements change.

Figure 2 shows an intelligent SDP power switch at the top of a rack, which can route input power from both incoming feeds to power distribution units (PDUs) in the rack. Each module (compute, storage, network, etc.) in the rack is connected to two PDUs (one for cord A and one for chord B). Power redundancy can then be controlled dynamically under software control via the power flow switch. The switch uses software-defined policies to control redundancy based on the priorities of the workload in the rack, and predictive algorithms based on historical load patterns in the data center.

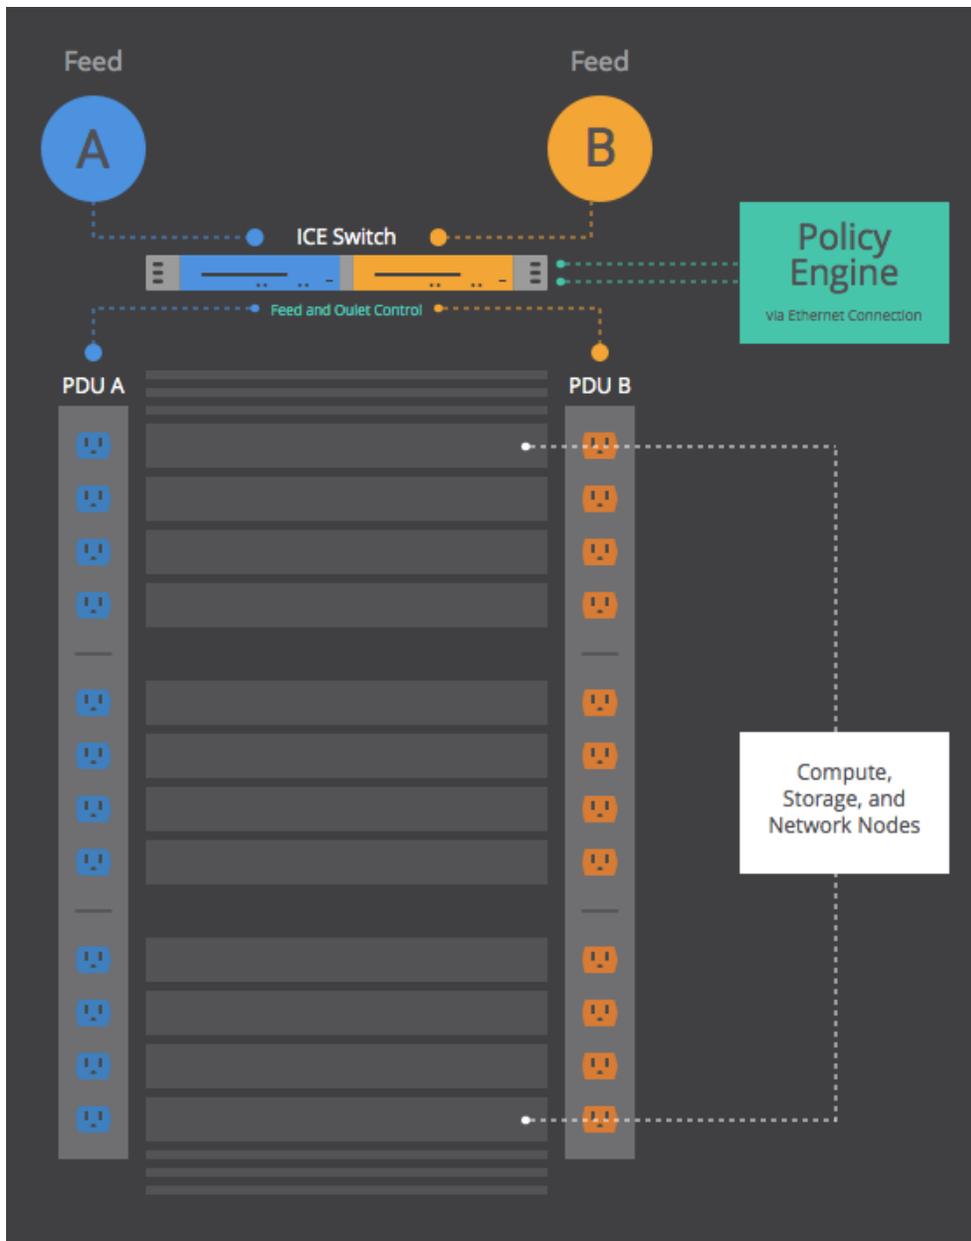


Figure 2. Power switch at the top of a rack controlling redundant power feeds.

At any point in time, the switch can decide to power a module (sled, chassis) using both PDUs, a single PDU or none at all. This determines the power redundancy—2N or 1N—for each module. When operating in 2N mode, a feed failure will automatically cause power to be drawn from the other PDU. In 1N mode, a feed failure will cause the switch to make an instantaneous decision based on the current power situation and workload priority to either keep a module on the current feed, switch to the alternate feed, or turn the module off completely to operate within available power limits. Switches can also be configured to force a 1N module to prioritize power draw from a particular feed.

Through software-driven control of these switches, VPS SourceMix dynamic redundancy provides two distinct benefits:

1. **Unlocking Redundant Power:** Colocation providers and enterprise data centers can unlock additional power from their redundant infrastructure for active ongoing use.
2. **On-demand Redundancy:** Intelligent provisioning systems like Intel® RSD POD Manager can dynamically control the redundancy of their compute and storage nodes to ensure high-availability of critical systems as configurations change.

### 3. Unlocking Redundant Power

Data centers provide power redundancy to make workloads fault-tolerant to utility and power equipment failures. They create alternate power paths that can take on additional load in the event of failures. Figure 3 illustrates a representative 2 Megawatt (MW) data center with two fully redundant power paths from the substation (or from two different utilities). Its high-priority workloads are connected to both power paths in a 2N configuration. If one feed were to fail, the other feed is fully capable of handling the entire workload.

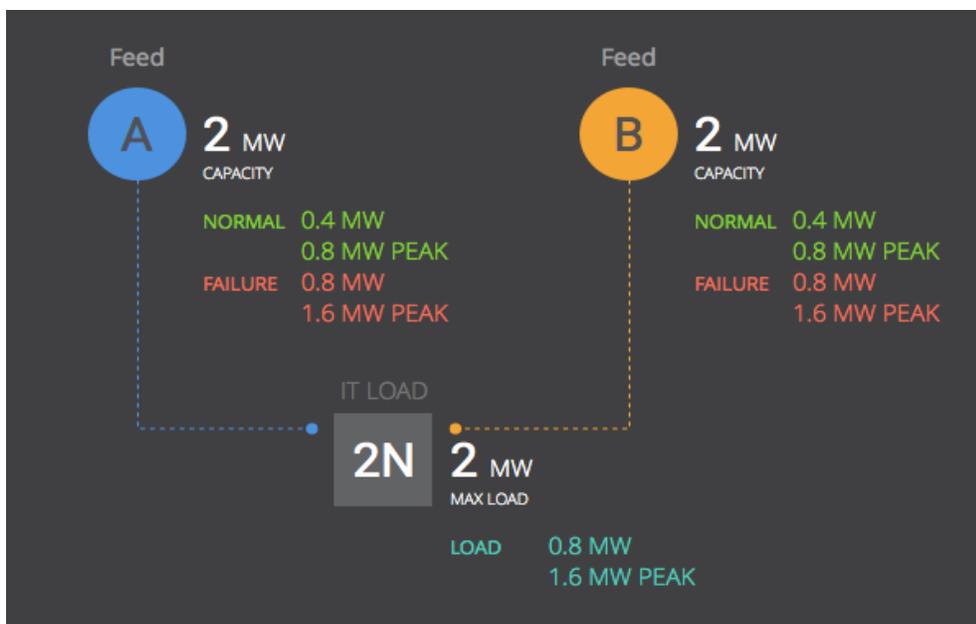


Figure 3. Redundant power paths to 2N workloads in a data center.

In this example, the power system is designed for a 2 MW maximum load, but most data centers operate at a much lower utilization—often in the range of 20% to 40%. For our illustrative data center, we assume a 40% normal utilization with peaks going up to 80% about 10% of the time—a very conservative estimate. Figure 3 shows power utilization during “Normal” operation (both feeds supplying power normally) and “Failure” operation (when one of the power feeds is not operating). Each feed contributes only 0.4MW for 90% of the time, and 0.8 MW for roughly 10% of the time. In the rare occurrence of a failure of one of the feeds

during peak load, the other feed would need to provide 1.6MW, but this is expected less than 0.01% of the time. So 99.9% of the time, total power capacity is significantly underutilized.

This presents an opportunity for colocation providers and enterprise data centers to use some of this overprovisioned capacity for lower priority workloads (we call them “less than 2N” or “<2N”). Figure 4 shows the power that can be unlocked in our example data center. There’s a potential to add almost 2.4MW of workload 99.9% of the time! However, if you were to add more workload in the existing power architecture as is, you run the risk of overloading either feed in case of a failure. This would have the disastrous effect of bringing the whole data center down, including the 2N workloads.

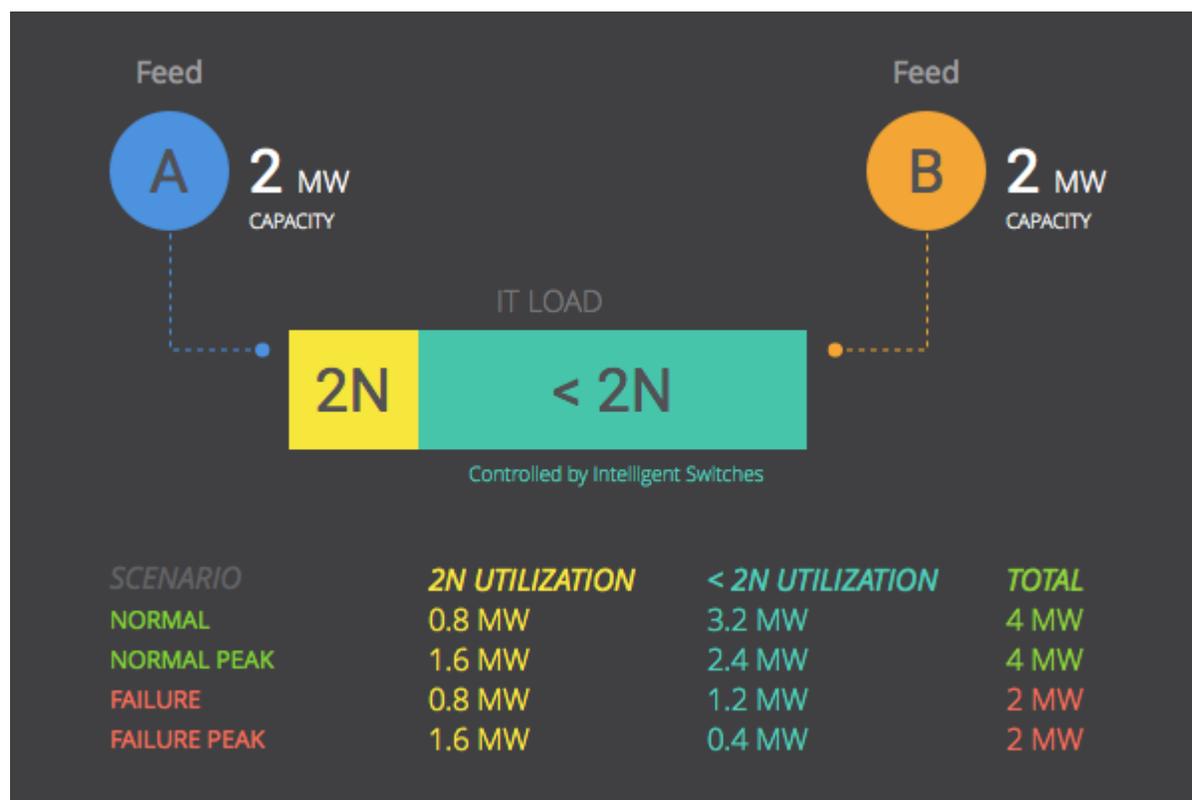


Figure 4: Capacity unlocked by dynamic redundancy.

VPS SourceMix can unlock power capacity while mitigating the risk of bringing critical workloads down due to a power failure. SourceMix uses information on power consumed by power switches across the entire data center, allowing the switches to coordinate with each other. The system reacts quickly when there is failure by cutting power to non-critical workloads and leaving critical loads fully operational within the limits of the available feeds.

To ensure fast, correct and coordinated reaction from the switches, VPS’s patent-pending algorithms analyze data center-wide power consumption and compress this information into a few bits of data, which are then distributed to all the intelligent switches via a managed network. This information is updated every 10 to 20 seconds, and takes into account overall trends to

anticipate critical situations. The exact information sent to each switch is different based on the nature of the power problem, where the switch is located in the data center topology and the priority of the loads it controls. This allows SourceMix to deliver millisecond reaction times at the switches to maintain critical workloads without exceeding maximum power constraints.

## 4. On-Demand Redundancy

VPS SourceMix also allows third-party provisioning and orchestration systems to dynamically designate the power redundancy needed at a particular rack or node. One designated, the SourceMix algorithms re-compute policies for the power switches to ensure that the requested redundancy is honored. These designations can be changed as often as needed.

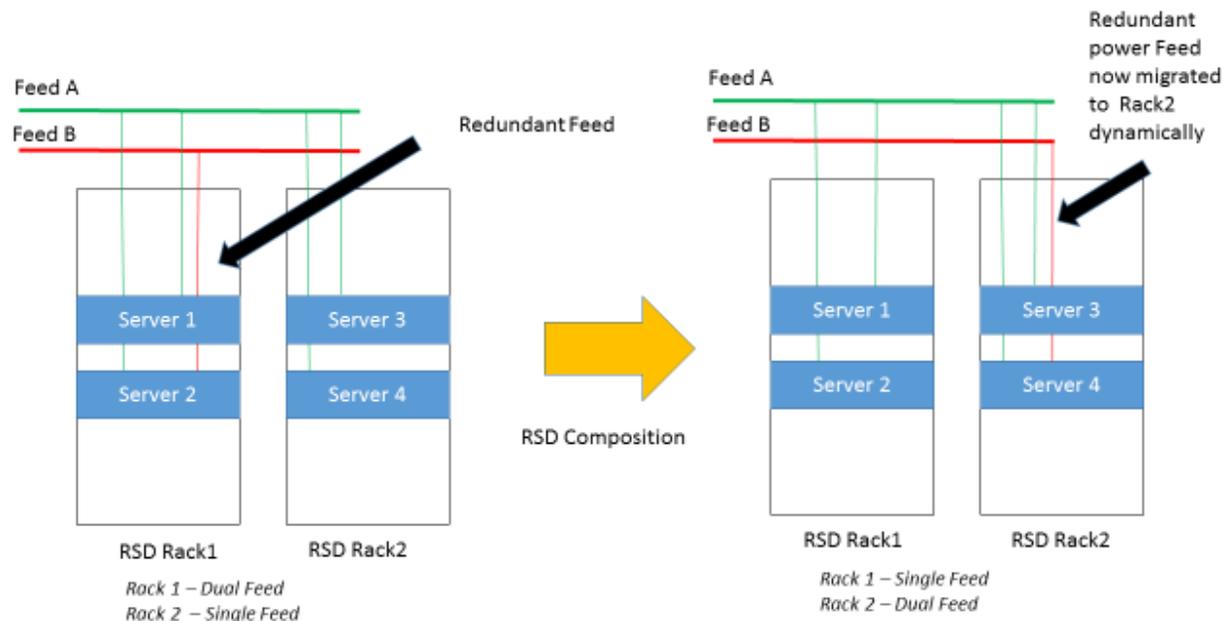


Figure 5. Intel® RSD Composable power.

This complements the Intel® RSD design philosophy because modules don't need to be permanently designated as 2N by virtue of physical cabling. Redundancy can be allocated dynamically when a system is composed. If a 2N compute or storage module fails, a new system can be composed and the workload migrated to it. VPS SourceMix allows the power redundancy configuration to also move to the newly composed system, maintaining 2N power provisioning for the critical workload.

By integrating SourceMix into the Intel® RSD infrastructure, the on-demand power redundancy capabilities of VPS can be used by orchestration software to deliver dynamic redundancy across racks or within nodes in the same rack as requirements change.

## 5. Handling Cooling and Other Realities

When new power consumption is enabled by VPS, it also affects the cooling load. In our example, data center the cooling plant might be sized to a maximum of 2MW. What happens when we use SourceMix to enable more power consumption?

In most data centers, just as the actual power utilization is often significantly below the provisioned capacity, the cooling capacity is also underutilized to a similar extent. So additional power consumption enabled by SourceMix is within the maximum cooling load limit a vast majority of the time. VPS SourceMix gathers power consumption information at sub-second intervals to fuel its decisions, so it can react quickly to any normal or abnormal increases in utilization, keeping total power consumption within the limits of cooling and other physical constraints.

## 6. Adaptive Intelligence

The VPS solution incorporates learning algorithms that analyze and predict power consumption to improve capacity utilization over time. Multiple levels of monitoring, control and optimization intelligence, as shown in Figure 5, ensure that short-term and long-term power consumption trends are taken into account in managing dynamic redundancy.

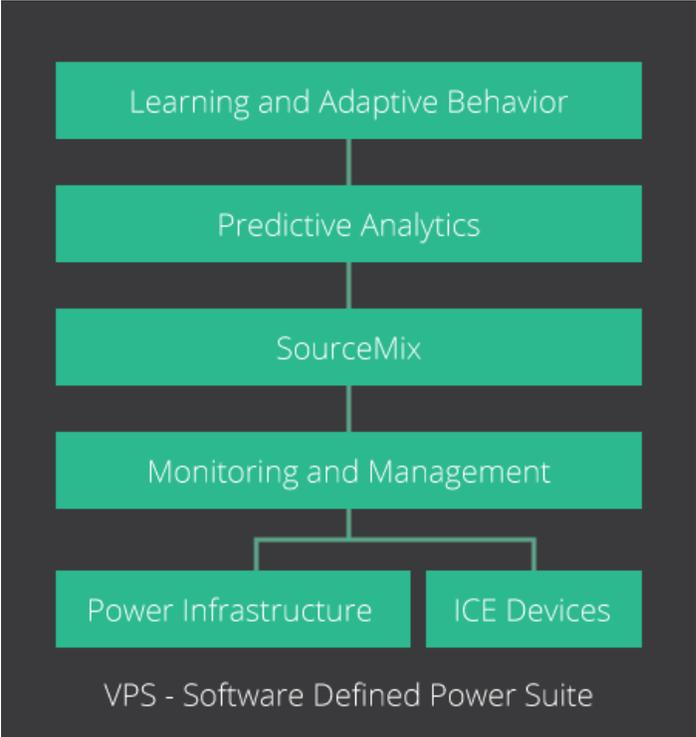


Figure 5. SDP Hierarchy of Optimization and Intelligence

VPS's solutions are available for commercial use via their SDP suite, and are ready for integration with other data center management systems via well-defined APIs. Implementing an

SDP solution involves installing the hardware anywhere in a rack, and installing the software on any server that is network accessible, enabling fast turn-key deployments.

## 7. Conclusions

Integration with VPS SourceMix technology increases the effectiveness of Intel® RSD implementations by enabling dynamic control of power redundancy by the Intel® RSD Pod Manager to ensure high availability of critical workloads in case of power or power equipment failure. SourceMix can also unlock underutilized power capacity by optimizing the use of power redundancy in the data center.